

Ivo Martinac · Bo Nørregaard Jørgensen ·
Zheng Grace Ma · Rúnar Unnþórsson ·
Chiara Bordin (Eds.)

LNCS 16095

Energy Informatics

First Nordic Energy Informatics Academy
Conference, EIA Nordic 2025,
Stockholm, Sweden, August 20–22, 2025
Proceedings, Part I

1
Part I



Springer

Founding Editors

Gerhard Goos

Juris Hartmanis

Editorial Board Members

Elisa Bertino, *Purdue University, West Lafayette, IN, USA*

Wen Gao, *Peking University, Beijing, China*

Bernhard Steffen , *TU Dortmund University, Dortmund, Germany*

Moti Yung , *Columbia University, New York, NY, USA*

The series Lecture Notes in Computer Science (LNCS), including its subseries Lecture Notes in Artificial Intelligence (LNAI) and Lecture Notes in Bioinformatics (LNBI), has established itself as a medium for the publication of new developments in computer science and information technology research, teaching, and education.

LNCS enjoys close cooperation with the computer science R & D community, the series counts many renowned academics among its volume editors and paper authors, and collaborates with prestigious societies. Its mission is to serve this international community by providing an invaluable service, mainly focused on the publication of conference and workshop proceedings and postproceedings. LNCS commenced publication in 1973.

Ivo Martinac · Bo Nørregaard Jørgensen ·
Zheng Grace Ma · Rúnar Unnþórsson ·
Chiara Bordin
Editors

Energy Informatics

First Nordic Energy Informatics Academy
Conference, EIA Nordic 2025
Stockholm, Sweden, August 20–22, 2025
Proceedings, Part I

Editors

Ivo Martinac 
KTH Royal Institute of Technology
Stockholm, Sweden

Zheng Grace Ma 
University of Southern Denmark
Odense, Denmark

Chiara Bordin 
Arctic University of Norway
Tromsø, Norway

Bo Nørregaard Jørgensen 
University of Southern Denmark
Odense, Denmark

Rúnar Unnþórsson 
University of Iceland
Reykjavík, Iceland

ISSN 0302-9743

Lecture Notes in Computer Science

ISBN 978-3-032-03100-6

<https://doi.org/10.1007/978-3-032-03101-3>

ISSN 1611-3349 (electronic)

ISBN 978-3-032-03101-3 (eBook)

© The Editor(s) (if applicable) and The Author(s), under exclusive license to Springer Nature Switzerland AG 2026

Chapters “Towards ICT-Enabled Multi-Agent based Operations in Local Energy Communities: A Proof of Concept” and “Decentralized Reinforcement Learning for Adaptive Power Sharing in Hybrid DC Microgrids” are licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>). For further details see license information in the chapters.

This work is subject to copyright. All rights are solely and exclusively licensed by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, expressed or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

If disposing of this product, please recycle the paper.

Preface

The Nordic Energy Informatics Academy Conference 2025 (EIA Nordic 2025) brought together researchers, innovators, and professionals across Europe and beyond to explore how digitalization and informatics can accelerate the sustainable transformation of the energy sector. Held over three days, the conference served as a vibrant forum for knowledge exchange, collaboration, and innovation at the intersection of energy systems and digital technologies.

Thanks to the dedicated efforts of the organizing and technical program committees, a total of sixty-five (65) submissions were received for EIA Nordic 2025. Each submission was evaluated through a rigorous single-blind peer-review process, with three expert reviewers assigned to each paper. Reviews were based on relevance, scientific quality, originality, and clarity. According to the review policy, a manuscript was accepted only if it received at least two positive recommendations.

As a result, fifty-one (51) high-quality research papers were accepted and presented across eleven thematic sections, including 43 full papers and 8 short papers. The strong alignment of submissions with the conference's specific scope, as well as the high standard of contributions from experienced research groups and ongoing collaborations within the Nordic energy informatics community, contributed to this acceptance outcome. All accepted papers are published in the Springer Lecture Notes in Computer Science (LNCS) series, ensuring global visibility and scientific impact.

The conference opened with keynote speeches that set the tone for deep discussion on emerging energy informatics challenges and opportunities. The technical program included insightful presentations organized around the following focused themes:

- Energy Forecasting and Intelligent Control Systems
- District Heating, Thermal Systems, and Retrofit Strategies
- Building Simulation, Urban Energy, and Environmental Sensing
- Industrial Process Efficiency and Biomass Utilization
- Energy Informatics for Electric Vehicles and Mobility Systems
- Multi-Agent Systems and Local Market Coordination
- Policy, Metrics, and Infrastructure Performance
- Smart Building Systems and Semantic Data Integration
- Prosumer Optimization and Energy Storage in Local Energy Communities
- Grid-Oriented AI, Simulation, and Resilience
- Non-Intrusive Load Monitoring and Data Competitions

Each session provided a platform to present cutting-edge methods, innovative applications, and empirical case studies that address real-world energy challenges—from grid flexibility and urban energy sensing to semantic data modeling, battery storage, and AI-based resilience strategies.

In addition to the paper sessions, the conference featured a technical site visit to the KTH Live-in Lab, offering participants further opportunities to network and experience real-world energy innovation in action.

EIA Nordic 2025 once again demonstrated the value of interdisciplinary research and the power of digital tools to drive impactful solutions in energy systems. The conference fostered new ideas, strengthened collaborations, and inspired the community toward the continued advancement of energy informatics.

July 2025

Ivo Martinac
Bo Nørregaard Jørgensen
Zheng Grace Ma
Rúnar Unnþórsson
Chiara Bordin
Zheng Grace Ma
Signe Rude Madsen

Organization

General Chairs

Ivo Martinac
Bo Nørregaard Jørgensen

KTH Royal Institute of Technology, Sweden
University of Southern Denmark, Denmark

Program Chairs

Zheng Grace Ma
Rúnar Unnþórsson
Chiara Bordin

University of Southern Denmark, Denmark
University of Iceland, Iceland
Arctic University of Norway, Norway

Publication Chair

Zheng Grace Ma

University of Southern Denmark, Denmark

Operations Officer

Signe Rude Madsen

University of Southern Denmark, Denmark

Program Committee

Abbas Fotouhi
Abdulsalam Yassine
Agus Hasan
Ahmed Yousif Mohamed Idries
Aikaterini Chatzivasileiad
Alexander Kies
Alireza Afshari
Amjad Anvari-Moghaddam
Aoife M. Foley
Arianna Fonsati

Cranfield University, UK
Lakehead University, Canada
Norwegian University of Science and Technology,
Norway
Norwegian University of Science and Technology,
Norway
Cardiff University, UK
Aarhus University, Denmark
Aalborg University, Denmark
Aalborg University, Denmark
University of Manchester, UK
KTH Royal Institute of Technology, Sweden

Bjørn Austbø	Norwegian University of Science and Technology, Norway
Carlos Fernandez	Robert Gordon University, UK
Cathryn Peoples	Ulster University, UK
Daniel Donaldson	University of Birmingham, UK
Di Wu	Norwegian University of Science and Technology, Norway
Eduard Petlenkov	Tallinn University of Technology, Estonia
Etienne Saloux	Natural Resources Canada, Canada
Hamid Reza Shaker	University of Southern Denmark, Denmark
Hans-Peter Schwefel	Aalborg University, Denmark
Hassam ur Rehman	VTT Technical Research Centre of Finland, Finland
Hicham Johra	SINTEF, Norway
Himanshu Nagpal	EURAC Research, Italy
Hossam Gaber	Ontario Tech University, Canada
Hung Cao	University of New Brunswick, Canada
Igor Sartori	SINTEF, Norway
Ivan Korolija	University College London, UK
Jarek Kurnitski	Tallinn University of Technology, Estonia
Jianguo Ding	Blekinge Institute of Technology, Sweden
Joy Dalmacio Billanes	University of Southern Denmark, Denmark
Joyjit Chatterjee	University of Hull, UK
Juri Belikov	Tallinn University of Technology, Estonia
Kai Heussen	Technical University of Denmark, Denmark
Katja Sirviö	VTT Technical Research Centre of Finland, Finland
Kristoffer Christensen	University of Southern Denmark, Denmark
Kun Zhang	Université du Québec, Canada
Laurent Georges	Norwegian University of Science and Technology, Norway
Linas Martišauskas	Lithuanian Energy Institute, Lithuania
Louis Gosselin	Université Laval, Canada
Magnus Værbæk	University of Southern Denmark, Denmark
Mahshid Javidsharifi	Aalborg University, Denmark
Majid Ali	Aalborg University, Denmark
Manuel Llorca	Copenhagen Business School, Denmark
Mohamed Emad Farrag	Glasgow Caledonian University, UK
Najmeh Bazmohammadi	Aalborg University, Denmark
Niklas Lavesson	Blekinge Institute of Technology, Sweden
Ólafur Pétur Pálsson	University of Iceland, Iceland
Paula Carroll	University College Dublin, Ireland

Pei Huang	Dalarna University, Sweden
Qian Wang	KTH Royal Institute of Technology, Sweden
Ravi Pandit	Cranfield University, UK
Risto Kosonen	Aalto University, Finland
Saleh Abdel-Afou Alaliyat	Norwegian University of Science and Technology, Norway
Sebastian Büttrich	IT University of Copenhagen, Denmark
Seyed Shahabaldin Tohidi	Technical University of Denmark, Denmark
Stefán Thor Smith	University of Reading, UK
Touraj Ashrafian	Northumbria University, UK
Umit Cali	University of York, UK
Vahid Nik	Lund University, Sweden
Weiqi Hua	University of Birmingham, UK
Yasunori Akashi	University of Tokyo, Japan
Yushuai Li	Aalborg University, Denmark
Zeeshan Afzal	Linköping University, Sweden

Contents – Part I

Energy Forecasting and Intelligent Control Systems

A Multi-stage Deep Learning Framework for Short-Term Electricity Load Forecasting 3
Ege Kandemir, Agus Hasan, and Saleh Abdel-Afou Alaliyat

Solaris AI: Enhancing Solar Energy Forecasting with Generative AI and Deep Learning 15
Mohammed Farhan Faisal, Nimisha Nixon, and Pamba Raja Varma

A Tool for Synthesizing and Implementing Medium Voltage Load Profiles 32
*Hendrik Plompen, Ranier Alexander Arruda Moura,
Khawaja Khalid Mehmood, Anne van der Molen, Peter van der Wielen,
and Phuong Hong Nguyen*

Decentralized Reinforcement Learning for Adaptive Power Sharing in Hybrid DC Microgrids 60
Abd Alelah Derbas, Chiara Bordin, Sambeet Mishra, and Frede Blaabjerg

District Heating, Thermal Systems, and Retrofit Strategies

Digitalization in District Heating: Comparative Insights from Denmark and Sweden on Adoption, Barriers, and Value Creation 79
Zheng Ma, Joy Dalmacio Billanes, and Kristina Lygnerud

Evaluating Retrofit Strategies and Decentralized Systems for the Transition to Low-Temperature District Heating: A Simulation-Based Case Study in Borlänge, Sweden 95
Vignesh Pechiappan Ayyathurai and Abdelmomen Najmadin

Parametric Study Model for Observing Exergy Balance Through the Series of Subsystems in District Heating 110
Genku Kayo

Demand-Side Frequency Response Based on District Heating System Integrated with Heat Pump 119
Hui Yan and Sara Walker

A Parametric Optimization Approach for Enviro-Economic Evaluation of Energy Renovation Strategies – A Case Study on a Congregation House in Southern Sweden	131
---	-----

*Md Parvaz, Einar Örn Þorvaldsson, Jesper Engström,
Dennis Johansson, and Henrik Davidsson*

Building Simulation, Urban Energy, and Environmental Sensing

Development of an Automated 3D Building Modeling Method and Urban-Scale Analysis of Heating and Cooling Loads	149
---	-----

Hisato Osawa, Taro Mori, and Konatsu Suzuki

Research on 3D Modeling for Thermal Environment Simulation Using Publicly Available Map Data	164
--	-----

Kazuma Otani, Taro Mori, Hisato Osawa, and Konatsu Suzuki

Rooftop Irregularity Segmentation in Aerial Imagery Using Deep Learning	173
---	-----

Konatsu Suzuki, Taro Mori, Hisato Osawa, and Kazuma Otani

Autonomous Air Quality Monitoring System with Photovoltaic Energy Harvesting: A Sustainable Approach for Public Policies	185
--	-----

Emerson Santana and Eduardo Liberado

How Clean is the Air You Breathe During Urban Walk? A Case Study of Central London	198
--	-----

Nikhil Ravindra, Amin Al-Habaibeh, and Benachir Medjdoub

Industrial Process Efficiency and Biomass Utilization

Examining the Role of Digital Technologies and Artificial Intelligence in Climate Resilience and Energy Adaptation Within Energy-Intensive Industries	215
---	-----

Joy Dalmacio Billanes, Bo Nørregaard Jørgensen, and Zheng Grace Ma

Energy Efficiency Optimization in Plastic Pyrolysis: A Data-Driven Modeling Study	230
---	-----

Aysan Safavi, Christiaan Richter, and Runar Unnþorsson

Globally Optimal Scheduling for Industrial Energy Cost Reduction Under Dynamic Electricity Pricing	237
--	-----

Lu Cong, Bo Nørregaard Jørgensen, and Zheng Grace Ma

Energy Informatics for Electric Vehicles and Mobility Systems

The Irish Highway Network: A Novel Test Instance for the Charging Station Location Problem 255
Jingyu Xiang, Paula Carroll, and Annunziata Esposito Amideo

Enhancing EVRP Benchmark Instances with Energy Estimates 267
Clíodhna Ní Shé and Paula Carroll

Economic and Environmental Benefits of Centralized MILP Optimization of EV Fleet Charging 283
Lucija Hajsok and Tea Žakula

Electric Vehicle Based Virtual Electricity Network (EVEN) Solution for Performance Enhancement in Distribution Networks 299
Pei Huang and Rehman Zafar

Variational Quantum Eigensolver-Based CaaS Business Model for V2G 309
Desh Deepak Sharma, Ramesh C. Bansal, and Jeremy Lin

Multi-Agent Systems and Local Market Coordination

Towards ICT-Enabled Multi-agent Based Operations in Local Energy Communities: A Proof of Concept 323
Haoyu Huang, Natascha Fernengel, André Xhonneux, Alexander Holtwerth, Michael Hehemann, Eugen Hoppe, Simon Waczowicz, Kevin Förderer, Veit Hagenmeyer, and Dirk Müller

Peer-to-Peer Energy Management Model for Residential Homes 341
Najmeh Khajoei, Runar Unnþorsson, and Steinn Gudmundsson

Agent-Based Flexibility Aggregation for a Distributed Redispatch 354
Malin Radtke, Sanja Stark, and Stefanie Holly

A Visualization Framework for Exploring Multi-agent-Based Simulations: Case Study of an Electric Vehicle Home Charging Ecosystem 371
Kristoffer Christensen, Bo Nørregaard Jørgensen, and Zheng Grace Ma

Author Index 387

Contents – Part II

Policy, Metrics, and Infrastructure Performance

Proper Definitions of Micro Grid Metrics are Needed! – A Generalizable Framework	3
<i>Jonathan Fellerer, Jana Helbrecht, and Reinhard German</i>	
Solar-Geothermal Power “HGS-ORC” System for Energy Co-generation: Energy, Economic and Environmental Analysis: Algerian Case	24
<i>Mohamed El-Amine Slimani and Maria Sigríður Guðjónsdóttir</i>	
Towards the Integration of Data Space Technology in Hydrogen Research Workflows	36
<i>Ryan Ford, Rita Beigaité, and Lius Daniel</i>	
Comparison of Outages Trends and Statistics in Nordic Countries Across Distribution Networks and Their Impacts	51
<i>Seema, Andreas Theocharis, and Reza Sirjani</i>	
Managing Risk in Distribution Systems with Solar Generation: A Case Study Using the MATPOWER Optimal Scheduling Tool	67
<i>Johanna Bolaños-Zuñiga and Alberto J. Lamadrid L.</i>	
Smart Energy Management System with Individual Load Monitoring	78
<i>A. Alex, Jayarama Pradeep, T. S. Aiswarya, N. Harini, and Nikith Jude Serrao</i>	

Smart Building Systems and Semantic Data Integration

Towards a Taxonomy for Application of Machine Learning and Artificial Intelligence in Building and District Energy Management Systems	95
<i>Klaus Lichtenegger, Florian Ahammer, Fabian Schopper, Daniel Muschick, and Markus Gölles</i>	
A Dynamic Semantic Data Modeling Approach: Application to Flexible HVAC Zones	114
<i>Shoya Marumoto, Shohei Miyata, Keiichiro Taniguchi, and Yasunori Akashi</i>	

Leveraging Generative AI and Semantic Data for Improved Operation of a Real-Life Building	130
<i>John Clauß, Luis Caetano, Knut Nordanger, Thomas Elvrum Lassen, and Reidar Kind</i>	
Development of an LSTM-Based Model for High-Resolution Downsampling and Reconstruction of HVAC Chiller Flow Data	144
<i>Yue He, Shanrui Shi, Shohei Miyata, and Yasunori Akashi</i>	
Data-Driven Optimal Air-Balancing Control for Multizone Ventilation Systems with Design-to-Operation Adaptation	154
<i>Shanrui Shi, Shohei Miyata, and Yasunori Akashi</i>	
Prosumer Optimization and Energy Storage in Local Energy Communities	
Evaluating the Potential for Developing Local Energy Communities in Sweden: Case Studies at Jättesten and Chalmers Campus	173
<i>Mohammadreza Mazidi, Araavind Sridhar, David Steen, Elena Malakhatka, Sara Abouebeid, Felix Niklasson, Le Anh Tuan, and Holger Wallbaum</i>	
Data-Driven Correlated Uncertainty Sets for PV Generation and Electricity Demand	187
<i>Debjayoti Biswas, Cristian Aguayo, Anna Mutule, and Paula Carroll</i>	
Scheduling Heat Pumps for Balancing Thermal Storage and Grid Export	204
<i>Aditya Somawanshi and Anupama Kowli</i>	
Battery Energy Storage Integration with BIPV Systems: A Multi-scenario Economic Analysis and Optimization	221
<i>Hashem Amini Toosi</i>	
Grid-Oriented AI, Simulation, and Resilience	
A Data-Driven Analysis of Unscheduled Flows in the European Power System	237
<i>Maurizio Titz and Dirk Witthaut</i>	
Green Hydrogen Under Uncertainty: Evaluating Power-to-X Strategies Using Agent-Based Simulation and Multi-criteria Decision Framework	254
<i>Frederik Wagner Madsen, Joy Dalmacio Billanes, Bo Nørregaard Jørgensen, and Zheng Ma</i>	

Synthesizing Fault Localization Datasets	271
<i>Zhonghe Chen, Adi Botea, Paula Carroll, and Deepak Ajwani</i>	
Machine Learning-Based Cyberattack Detection in Power Data	285
<i>Robert A. Becker, Nikolai Kamenev, Celina Koelsch, Aashay Kulkarni, and Thomas Bleistein</i>	
Optimization of Second-Life Battery Energy Storage System in Buildings with Photovoltaic Panels: A Norwegian Case Study	300
<i>Italo Aldo Campodonico-Avendano, Amin Moazami, Aileen Yang, and Hicham Johra</i>	
Non-Intrusive Load Monitoring and Data Competitions	
ADRENALIN: Energy Data Preparation and Validation for HVAC Load Disaggregation in Commercial Buildings	321
<i>Balázs András Tolnai, Zheng Grace Ma, Igor Sartori, Surya Venkatesh Pandiyan, Matt Amos, Gustaf Bengtsson, Synne Krekling Lien, Harald Taxt Walnum, Akram Hameed, Jayaprakash Rajasekharan, Rafeal Gomez Garcia, and Bo Nørregaard Jørgensen</i>	
Advancing Non-intrusive Load Monitoring: Insights from the Winning Algorithms in the ADRENALIN 2024 Load Disaggregation Competition	338
<i>Balázs András Tolnai, Rafael Sudbrack Zimmermann, Yangxinyu Xie, Ngoc Tran, Cihat Emre Çeliker, Zheng Grace Ma, Igor Sartori, Matt Amos, Gustaf Bengtsson, Synne Krekling Lien, Clayton Miller, Akram Hameed, and Bo Nørregaard Jørgensen</i>	
Comparison of Three Algorithms for Low-Frequency Temperature-Dependent Load Disaggregation in Buildings Without Submetering	355
<i>Balázs András Tolnai, Zheng Ma, and Bo Nørregaard Jørgensen</i>	
Lessons Learned from the ADRENALIN Load Disaggregation Challenge	371
<i>Balázs András Tolnai, Zheng Ma, Igor Sartori, Clayton Miller, Stephen White, Matt Amos, Gustaf Bengtsson, Akram Hameed, and Bo Nørregaard Jørgensen</i>	
Business Model Innovation in Data Competitions: Insights from the 2024 ADRENALIN Load Disaggregation Challenge	388
<i>Zheng Grace Ma, Balázs András Tolnai, and Bo Nørregaard Jørgensen</i>	
Author Index	407

Energy Forecasting and Intelligent Control Systems



A Multi-stage Deep Learning Framework for Short-Term Electricity Load Forecasting

Ege Kandemir^(✉), Agus Hasan, and Saleh Abdel-Afou Alaliyat

Norwegian University of Science and Technology, Ålesund, Norway
ege.kandemir@ntnu.no

Abstract. Electricity load forecasting is essential for economic viability and reliable operations in modern, complex power systems. Although traditional statistical methods and deep learning methods can be utilized for predictions, they tend to fall short in capturing complex non-linear behaviors with increasing demand. For this reason, in this study, a new cascaded method, Multi-Stage Deep Learning Framework (MSDL) is proposed which can capture these behaviors more accurately. MSDL consists of three different architectures, namely Seasonal Decomposition (SD), Empirical Mode Decomposition (EMD), and Long Short-Term Memory (LSTM). In MSDL, the time series data is first decomposed into trend, seasonal, and residual components using SD. The residual component is then further decomposed into Intrinsic Mode Functions (IMFs) using EMD to extract intricate, high-frequency patterns. These components are subsequently used as inputs to LSTM models for forecasting. Benchmarking results demonstrate that the proposed method outperforms several reference models in terms of RMSE, MAE, and MAPE. The study also goes beyond performance comparison by investigating the influence of the quantity of IMFs on forecast accuracy. It has been shown that for the selected case study, including more than five IMFs in forecasting does not significantly increase prediction performance.

Keywords: Electricity Load Forecasting · Cascaded Models · Seasonal Decomposition · Empirical Mode Decomposition (EMD) · Long Short-Term Memory (LSTM)

1 Introduction

Electricity load forecasting is used to predict the electricity demand over some time horizons. This can vary from minutes to years depending on the application purpose. Utilization of historical data and other related variables enhances accurate forecasting for efficient operation in power systems for energy generation, distribution, and pricing [14]. Forecasting the load significantly influences production amount, demand locations, and resource allocation for reliable and

cost-effective energy supply. The increase in complexity of modern power systems such as the integration of renewable energy sources, decentralization, and the alteration in consumption patterns has increased the importance of load forecasting. The varying forecast intervals aim at different purposes; for instance, short-term forecasting may aim to maintain real-time system stability, while medium or long-term forecasting supports maintenance planning and investment decisions [3]. The lack of accurate predictions may lead to overproduction or underproduction, causing economic losses or grid instability.

The complex energy systems have increased the need for improved forecasting methods. Traditional statistical models are useful to some extent; however, most of the time, they fall short in capturing the non-linear behavior of demand, influenced by external variables such as weather conditions, economic activities, or social behaviors. This has led to a growing interest in data-driven and cascaded models [5]. Cascaded models combine multiple approaches like statistical models with machine learning algorithms, or multiple machine learning algorithms to overcome the limitations of standalone models. This approach leverages the strength of each technique for more accurate capture of temporal and seasonal patterns [7].

The main objective of the study is to improve the forecasting accuracy of electric load through a novel cascaded model. In this extension, the study aims to investigate historical load data using various decomposition methods to provide more relevant and condensed information for the predictive model. This approach seeks to address the limitations of traditional methods and deep learning models by using a cascaded approach to contribute to more reliable energy management systems.

1.1 Related Works

Energy load forecasting is essential to support maintenance planning, grid stability, and economic viability. Therefore, there have been several studies available in literature in this respect. In study [1], an analysis of different forecasting techniques is conducted. The study compares different machine learning and deep learning methods besides statistical methods. The analysis shows that the cascaded predictive models perform better than standalone models. Although there have been some studies showcasing the importance of single deep learning algorithms in terms of effectiveness in prediction [2,9], most of the recent research focus on cascaded methods for short-term predictions. In study [21], a cascaded method combining convolutional neural networks (CNN) and sequence-to-sequence (Seq2Seq) method with attention mechanisms is implemented for short-term multi-energy load forecasting. Compared to CNN and long short-term memory (LSTM), the proposed method has higher accuracy. Similar to [21], in the study conducted by [13], a cascaded method is utilized for short-term energy load forecasting using CNN and attention mechanism. However, in this study, instead of the Seq2Seq algorithm, bidirectional gated recurrent units (BiGRU) are implemented. Another similar study is made by [18], implementing LSTM instead of BiGRU. Even though there has not been any comparison between

these three similar studies in terms of performance, in all of them it has been shown that the proposed cascaded models perform better than the standalone algorithms. In the study [8], a cascaded method of CNN and stacked LSTM is used to forecast short-term load. The study compares accuracy with several machine learning and deep learning algorithms, showing the lowest error metric in the proposed method. Another important method used in load forecasting is empirical mode decomposition (EMD). The paper [12] uses EMD and BiLSTM to forecast the electric load by resolving time series into different components. Similarly, in study [10], the EMD is applied to obtain intrinsic modes, then the modes are implemented into CNN and LSTM models with Bayesian optimization to increase accuracy. The study shows that the integration of EMD greatly improves the forecast accuracy. Another study that implemented EMD as a part of a cascaded model is done by [19]. Although forecasting is mainly aimed at wind speed prediction, the proposed method can be applied for load forecasting. The paper suggests using seasonal autoregressive integrated moving average (SARIMA) to extract linear features and nonlinear sequences before implementing EMD and LSTM for prediction.

1.2 Contribution

Based on the literature review, the most commonly used methods for electric load forecasting include CNN, LSTM, GRU, EMD, and attention mechanisms. In this paper, a novel cascaded forecasting framework, Multi-Stage Deep Learning Framework (MSDL), that integrates seasonal decomposition, empirical mode decomposition (EMD), and long short-term memory networks (LSTM) is proposed to enhance the accuracy of hourly electricity load predictions. The aim of this approach is to better inform the deep learning model by leveraging residual components from seasonal decomposition as input to EMD, which then generates multiple informative modes. These modes are subsequently used to train the LSTM model. The effectiveness of the MSDL method is validated by benchmarking it against several forecasting models, demonstrating its superior performance.

1.3 Outline

In Sect. 2, the methodology implemented in this paper is explained. The section includes the architecture of the proposed model (MSDL), along with the other reference models that have been used in this concept. The section also presents the real-world data used in the study with visual representation. In Sect. 3, the results of the MSDL model, along with a comparison of the reference models, are given. The section includes the performance of each model, the required computational time to train each model, and the visual results of forecasting for some instances. Besides the results, a deeper analysis of the findings is discussed in the chapter. In Sect. 4, the key insights of the study are summarized with suggestions for future research.

2 Methodology

2.1 Forecasting Algorithm

This section explains the forecasting algorithms implemented in this study, as listed in Table 1. To evaluate the performance of each model, root mean square error (RMSE), mean absolute error (MAE), and mean absolute percentage error (MAPE) are used [13].

Table 1. Algorithms with Corresponding Architectural Descriptions

Algorithm	Model Architecture
RNN	RNN Layer (20 Units) + Dense Layer (1 Unit)
LSTM	LSTM Layer (20 Units) + Dense Layer (1 Unit)
GRU	GRU Layer (20 Units) + Dense Layer (1 Unit)
SVR	Kernel = rbf, C = 1.0, epsilon = 0.1
MSDL Model	SD + EMD + LSTM Layer (20 Units) + Dense Layer (1 Unit)

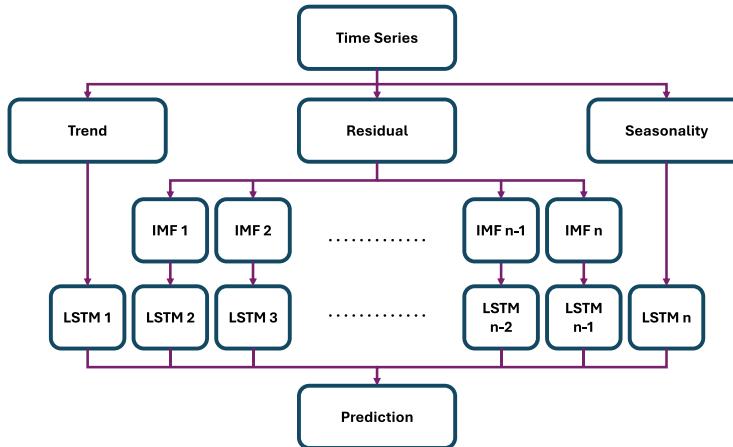


Fig. 1. Proposed Cascaded Model - Multi-Stage Deep Learning Framework (MSDL). The model is seasonally decomposed into trend, residual, and seasonality components. The residual component is further decomposed into intrinsic mode functions (IMFs) using empirical mode decomposition. Finally, the trend, seasonality, and IMFs are modeled with LSTM network for prediction.

Figure 1 presents the architecture of the proposed Multi-Stage Deep Learning (MSDL) algorithm. The proposed cascaded model (MSDL) consists of three main components: Seasonal Decomposition, Empirical Mode Decomposition, and

Long Short-Term Memory (LSTM). The model first decomposes the time series into three components: trend, seasonality, and residual. The residual component is further decomposed into several Intrinsic Mode Functions. The trend, seasonality, and IMFs are then used as inputs to an LSTM model for prediction. Finally, the predictions from each model are combined to generate the final load forecast. This cascaded approach aims to enhance prediction accuracy by decomposing the time series to extract more detailed information from the data.

Recurrent Neural Networks (RNN) are mainly specialized for sequential data. Despite the feedforward neural networks, RNNs have a hidden state h_t that contains the information from previous elements in the sequence. In Eq. 1, the architecture of RNN is given for time t . x_t is the input at time t , h_{t-1} is the hidden state from the previous time step, W_{xh} , W_{hh} , and W_{hy} are weight matrices, b_h and b_y are bias vectors, and σ denotes an activation function. RNNs can capture temporal dependencies, but they often suffer from the vanishing and exploding gradient problems [11].

$$\begin{aligned} h_t &= \sigma(W_{xh}x_t + W_{hh}h_{t-1} + b_h) \\ y_t &= W_{hy}h_t + b_y \end{aligned} \quad (1)$$

Long Short-Term Memory Networks (LSTM) aim to overcome the vanishing gradient problem in RNNs with the introduction of memory cells to keep the information over long periods. LSTM mainly consists of three main gates, which are input, forget, and output. In Eq. 2, a detailed mathematical explanation is given, where \odot denotes element-wise multiplication, σ is the sigmoid activation function, and \tanh is the hyperbolic tangent function. The terms f_t , i_t , C_t , o_t , and h_t refer to the forget gate, input gate, cell state update, output gate, and hidden state, respectively [20].

$$\begin{aligned} f_t &= \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \\ i_t &= \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \\ \tilde{C}_t &= \tanh(W_C \cdot [h_{t-1}, x_t] + b_C) \\ C_t &= f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \\ o_t &= \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \\ h_t &= o_t \odot \tanh(C_t) \end{aligned} \quad (2)$$

Gated Recurrent Unit (GRU) is similar to LSTM, where the forget and input gates merge into a single update gate. This allows computational efficiency compared to LSTM. In Eq. 3, the details of the method are shown. z_t , r_t , \tilde{h}_t , and h_t are the update gate, reset gate, candidate hidden state, and hidden states update, respectively [16].

$$\begin{aligned} z_t &= \sigma(W_z[h_{t-1}, x_t] + b_z) \\ r_t &= \sigma(W_r[h_{t-1}, x_t] + b_r) \\ \tilde{h}_t &= \tanh(W_h[r_t * h_{t-1}, x_t] + b_h) \\ h_t &= (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t \end{aligned} \quad (3)$$

Support Vector Regression (SVR) is the equivalent of support vector machines for regression problems. The idea behind SVR is to find a function that best fits a hyperplane with a certain margin of tolerance ϵ . In Eq. 4, the objective function of SVR is given, and constraints are in Eq. 5. w , b , ϵ , C , ξ_i , and ξ_i^* are the weight vector, bias term, tolerance margin, penalty parameter, and slack variables, respectively [17].

$$\frac{1}{2} \|w\|^2 + C \sum_{i=1}^n (\xi_i + \xi_i^*) \quad (4)$$

$$\begin{aligned} y_i - w^T x_i - b &\leq \epsilon + \xi_i \\ w^T x_i + b - y_i &\leq \epsilon + \xi_i^* \\ \xi_i, \xi_i^* &\geq 0 \quad \text{for all } i = 1, \dots, n \end{aligned} \quad (5)$$

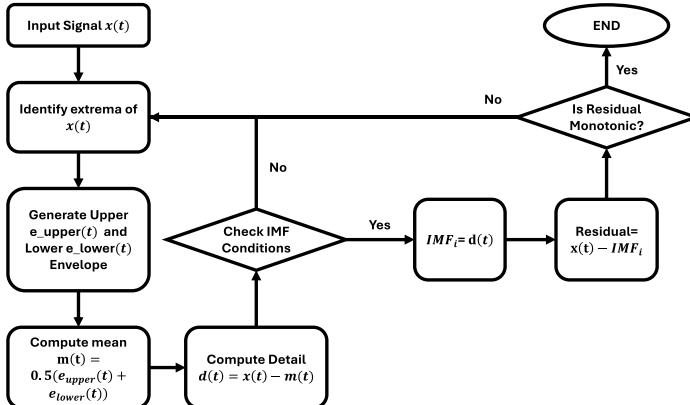


Fig. 2. Empirical Mode Decomposition Flowchart. The method starts with identifying local maxima and minima, from which upper and lower envelopes are created. The average of these envelopes is subtracted to produce a candidate IMF (or detail). This process repeats until the component meets IMF criteria. Once an IMF is extracted, it is removed from the original signal, and the residue is decomposed further.

Seasonal Decomposition (SD) separates the time series into three components: Trend T_t , Seasonality S_t , and Residual R_t . The Trend component computes long-term progression using a centered moving average, while Seasonality focuses on capturing repeating short-term cycles. The Residual is the remaining component after subtracting seasonality and trend from the original data. Seasonal decomposition can have two models: additive and multiplicative. In Eq. 6, the additive model is given, and in Eq. 7, the multiplicative model is presented [6].

$$y_t = T_t + S_t + R_t \quad (6)$$

$$y_t = T_t \times S_t \times R_t \quad (7)$$

Empirical Mode Decomposition (EMD) is a data-driven signal processing technique that decomposes nonlinear and non-stationary time series into a finite set of Intrinsic Mode Functions (IMFs) and a residual. Each IMF needs to satisfy two criteria: (1) the number of extrema and the number of zero crossings must either be equal or differ by at most one, and (2) the mean value of the envelope defined by the local maxima and minima must be close to zero. The flowchart of Empirical Mode Decomposition is shown in Fig. 2. In Eq. 8, the signal decomposition is illustrated, where $IMF_i(t)$ represents the i^{th} intrinsic mode function, $r_n(t)$ is the final residual, and $x(t)$ is the signal of interest [15].

$$x(t) = \sum_{i=1}^n IMF_i(t) + r_n(t) \quad (8)$$

2.2 Data Description

In this study, hourly aggregated electricity load values from the European Network of Transmission System Operators for Electricity (ENTSO-E) are utilized [4]. The dataset is univariate, consisting of hourly load values from various European countries. However, this paper focuses solely on the electricity load statistics for Norway, covering the period from January 1st to December 31st, 2024. The load values, measured in megawatts (MW), range from a minimum of 9,862 MW to a maximum of 24,930 MW throughout the year. Figure 3 displays the time series of hourly load data for Norway in 2024. Prior to analysis, the data was normalized to a 0–1 scale. The time series is structured as sequences of 24 hourly instances, with each sequence representing a single day load pattern.

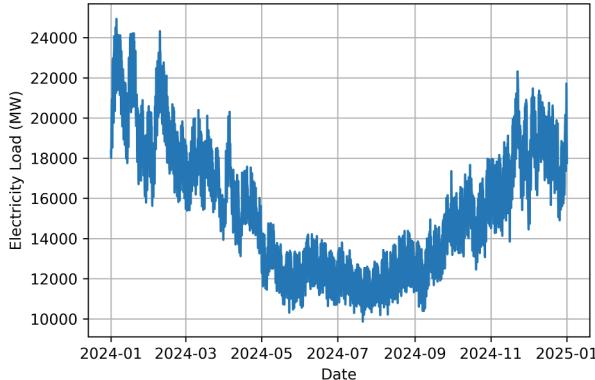


Fig. 3. Time series plot showing hourly electricity load (in MW) for Norway from January 1st to December 31st, 2024. The data sourced from ENTSO-E each point represents the total load recorded at a given hour.

3 Results and Discussion

In this section, the results of the study are discussed through detailed analysis and comparison. As shown in Table 1, five different architectures were implemented: RNN, LSTM, GRU, SVR, and a MSDL. Each model was trained on the same dataset, which was split into training and testing subsets. The goal of this analysis is to evaluate forecast performance and assess the impact of the MSDL on prediction accuracy. The four standalone models, which are RNN, LSTM, GRU, and SVR, were used as baseline references. Among them, as shown in Table 2 the LSTM model achieved the best performance, with RMSE, MAE, and MAPE values of 0.017671, 0.013299, and 0.027436, respectively. This aligns with findings on the effectiveness of LSTM in capturing long-term dependencies in time-series data. On the other hand, SVR showed the worst performance, with RMSE, MAE, and MAPE values of 0.045474, 0.036862, and 0.085639, respectively. GRU and RNN performed moderately but did not outperform LSTM. These results confirm that deep learning architectures are more capable of modeling complex temporal dependencies in electric load forecasting compared to traditional machine learning models like SVR.

Table 2. Performance Comparison of Models on Test Data Based on RMSE, MAE, and MAPE

Model Number	Model Name	RMSE	MAE	MAPE	Training Time (s)
Model 1	RNN	0.019070	0.014231	0.028882	34.044880
Model 2	LSTM	0.017671	0.013299	0.027436	142.583476
Model 3	GRU	0.018305	0.014054	0.028653	127.845189
Model 4	SVR	0.045474	0.036862	0.085639	0.067142
Model 5	MSDL Model (7 IMF)	0.009097	0.007255	0.014979	324.863204
Model 6	MSDL Model (6 IMF)	0.009102	0.007258	0.014997	311.788076
Model 7	MSDL Model (5 IMF)	0.009440	0.007518	0.015567	293.982230
Model 8	MSDL Model (4 IMF)	0.010761	0.008559	0.017765	270.970463
Model 9	MSDL Model (3 IMF)	0.016582	0.012971	0.026944	234.412794
Model 10	MSDL Model (2 IMF)	0.023838	0.019532	0.040799	191.949839

The cascaded model was constructed by combining the Seasonal Decomposition and EMD of the time series with an LSTM-based learning algorithm. Six variations of the cascaded model were trained, each using a different number of IMFs to investigate their contribution to forecasting. The best results were obtained when the first seven IMFs were included as given in Table 2, achieving RMSE, MAE, and MAPE values of 0.009097, 0.007255, and 0.014979, respectively. In contrast, using only the first two IMFs resulted in the highest error metrics among the other cascaded models, with corresponding values of 0.023838, 0.019532, and 0.040799. These results demonstrate that increasing the

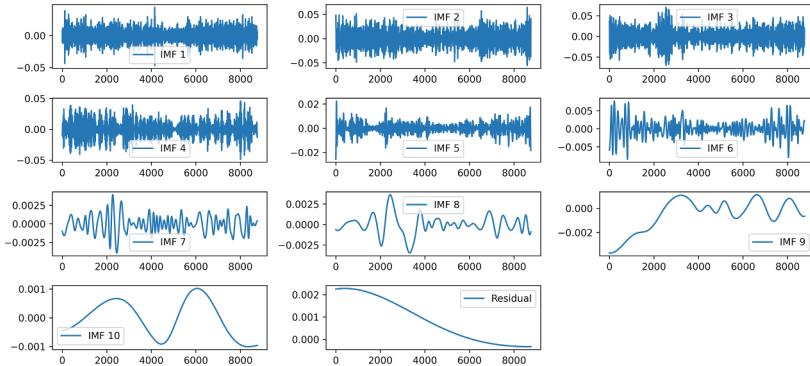


Fig. 4. Intrinsic Mode Functions (IMFs) extracted using Empirical Mode Decomposition (EMD) applied to the residual component of a seasonal decomposition of time series data. The ten resulting IMFs and residual capture distinct temporal patterns aiming to improve forecasting performance.

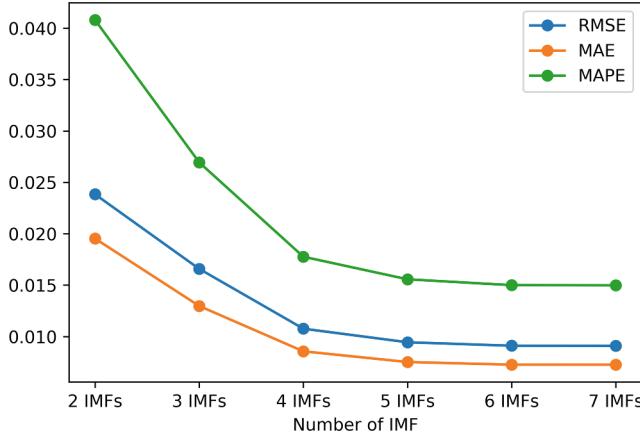


Fig. 5. Comparison of the number of Intrinsic Mode Functions (IMFs) and model performance. The variation in three error metrics is shown: RMSE is represented by the blue curve, MAE by the orange curve, and MAPE by the green curve. The x-axis indicates the number of IMFs used in the model, and the y-axis shows the corresponding error values for model predictions on the test data. (Color figure online)

number of IMFs initially enhances accuracy by providing the model with more information. However, the marginal benefit declines after a certain point, as higher-order IMFs contain less information. Figure 4 shows the extracted IMFs and the residual component of the signal. As the decomposition progresses, the IMFs become smoother and less informative, indicating a decreasing influence of returns in further decomposition. Notably, as in Table 2 model performance improves significantly as IMFs are added up to the fifth component. Beyond that point (from Cascaded Model - 5 IMF), the computational cost increases more

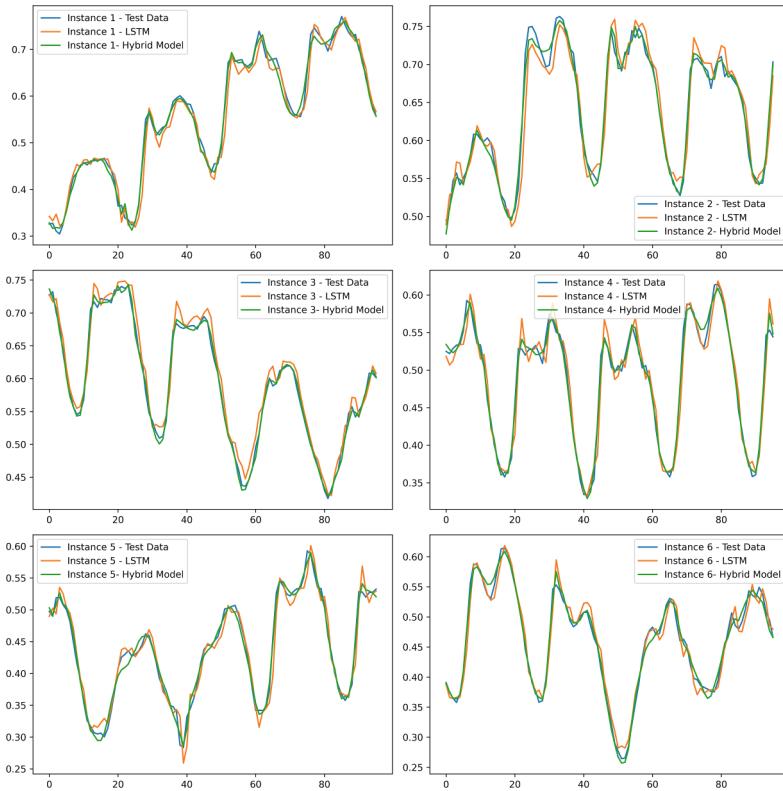


Fig. 6. Comparison of LSTM, Cascaded Model, and test data over 96 h for six randomly selected cases. Blue curve shows test data, orange is the LSTM forecast, and green is the Proposed Cascaded model (MSDL). The x-axis represents hours, and the y-axis shows scaled electric load. (Color figure online)

than the marginal gain in accuracy as shown in Fig. 5. In all cases, using varying numbers of IMFs in cascaded models resulted in additional computational expense, as shown in Table 2, due to increased model complexity. For instance, Model 10, which used two IMFs, did not improve prediction accuracy compared to Model 2, which employed a standalone LSTM architecture. However, the inclusion of a third IMF led to increased accuracy, as evidenced by improvements across three different error metrics. This observation is consistent with the findings of [10], who showed the correlation between the number of IMFs. It can support the search for an optimal balance between accuracy and resource usage.

In addition to the numerical analysis, Fig. 6 presents six randomly selected, 96-hour scaled prediction instances. Visually, the cascaded model follows the test data more closely than the standalone LSTM. While minor prediction overshoots are observed in the cascaded model for some time steps, overall, it shows a bet-

ter fit to actual values. In contrast, LSTM alone exhibits more deviations. This visual inspection supports the quantitative findings and highlights the advantage of decomposing complex time-series data into more detailed components before training. Besides improved accuracy, the cascaded model supports greater interpretability. By examining the influence of individual IMFs, it provides insights into which components contribute most significantly to the prediction. This is especially beneficial for energy demand forecasting that requires explainability. However, it should be noted that including more IMFs for forecasting increases the computational cost of training the model. Therefore, depending on the application and dataset, the most efficient number of IMFs should be determined.

4 Conclusion

In this study, a comprehensive analysis of various forecasting algorithms for hourly electricity load prediction was conducted. Among the standalone models, LSTM achieved the best performance due to its ability to capture long-term temporal dependencies. However, MSDL significantly outperformed all baseline models, including LSTM. Notably, incorporating the first seven Intrinsic Mode Functions (IMFs) led to the highest accuracy with the lowest RMSE, MAE, and MAPE values. The study also highlights a trade-off between predictive accuracy and computational cost when including a larger number of IMFs. The forecasting results were further validated through visual inspection, which aligned with quantitative metrics. Future research could explore alternative time series decomposition methods, such as Variational Mode Decomposition or Wavelet Transform, to enhance interpretability and further improve forecasting performance.

Acknowledgement. This work is partially supported by the Horizon Europe TWINVEST project (Grant No. 101146936).

References

1. Ahmad, N., Ghadi, Y., Adnan, M., Ali, M.: Load forecasting techniques for power system: research challenges and survey. *IEEE Access* **10**, 71054–71090 (2022)
2. Amarasinghe, K., Marino, D.L., Manic, M.: Deep neural networks for energy load forecasting. In: 2017 IEEE 26th International Symposium on Industrial Electronics (ISIE), pp. 1483–1488 (2017)
3. Azeem, A., Ismail, I., Jameel, S.M., Harindran, V.R.: Electrical load forecasting models for different generation modalities: a review. *IEEE Access* **9**, 142239–142263 (2021)
4. European Network of Transmission System Operators for Electricity: Monthly hourly load values. <https://www.entsoe.eu/data/power-stats>
5. Holderbaum, W., Alasali, F., Sinha, A.: Short Term Load Forecasting (STLF), pp. 13–56. Springer, Cham (2023)
6. Hyndman, R.J., Athanasopoulos, G.: Classical decomposition, 2nd edn., Chap. 6.3. OTexts, Melbourne, Australia (2018)

7. Kandemir, E., Hasan, A., Kvamsdal, T., Abdel-Afou Alaliyat, S.: Predictive digital twin for wind energy systems: a literature review. *Energy Inform.* **7**(1), 68 (2024)
8. Khan, Z.A., et al.: Efficient short-term electricity load forecasting for effective energy management. *Sustain. Energy Technol. Assess.* **53**, 102337 (2022)
9. Kuo, P.H., Huang, C.J.: A high precision artificial neural networks model for short-term energy load forecasting. *Energies* **11**(1) (2018)
10. Lotfipoor, A., Patidar, S., Jenkins, D.P.: Deep neural network with empirical mode decomposition and Bayesian optimisation for residential load forecasting. *Exp. Syst. Appl.* **237**, 121355 (2024)
11. Mienye, I.D., Swart, T.G., Obaido, G.: Recurrent neural networks: a comprehensive review of architectures, variants, and applications. *Information* **15**(9) (2024)
12. Mounir, N., Ouadi, H., Jrhilifa, I.: Short-term electric load forecasting using an EMD-BI-LSTM approach for smart grid energy management system. *Energy Build.* **288**, 113022 (2023)
13. Niu, D., Yu, M., Sun, L., Gao, T., Wang, K.: Short-term multi-energy load forecasting for integrated energy systems based on CNN-BiGRU optimized by attention mechanism. *Appl. Energy* **313**, 118801 (2022)
14. Nti, I.K., Teimeh, M., Nyarko-Boateng, O., Adekoya, A.F.: Electricity load forecasting: a systematic review. *J. Electr. Syst. Inf. Technol.* **7**(1), 13 (2020)
15. Rilling, G., Flandrin, P., Gonçalves, P.: On empirical mode decomposition and its algorithms. In: IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing NSIP-03, Grado, Italy, June 2003
16. Shiri, F.M., Perumal, T., Mustapha, N., Mohamed, R.: A comprehensive overview and comparative analysis on deep learning models. *J. Artif. Intell.* **6**(1) (2024)
17. Smola, A.J., Schölkopf, B.: A tutorial on support vector regression. *Stat. Comput.* **14**(3), 199–222 (2004)
18. Wan, A., Chang, Q., AL-Bukhaiti, K., He, J.: Short-term power load forecasting for combined heat and power using CNN-LSTM enhanced by attention mechanism. *Energy* **282**, 128274 (2023)
19. Yan, Y., Wang, X., Ren, F., Shao, Z., Tian, C.: Wind speed prediction using a hybrid model of EEMD and LSTM considering seasonal features. *Energy Rep.* **8**, 8965–8980 (2022)
20. Yu, S., Hur, J.: An enhanced performance evaluation metrics for wind power ramp event forecasting. *IEEE Access* **11**, 100195–100206 (2023)
21. Zhang, G., Bai, X., Wang, Y.: Short-time multi-energy load forecasting method based on CNN-Seq2Seq model with attention mechanism. *Mach. Learn. Appl.* **5**, 100064 (2021)



Solaris AI: Enhancing Solar Energy Forecasting with Generative AI and Deep Learning

Mohammed Farhan Faisal¹(✉), Nimisha Nixon², and Pamba Raja Varma³

¹ Department of Computer Science and Engineering, Manipal Academy of Higher Education, Dubai, UAE

farhan.faisal@dxb.manipal.edu

² Department of Data Science and Engineering, Manipal Academy of Higher Education, Dubai, UAE

nimisha.nixon@dxb.manipal.edu

³ School of Engineering and IT, Manipal Academy of Higher Education, Dubai, UAE
pamba.rajavarma@manipaldubai.com

Abstract. The growing adoption of solar energy presents both an opportunity and a challenge. While clean and sustainable, solar power generation is inherently variable due to changing weather conditions, making reliable forecasting critical for grid integration, infrastructure planning, optimal solar panel placement, and overall energy stability. Solaris AI addresses these challenges by integrating Generative AI with traditional Deep Learning models, creating an intelligent and explainable expert system for solar energy analysis. The system initially featured an Artificial Neural Network (ANN) model trained on key meteorological parameters such as cloud cover, humidity, temperature, and wind speed to predict solar power output. It was later enhanced with a Long Short-Term Memory (LSTM) model, which is better suited for handling time-series patterns in solar power output. To make these insights accessible to a wider audience, Solaris AI incorporates Retrieval-Augmented Generation (RAG) using OpenAI's language models and LangChain. The model, along with the dataset, is integrated into the system to generate explainable and accurate forecasts through natural language interaction. This enables stakeholders, including government planners, policymakers, businesses, grid operators, and citizens, to interact with the system conversationally, retrieve relevant data, and receive model predictions. A clean web interface with cloud deployment ensures scalability and usability. By bridging the gap between technical forecasts and practical decision-making, Solaris AI supports the broader adoption and integration of solar energy worldwide.

Keywords: Solar forecasting · Deep Learning · ANN · LSTM · Generative AI · RAG · Explainable AI · Clean energy · Expert system · OpenAI

1 Introduction

In response to the increasing global energy demands and the urgent shift toward sustainable alternatives, solar energy has emerged as the cornerstone of the renewable energy revolution. It is an environmentally sustainable and viable source of energy with significant economic benefits. Optimizing solar energy usage and storage for future requires efficient prediction of solar power output and this is where solar forecasting methods play a crucial role [1]. It has been estimated that energy demand will increase by 48% between 2012–2040 [2,3]. With the global shift towards clean energy, efficient solar power generation and forecasting have become crucial to optimize energy distribution, minimize power losses, and improve grid stability, especially in Middle Eastern countries, where sunlight is radiant almost throughout the year. Although multiple renewable energy sources are being explored in the global market, solar energy and wind energy are most preferred [1].

Despite the growing adoption of solar energy, much of the research in this space has traditionally been rooted in electrical and mechanical engineering, focusing on hardware efficiency and photovoltaic material science. As a result, solar forecasting remains underexplored in the context of artificial intelligence [4]. When AI techniques are used, they often result in black-box models that lack interpretability, making it difficult for stakeholders to trust and act on predictions. In recent years, however, deep learning has shown promise in modelling complex, nonlinear relationships within time-series data, leading to more accurate and data-driven solar energy forecasting.

Solaris AI is designed to address these challenges. It combines the predictive capabilities of Deep Learning with the interpretability of Generative AI to create an intelligent and explainable solar forecasting system. The system initially employed an Artificial Neural Network (ANN) trained on historical meteorological features. It was later extended to a Long Short-Term Memory (LSTM) model, which is better suited for capturing long-range temporal dependencies in solar generation patterns. To make this forecasting process transparent and accessible, we incorporate a Retrieval-Augmented Generation (RAG) pipeline that allows the large language model (LLM) to interact not only with the historical dataset but also with the trained model.

This means that Solaris AI can answer questions based on two sources: the underlying CSV data set containing detailed solar parameters such as temperature, irradiance, and humidity, and the model's real-time output. When a user asks, for example, "Why was solar output low yesterday?" or "How much power was generated this week?", the system retrieves relevant data, interprets it, and presents an accurate, human-readable answer grounded in both empirical data and learned predictions while also preventing hallucinations. The RAG system was implemented using LangChain, an open-source framework designed for building applications with reliable large language models.

The deployment of Solaris AI is structured to be scalable, modular, and accessible through an API built using FastAPI, and is containerized using Docker to ensure consistent deployment across environments. The application is hosted

serverlessly via AWS Lambda, enabling cost-efficient and on-demand operation without maintaining persistent infrastructure. DynamoDB is used to store queries and responses for auditing and future analysis, while asynchronous processing via secondary Lambda functions ensures a responsive user experience.

By combining deep learning forecasting with generative reasoning, and deploying it through a robust cloud infrastructure, Solaris AI ensures that users receive both reliable predictions and meaningful explanations, eliminating the black-box nature of typical AI models. The result is a truly intelligent assistant that not only forecasts but also understands and communicates solar energy insights effectively.

In the following sections, we explore the technical components of Solaris AI in detail, including the model architectures, dataset selection, RAG pipeline, and deployment strategy. We also present a comparative evaluation of ANN and LSTM models to highlight the evolution and performance gains of the system.

2 Related Work

2.1 Solar Energy Forecasting Using Machine Learning and Deep Learning Techniques

Solar energy is a widely available and environmentally friendly resource, making it a key focus in the global transition toward renewable energy. Forecasting solar energy accurately is essential to ensure its efficient use and integration into modern systems. Initially, global solar radiation (GSR) forecasting was performed using traditional mathematical and statistical models. However, recent advances in artificial intelligence (AI), with machine learning and deep learning techniques, have proven to be more effective in modelling complex, nonlinear relationships between weather variables and solar power output. As a result, solar forecasting has become more reliable and data-driven, supporting better planning and decision-making in energy systems [1].

2.2 Hybrid Deep Learning Models for Time Series Forecasting of Solar Power

Accurate forecasting of solar power generation is vital for managing renewable energy systems. One study explored the use of hybrid deep learning models that combine Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM), and Transformer architectures for time-series forecasting.

The research compared various combinations of these models, such as CNN-LSTM-Transformer and Transformer-LSTM, against standalone versions. Performance was evaluated using metrics, such as Mean Absolute Error (MAE), and optimizers such as Nadam were tested. The best results were achieved using the CNN-LSTM-Transformer hybrid model, which achieved an MAE of just 0.551%. In contrast, the Transformer-LSTM model performed poorly (MAE of 16.17%), highlighting the importance of architecture choice in solar forecasting. This study is among the first to examine transformer networks in hybrid forecasting models for solar energy [4].

2.3 Investigating Photovoltaic Solar Power Output Forecasting Using Machine Learning Algorithms

Integration of solar power into electrical grids can be a complicated process as it is highly dependent on daily weather conditions. To acknowledge this problem, there has been constant research and development to determine the best machine learning algorithms for PV solar power forecasting. Existing traditional models, such as artificial neural networks(ANN) and random forest(RF) algorithms, have shown strong performance. More recent studies have explored alternative algorithms, including Decision Trees (DT), Extreme Gradient Boosting (XGB), and Long Short-Term Memory (LSTM) networks. These newer models offer promising results for photovoltaic (PV) forecasting and highlight the growing role of AI in managing renewable energy systems [5].

2.4 EF-LLM: Energy Forecasting LLM with AI-Assisted Automation, Enhanced Sparse Prediction, Hallucination Detection

Recent advancements in large language models (LLMs) have also been applied to energy forecasting. The EF-LLM (Energy Forecasting LLM) system was proposed to overcome the limitations of traditional forecasting models, which often require expert intervention and struggle with sparse data. EF-LLM combines temporal data with domain-specific knowledge to support both pre-forecasting and post-forecasting decision-making. It uses a multichannel architecture to handle different types of input (such as text, numerical, and visual) and uses LoRA-based modules for continuous learning. By enabling human-AI interaction and incorporating multimodal data sources, EF-LLM enhances the accuracy and accessibility of energy forecasting, particularly in data-scarce scenarios [6].

3 Methodology

3.1 Data Selection and Preprocessing

Two datasets were used in different stages of development. The initial ANN model was trained on a GitHub dataset [7] lacking irradiance parameters. For the LSTM model, we used a richer, publicly available dataset [8] that includes key irradiance features (GHI, GTI) along with hourly solar power output and weather attributes such as temperature, humidity, wind speed, and UV index.

The power data was originally recorded at 15-min intervals, while weather data was hourly. To align these sources, the power data was resampled to hourly intervals using aggregation, and then merged with the weather data by timestamp.

Missing values were handled by dropping columns with high null counts and imputing remaining missing values using the column median. Negative or placeholder values were also cleaned using median substitution. Outliers in the power

column were replaced using values from the previous matching time period to preserve temporal consistency.

Categorical features (e.g., weather icon) were encoded using one-hot encoding, and all numerical features were standardized to ensure uniform scaling during LSTM training. The dataset was split chronologically into training (75%) and testing (25%) sets to simulate realistic forecasting and avoid data leakage.

This cleaned and engineered dataset was used to train both the linear regression baseline and the LSTM models for historical and forecast solar power prediction.

3.2 Initial ANN Model

The initial ANN model was built using TensorFlow's Keras and trained on the earlier GitHub dataset. It comprised several fully connected layers with ReLU activations and a final dense output layer. The model was trained using Mean Squared Error (MSE) as the loss function and evaluated using MSE, Root Mean Squared Error (RMSE), and R^2 score.

While the ANN achieved an R^2 of 0.778 on the test set, it lacked the ability to capture temporal dependencies and was limited by the absence of irradiance features. As such, the ANN is treated as a baseline, and the system was expanded with a time-series-aware LSTM model trained on a more comprehensive dataset.

This model is developed using a sequential architecture implemented via TensorFlow's Keras. The network structure is dynamic, allowing for flexible configuration of the number of layers and neurons per layer. In its standard configuration, the model comprises multiple fully connected (dense) layers, each initialized with a specified kernel initializer and activated using a nonlinear activation function—typically ReLU. The first layer accepts an input dimension equivalent to the number of features in the dataset and is followed by one or more hidden layers with the same activation and kernel parameters.

The final output layer is a single-node dense layer, designed to predict the solar energy generation value. The model is compiled using the Mean Squared Error (MSE) loss function, with the Adam optimizer to facilitate efficient gradient-based optimization. Mean Squared Error (MSE) is used as the primary loss function during the training of the ANN model. It measures the average of the squares of the errors—that is, the average squared difference between the actual and predicted values. The formula is defined as:

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (1)$$

where y_i is the actual value, \hat{y}_i is the predicted value, and n is the number of observations.

Additionally, Root Mean Squared Error (RMSE) is tracked as a performance metric during training. The Root Mean Squared Error (RMSE) is used as a key evaluation metric to assess the performance of the ANN model. It measures

the average magnitude of the error between predicted and actual values, and is defined as:

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (2)$$

where y_i represents the actual value, \hat{y}_i is the predicted value, and n is the total number of observations.

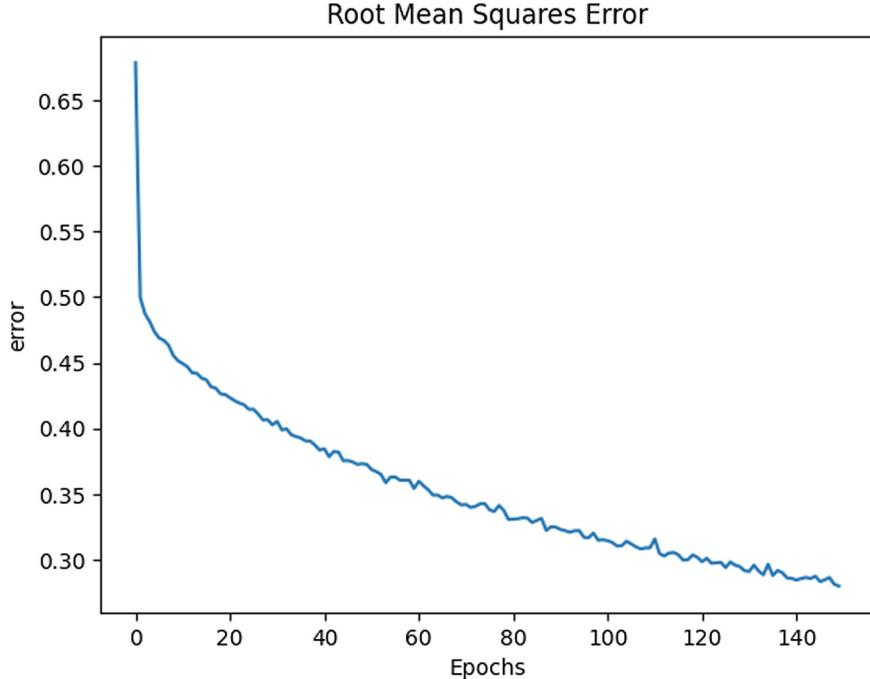


Fig. 1. Root Mean Squared Error (RMSE) curve during ANN model training

As shown in Fig. 1, the RMSE decreased steadily during training, indicating effective learning and convergence of the ANN model.

Dropout layers were optionally included during experimentation to explore regularization, although they were commented out in the final architecture due to acceptable training stability. This configurable and modular ANN architecture enables Solaris AI to adapt to various dataset sizes and complexity levels, offering a scalable solution for accurate solar energy forecasting.

Model Evaluation. After training the Artificial Neural Network (ANN), we evaluated its performance on both training and testing datasets. The dataset

was split in an 80:20 ratio. Predictions were generated for both sets using the trained model, and the outputs were inverse-transformed to their original scale for accurate error computation. We employed three evaluation metrics:

- **Mean Squared Error (MSE)** – measures the average squared difference between predicted and actual values.
- **Root Mean Squared Error (RMSE)** – the square root of MSE, providing error in the original units.
- **Coefficient of Determination (R^2 Score)** – indicates how well the predictions approximate the actual values.

On the test set, the model achieved an RMSE of approximately **434.02**, an MSE of **259.68**, and an R^2 score of **0.778**. For the training set, the R^2 score was **0.917**, suggesting that the model generalizes well without significant overfitting.

To visualize prediction quality, we generated scatter plots comparing the predicted versus actual values for both training and testing datasets (Fig. 2). The clustering of points around the diagonal line indicates strong model performance.

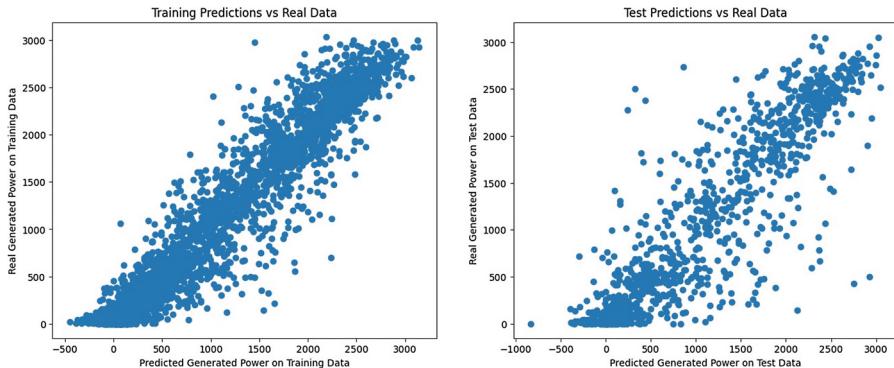


Fig. 2. Scatter plots of predicted vs. actual generated power for training (left) and testing (right) datasets.

3.3 Liner Regression Model for Comparison

Linear Regression is a statistical method used to model relationship between a dependent variable or target variable, which in our case is solar power, and one or more independent variables, again in our case is weather and irradiance features. We use this model to serve as a baseline for solar power prediction which we will use to compare against our LSTM model which we will train later.

Model Training. We used our prepared final dataset [8] to train the model. The dataset is split into a 75:25 ratio, where 75% of the dataset is used for training and 25% of the dataset is used for testing to evaluate model generalization. We train our Linear Regression model on the training set using the prepared features. The model learns the optimal coefficients to minimize the mean squared error between predicted and actual power values.

Model Evaluation. We evaluate the training model on test dataset. We then use metrics such as MSE (Mean Squared Error), RMSE (Root Mean Squared Error), MAE (Mean Absolute Error), R^2 and Adjusted R^2 . Our evaluation metrics results are as follows: Mean Square Error is 159.44, Root Mean Square Error is 12.62, Mean Absolute Error is 8.56, Median Absolute Error is 6.14, R^2 is 0.57, Adjusted R^2 is 0.57. We then visualize the results by plotting the predicted vs actual power values as seen in Fig. 3. Finally, the trained Linear Regression model is saved for future use and comparison with other models.

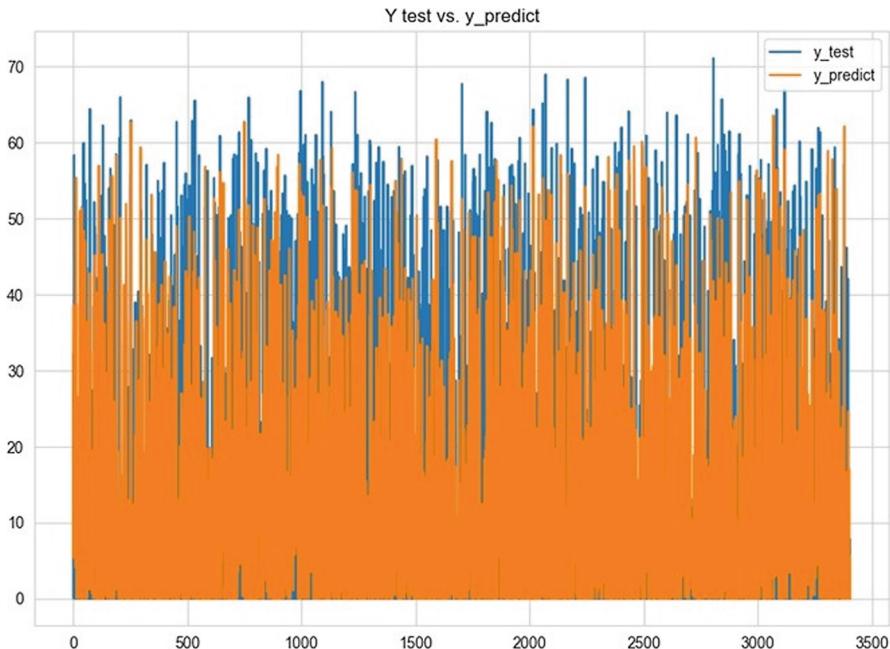


Fig. 3. Comparison plot of actual vs. predicted solar power values in the Linear Regression model

3.4 LSTM Model for Time-Series Forecasting

We use LSTM to model the sequential nature of solar power generation based on historical weather conditions and irradiance data. We use the same set of

features that were used for Linear Regression. This feature matrix is reshaped into a 3D array as required by LSTM, i.e. samples, time steps, features. In our case, each sample represents one hour, with a single time step. The dataset is split into a 75:25 ratio, where 75% of the dataset is used for training and 25% of the dataset is used for testing while maintaining the chronological order. The LSTM model consists of:

1. A LSTM layer with units equal to the lag value.
2. A Dropout layer to prevent overfitting.
3. A Dense output layer to prevent overfitting.

We trained two versions of the LSTM model using different lag values: 24 (1 day) and 672 (4 weeks). Each input sample consisted of historical weather and irradiance values across the lag window. The LSTM architecture included an LSTM layer with a number of units equal to the lag, a dropout layer (0.3–0.33), and a final dense output layer. The model was trained using the Adam optimizer (Adaptive Moment Estimation) to efficiently update model weights and MSE loss function.

Model Evaluation. After training our LSTM models, the model's predictions are compared to actual values using the same evaluation metrics as Linear Regression, i.e. MSE (Mean Squared Error), RMSE (Root Mean Squared Error), MAE (Mean Absolute Error), R^2 . Our evaluation metrics for training lag 24 are as follows: Mean Square Error is 134.53, Root Mean Square Error is 11.59, Mean Absolute Error is 6.95, Median Absolute Error is 2.54, R^2 is 0.643, Adjusted R^2 is 0.641. Similarly, our evaluation metrics for training lag 672 are as follows: Mean Square Error is 130.93, Root Mean Square Error is 11.44, Mean Absolute Error is 6.64, Median Absolute Error is 2.28, R^2 is 0.653, Adjusted R^2 is 0.641. We plot training and validation loss curves to assess model convergence as well as visualizing a predicted vs actual power values as seen in Fig. 4. Moreover, we also visualize the results by plotting the predicted vs actual power values for both lag 24 (Fig. 5) and lag 672 (Fig. 6).

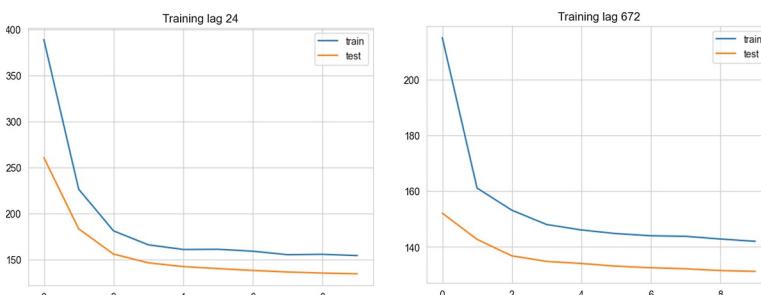


Fig. 4. RMSE training loss curves for LSTM with lag 24 vs lag 672

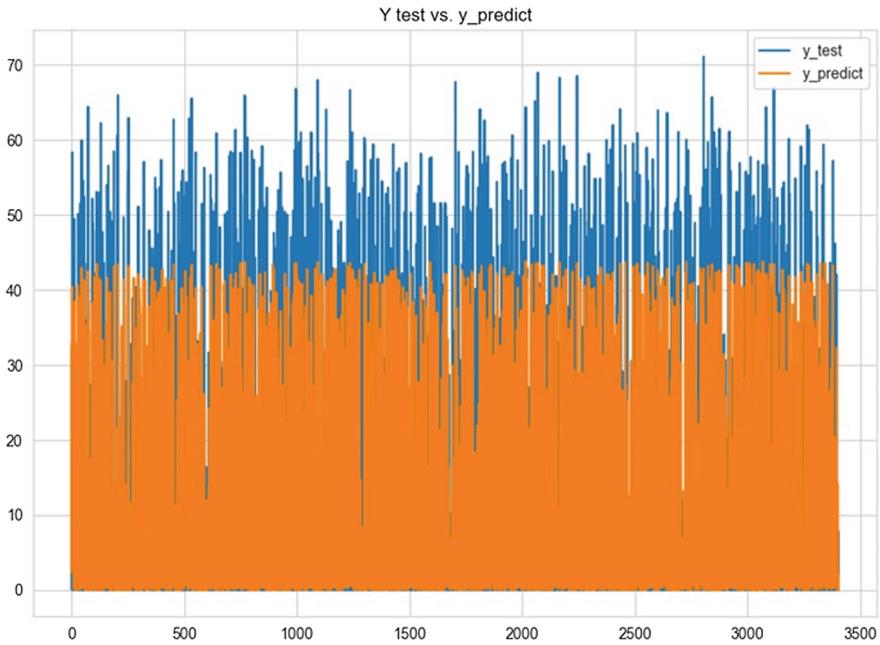


Fig. 5. Comparison plot of actual vs. predicted solar power values in the LSTM Lag 24 Model

This transition to LSTM reflects a key evolution in Solaris AI and addresses reviewer concerns about benchmarking with advanced forecasting models.

3.5 LangChain Framework for Retrieval-Augmented Generation (RAG)

LangChain is an open source framework designed to simplify the development of applications built by large language models (LLMs) by connecting them with external data sources and computational logic [9]. In Solaris AI, LangChain plays a central role in implementing the Retrieval Augmented Generation (RAG) architecture, enabling seamless integration between solar forecasting data, vector embeddings, and OpenAI's generative models.

While RAG pipelines are typically applied to unstructured data sources such as PDFs and documents, Solaris AI adapts this architecture to work with a structured CSV dataset containing historical solar parameters such as irradiance, humidity, temperature, and power output. LangChain facilitates this adaptation by offering modular components for document loading, text splitting, embedding generation, vector store management, and query chaining [10].

This process begins by preprocesing solar dataset using LangChain's 'TextSplitter', which breaks the data into smaller, meaningful chunks [11]. These chunks are transformed into high-dimensional vector embeddings using OpenAI's

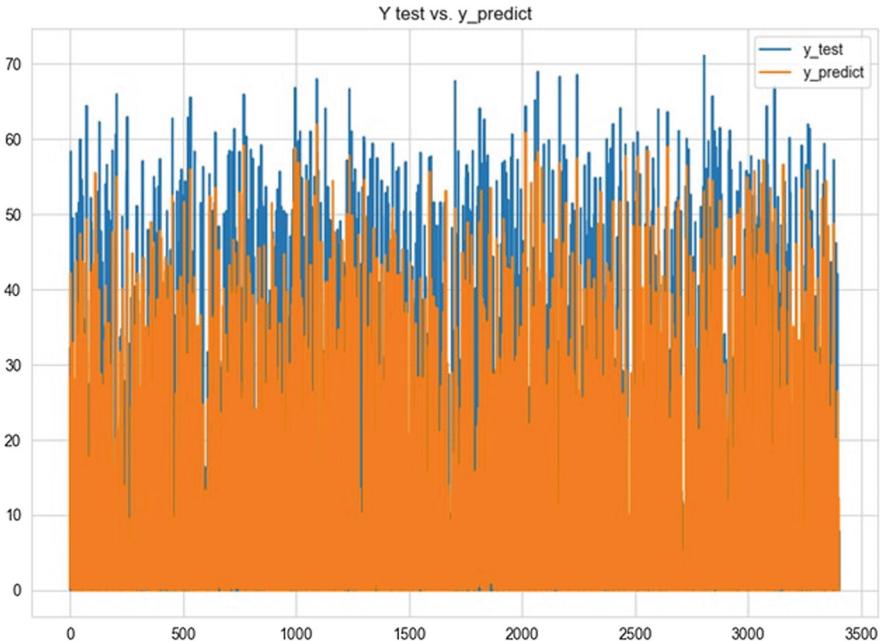


Fig. 6. Comparison plot of actual vs. predicted solar power values in the LSTM Lag 672 Model

embedding model. The embeddings are stored in ChromaDB, a vector database designed for fast semantic retrieval [13]. By structuring complex LLM workflows into reusable and modular components, LangChain ensures that Solaris AI’s RAG system remains both scalable and modular.

Upon receiving a user query, LangChain retrieves the most semantically relevant chunks from ChromaDB and combines them with the original query. This context is then passed to the LLM, which generates a response based on real data and the ANN model’s prediction. This retrieval-based architecture ensures the system’s reliability by tying each answer to factual, context-specific evidence, thereby reducing hallucinations.

Figure 7 illustrates the chunking and embedding workflow used in this RAG pipeline. Figure 8 displays the working RAG application.

3.6 Cloud Deployment Using AWS

To make the Solaris AI system accessible, scalable, and cost-efficient, the RAG based application was deployed using Amazon Web Services (AWS). The deployment strategy was chosen to support real-time user interaction with the model through a serverless architecture to ensure flexibility and minimize infrastructure management overhead [15].

The core application is served through a FastAPI framework, which provides a lightweight RESTful API interface for submitting user queries and receiving

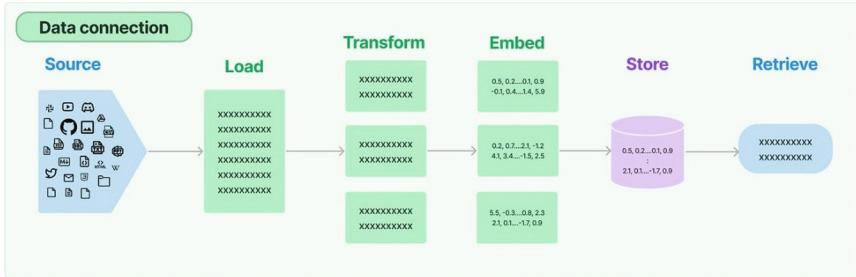


Fig. 7. Illustration of the chunking and embedding workflow used in this RAG pipeline. Taken from [14]

```
User question:
Why is solar power output lower on 2019-02-07 at 12:00:00. Wouldn't it be higher since solar radiation intensity is highest around noon?

Response:
To understand why the solar power output was lower than expected on 2019-02-07 at 12:00:00, let's look at the weather conditions recorded for that date and time:
{"datetime": "2019-02-07 12:00:00", "power": 16.41, "cloud_cover": 0.76, "apparent_temperature": 3.3, "temperature": 4, "humidity": 0.77, "dew_point": -3.04, "wind_bearing": 253.0, "wind_speed": 3.58, "wind_gust": 9.77, "pressure": 1028.89, "uv_index": 2.0, "ozone": 305.7, "precip_intensity": 0.0, "precip_probability": 0.0, "precip_type": null, "visibility": 16.09, "icon": "partly-cloudy-day", "summary": "partly cloudy", "updated_at": "2019-08-16 12:24:17"}
```

The key factor appears to be the high cloud cover of 0.76 (76% of the sky was covered by clouds) at that time. Clouds block incoming solar radiation, reducing the amount reaching the solar panels and lowering power output.

While solar intensity is typically highest around noon, the partly cloudy conditions on 2019-02-07 at 12:00:00 resulted in less than ideal solar irradiance hitting the panels, leading to the lower than expected 16.41 kW power output recorded.

Other factors like temperature (4°C), wind speed (3.58 m/s) and wind gusts (9.77 m/s) may have also played a minor role, but cloud cover was likely the dominant reason for the reduced midday solar generation on that date.

```
Sources: ['forecast.csv:group:326', 'forecast.csv:group:328', 'forecast.csv:group:343', 'forecast.csv:group:342', 'forecast.csv:group:348', 'forecast.csv:group:340']
```

Fig. 8. Working RAG Application

generated responses. RESTful API (Representational State Transfer) is a software architecture that enables interaction between clients and the application over HTTP, making the system easily accessible through web interfaces or other external application [18].

FastAPI is a modern, high performance web framework for building APIs with Python, based on standard Python type hints. It is known for its speed with performance on par with NodeJS and Go, and offers rapid development cycles- claiming up to 200–300% faster feature implementation compared to traditional frameworks [17]. In addition, it also provides automatic input validation, asynchronous support, and interactive documentation, making it a robust and scalable choice for deploying AI systems.

This FastAPI application is containerized using Docker to ensure a consistent runtime environment across local development, testing, and production. This process also helps simplify dependency management and allows the entire RAG application to be deployed as a portable unit [12].

The Docker container is deployed on AWS Lambda using it as a serverless model, allowing it to scale automatically based on incoming requests without the need for a dedicated server [15]. This architecture significantly reduces operational cost and improves scalability. Based on benchmark tests, the cost of running the application is estimated to be less than \$2 per 1,000 queries, mak-

ing it an affordable solution for educational, research, or large-scale production use [16].

DynamoDB serves as a back-end NoSQL database, keeping a record of user queries along with their responses. This allows easy access for future references, audits, or system performance checks. To handle long-running tasks and prevent latency in the user-facing API, complex queries are handled in the background by a second asynchronous Lambda function, also known as the Worker function [15]. This setup ensures smooth performance, even under varying computational loads.

Together, these cloud components form a robust production-ready pipeline that connects users to Solaris AI, enabling them to ask questions and receive intelligent data-backed insights on solar power generation.

Figure 9 illustrates the schematic of the system architecture for the RAG/AI application deployed on AWS. Figure 10 displays the query list records stored in the AWS DynamoDB backend.

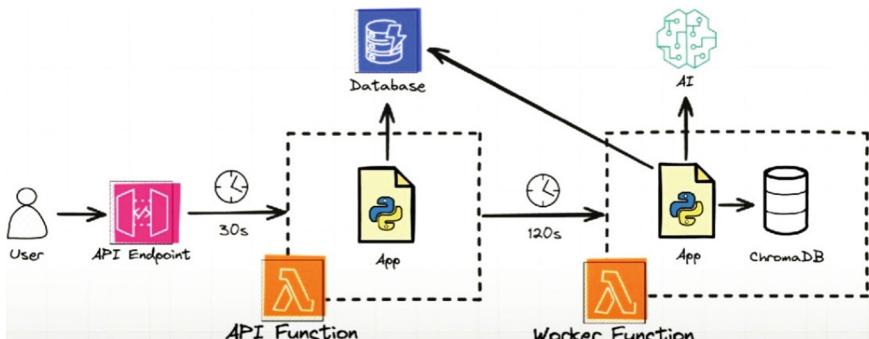


Fig. 9. System architecture for the RAG/AI application deployed on AWS. Taken from [12]

Table: RagCdkInfraStack-RagQueryTableB04F0F4B-124CA96UQ7ZIC - Items returned (10)

Scan started on June 23, 2025, 19:01:16		
< 1 ... >		
query_i...	answer_text	create_... is_com... query_text
<input type="checkbox"/>	e235936f...	To calculate the total solar pow...
<input type="checkbox"/>	d3956a53...	According to the context provid...
<input type="checkbox"/>	9d22e15a...	To calculate the total solar outp...

Fig. 10. Backend AWS DynamoDB Query Lists

4 Results and Discussion

4.1 ANN Model Performance

The Artificial Neural Network (ANN) model was trained on historical meteorological data, including temperature, irradiance, humidity, and atmospheric pressure, and produced promising results. The model achieved a Root Mean Square Error (RMSE) of approximately 16.11 and a Mean Squared Error (MSE) of 259.68. Furthermore, the R^2 score was 0.917 on the training set and 0.778 on the test set, indicating a strong ability to generalize without significant overfitting.

These results show that the ANN can capture the nonlinear relations between the input variables, and it can deliver a reliable solar power output. However, the dataset the ANN model was trained on lacked key features such as GHI and GTI. Lastly, ANN does not account for temporal dependencies as seen in solar irradiance patterns. LSTM is better suited for sequential data. It also aligns more closely with the physical characteristics of solar generation. Hence, we shift our attention from training our ANN model to developing our LSTM model, which is better suited for these tasks.

4.2 Long-Short Term Model Performance and its Comparison with Linear Regression

The LSTM model was trained on a richer meteorological dataset, as it contains GHI and GTI, and compared with the Linear Regression model, which was also trained on the same dataset. The Linear Regression model serves as a baseline performance for comparison purposes. The results of the evaluation are shown in Table 1.

Table 1. Comparison of Evaluation Metrics for Linear Regression and LSTM Models in Solar Energy Forecasting

Algorithm	Lag	Dropout	Epochs	Batch Size	MSE	RMSE (Power Gen)	R^2
LSTM	672	0.33	10	12	130.93	11.44	0.653
LSTM	24	0.30	10	12	134.53	11.60	0.643
Linear Regression	N/A	N/A	N/A	N/A	159.44	12.63	0.570

From the comparison, we can see that LSTM significantly outperforms the linear regression baseline in every single metric. From the LSTM configurations, a lag of 672 shows better performance compared to a lag of 24. This indicates that longer historical patterns play an important role in solar energy forecasting. Moreover, dropout values around 0.3 seem to be effective in preventing overfitting and also effective in maintaining generalization with no signs of major performance trade-offs. Lastly, LSTM's higher R^2 indicates the model explains more variance in the data which is an inherent nature of weather and solar energy.

However, Table 2 shows the comparison between the performance of ANN and LSTM.

Table 2. Comparison of Evaluation Metrics for ANN Model and LSTM (Lag 672) Model in Solar Energy Forecasting

Algorithm	MSE	RMSE	R^2
LSTM (lag 672)	130.93	11.44	0.653
ANN	259.68	16.11	0.778

Despite LSTM having a marginally lower R^2 than the ANN model, the ANN model did include critical features such as GHI and GTI. The LSTM model, however, was trained on a richer more physically meaningful feature set. Moreover, the LSTM model significantly outperformed the ANN model in RMSE and MSE, indicating better real world forecasting accuracy and consistency.

4.3 Generative AI with RAG Integration

The Retrieval-Augmented Generation (RAG) pipeline was developed using LangChain to integrate OpenAI’s large language model (LLM) with both the LSTM model and the solar dataset. This setup enabled the system to answer natural-language questions such as “Why was solar output low yesterday?” by retrieving relevant data from a ChromaDB vector database and combining it with the model predictions.

This architecture improved explainability and factual accuracy by confining the response to the data instead of on the model knowledge alone. This reduces hallucination and creates more transparency and making it more reliable and convenient for the user.

4.4 Cloud Deployment and System Scalability

The application pipeline was deployed using FastAPI for seamless API interactions and Docker for reliable containerization. It runs serverlessly on AWS Lambda, enabling automatic scaling without requiring back-end management. DynamoDB stores query-response data, while asynchronous Lambda functions ensure fast and efficient processing, even under heavy traffic.

This architecture ensures that Solaris AI is not only technically robust but also scalable, modular, and cost-effective. The estimated cost via Amazon Cost Calculator is less than \$2 per 1,000 queries, which is very affordable.

5 Conclusion and Future Work

Accurate prediction of solar power is vital for integrating renewable energy into grids and achieving a low-carbon future. Forecasting models serve as a crucial

tool in stabilizing grid operations, optimizing energy resource planning, and supporting data-driven policy decisions. In this research, we developed Solaris AI, a solar forecasting system that initially leveraged an Artificial Neural Network (ANN) model and was later enhanced with a Long Short-Term Memory (LSTM) model to improve time-series prediction accuracy. The system combines deep learning-based forecasting with the explainability and accessibility of a Retrieval-Augmented Generation (RAG) pipeline. Our approach integrates a deep learning forecasting model with OpenAI's language models via LangChain, allowing users to interact with both datasets and real-time ANN predictions through a conversational interface. The system is deployed using FastAPI, Docker, and AWS Lambda for scalable, serverless access, and includes DynamoDB for persistent query logging.

Our final output is a chatbot-like interface that enables users to ask intuitive, natural-language questions such as "How much power was generated this week?" or "Why was output low yesterday?" and receive context-aware, data-backed answers. This makes Solaris AI not only a forecasting tool but also a transparent, multilingual assistant for solar energy systems.

Future work includes integrating real-time data from sensors or weather APIs to improve the accuracy and responsiveness of predictions. This would allow Solaris AI to adapt to environmental changes and provide live forecasting capabilities. Further enhancements could include support for voice-based queries, making interactions even more natural and convenient. The system could also expand to cover other renewable energy sources, like wind and hybrid platforms. In addition, the use of video-based generative AI could enable automatic generation of visual summaries or forecast explanations, making the system more accessible and intuitive for a wider range of users.

References

1. Rajasundrapandian, T., Kumaresan, K., Murugan, S., Subathra, M.S.P., Sivakumar, M.: Solar energy forecasting using machine learning and deep learning techniques. *Int. J. Num. Meth. Eng.* (2023). <https://www.researchgate.net/publication/378395432>. Accessed 15 May 2025
2. EIA projects 48% increase in world energy consumption by 2040 (2016). <https://www.eia.gov/todayinenergy/detail.php?id=26212>. Accessed 27 Mar 2025
3. Erduman, A.: A smart short-term solar power output prediction by artificial neural network. *Electr. Eng.* **102**, 1441–1449 (2020). <https://doi.org/10.1007/s00202-020-00971-2>
4. Salman, D., Kusaf, M., Direkoglu, C.: Hybrid deep learning models for time series forecasting of solar power. *Neural Comput. Appl.* (2024). <https://doi.org/10.1007/s00521-024-09558-5>
5. Essam, Y., et al.: Investigating photovoltaic solar power output forecasting using machine learning algorithms. Informa UK Limited, Taylor and Francis Group (2022). <https://doi.org/10.1080/19942060.2022.2126528>
6. Qiu, Z., et al.: EF-LLM: energy forecasting LLM with AI-assisted automation, enhanced sparse prediction, hallucination detection. *arXiv preprint*

arXiv:2411.00852v2 (2024). <https://arxiv.org/pdf/2411.00852v2.pdf>. Accessed 15 May 2025

7. Gupta, A.: Forecasting the power generated by a solar plant. GitHub Repository (2023). <https://github.com/anantgupta129/Solar-Power-Generation-Forecasting>. Accessed 15 May 2025
8. Viraat, I.: Short-term power forecasting. Github repository (2020). <https://github.com/shubhamchouksey/Power-Prediction-LSTM/tree/master/Data>. Accessed 01 Jun 2025
9. Amazon Web Services: What is LangChain? (2025). <https://aws.amazon.com/what-is/langchain/>. Accessed 15 May 2025
10. LangChain: LangChain conceptual guide (2025). <https://github.com/langchain-ai/langchain/blob/master/docs/docs/concepts>. Accessed 15 May 2025
11. LangChain: Text splitters, LangChain (2025). https://python.langchain.com/docs/concepts/text_splitters/. Accessed 15 May 2025
12. Pixegami: Deploy RAG/AI app to AWS (2024). <https://github.com/pixegami/deploy-rag-to-aws>. Accessed 15 May 2025
13. Chroma: Chroma documentation: Introduction (2025). <https://docs.trychroma.com/docs/overview/introduction>. Accessed 15 Apr 2025
14. Garza, E.: Hands-on large language models with LangChain – Part 2: Building a robust document search engine (2023). <https://ernestodotnet.medium.com/part-2-hands-on-large-language-models-with-langchain-building-a-robust-document-search-engine-95396195326>. Accessed 15 May 2025
15. Amazon Web Services: AWS documentation (2025). <https://docs.aws.amazon.com>. Accessed 21 Apr 2025
16. Amazon Web Services: AWS pricing calculator (2025). <https://aws.amazon.com/pricing/>. Accessed 29 Apr 2025
17. Ramírez, S.: FastAPI documentation (2025). <https://fastapi.tiangolo.com/>. Accessed 02 May 2025
18. Amazon Web Services: What is a RESTful API? (2025). <https://aws.amazon.com/what-is/restful-api/>. Accessed 02 May 2025



A Tool for Synthesizing and Implementing Medium Voltage Load Profiles

Hendrik Plompen^{1,2} , Ranier Alexander Arruda Moura² , Khawaja Khalid Mehmood² , Anne van der Molen² , Peter van der Wielen² , and Phuong Hong Nguyen²

¹ TenneT, 6812 Arnhem, AR, The Netherlands

² Eindhoven University of Technology, 5612 Eindhoven, AZ, The Netherlands

r.a.a.m.arruda.moura@tue.nl

Abstract. To support energy planning and grid investigations, this paper presents an open source tool to generate and integrate realistic synthetic load profiles (SLPs) into power system software. The tool assigns daily SLPs to all loads within a network and is compatible with Vision Network Analysis and OpenDSS. It features a Graphical User Interface (GUI) and employs the probabilistic Multivariate Elliptical Copula (MEC) method, using active power consumption data as input. Additional functionalities include preprocessing, clustering, and validation. Preprocessing handles outlier and duplicate removal via the interquartile range (IQR) and imputes missing values using linear interpolation or k-Nearest Neighbors (kNN). Clustering groups data by month and day type, and later by K-Means to identify different consumer types. A selected cluster serves as input to the MEC model, which can be conditioned on annual consumption, mean load, and peak load. Validation metrics include Mean Absolute Percentage Error (MAPE), Root Mean Square Error (RMSE), and Jensen-Shannon Divergence (JS-D). In a case study using real medium voltage (MV) load data from a Dutch city, the tool achieved an average MAPE of 0.56% for industrial and 0.51% for aggregated residential consumers. It also successfully assigned SLPs, performed load flow calculations, and extracted congestion-related results in both software environments.

Keywords: Load profile · Synthetic data generation · Power system software · Congestion · Preprocessing

1 Introduction

The electricity grid has been one of the primary energy carriers in society for many years, and its importance is expected to grow even further in the future [14]. This grid supplies power to a diverse range of users, including households, businesses, and industries. Grid operators are responsible for ensuring that the network remains safe, reliable, stable, and efficient at all times. This is achieved

through careful planning, development, design, and management of the entire grid infrastructure.

However, challenges are arising due to the limited capacity of the existing grid. Operators are increasingly facing congestion issues multiple times a day, caused by high electricity demand and the penetration of renewable energy sources (RES). The variable and decentralized nature of them adds complexity to grid management and poses significant obstacles for network operation [19].

One way to support energy planners and grid operators is by generating SLPs. These synthetic yet reliable profiles aim to replicate the statistical properties and correlation structures of real-world load profiles (LPs), making them useful for power system analysis and decision-making [1]. In other words, they can help indicate whether different scenarios may cause congestion in a specific grid topology.

1.1 Related Works

In [8], a top-down probabilistic approach based on the Multivariate Elliptical Copula (MEC) is proposed to generate synthetic residential data. This method is capable of capturing the statistical properties of smart meter measurement datasets, considering the autocorrelation of their correspondent time series (TS) and their probability density functions (PDFs) across different seasons. In addition, conditioned variables are incorporated into the modeling process, such as annual energy consumption and daily weather profiles.

In [1], the MEC technique without conditioning was also applied to medium-voltage (MV) loads, focusing on industrial and aggregated residential profiles collected at the transformer level in a Dutch city. The synthesized scenarios were evaluated considering analyses of Spearman's rank correlation coefficients, average power behaviors, and histogram distributions. Despite this, numerical issues emerged during the simulations due to discrepancies between the PDFs. A similar work was done in [2], but specifically for transformer loads.

Since the synchronization of residential consumers can potentially lead to overloads, Gaussian Copula (GC) modeling was employed in [11] to capture this synchronization by analyzing the time-varying dependence structure between individual demands, specifically using the correlation parameter as an indicator. A GC was also used in [9] to model uncertainties in home energy management, including load parameters, outdoor temperatures, and energy prices, enabling the generation of correlated scenarios.

In [12] a SLP generator tool is combined with the power system software OpenDSS to automatically assign profiles to network components. This was done to investigate the effects that the profiles have on a modeled network. However, it only involves SLP generation for electric vehicles (EV).

1.2 Contribution

This research improves the use of SLPs in distribution network operations in two ways: by enhancing synthesizing techniques and improving their implementation

in power system software. A tool was developed to generate and assign daily SLPs to network loads in Vision Network Analysis and OpenDSS. It uses the top-down MEC method from [8] and includes preprocessing, clustering, and validation features. A case study tested the functionality and performance of the tool. Furthermore, the contributions of this work are:

- Creation of an open-source tool that is able to combine both generation and implementation as one
- The ability of the tool to implement LPs into a network within various power system software
- The possibility to only use the generation part of the tool, which compared to other open-source stand-alone generators is able to generate realistic SLPs based solely on smart meter data.

The remainder of this paper is organized as follows: In Sect. 2, the methodology employed is presented. Section 3 discusses how the complete tool is operated. The case study is presented in Sect. 4, with the obtained results presented in Sect. 5. These results are discussed in Sect. 6 and Sect. 7 will conclude the investigation. A step by step tool guide is also provided in Appendix A.

2 Methods

The following subsections discuss the various methods applied throughout the entire process of the tool. A complete overview of this process and its steps is further discussed in detail in Sect. 3.

2.1 Preprocessing

The first step in generating SLPs from historical data is preprocessing, which handles issues such as missing values, duplicates, and outliers. This study applies either linear interpolation or the kNN algorithm for imputing missing data, and uses the IQR method for detecting outliers.

Linear interpolation estimates missing values by assuming a linear relationship between neighboring data points within the same TS, while the kNN algorithm is based on similarities of neighboring values [20]. This similarity makes use of the distance between the missing and the neighboring values and for this the Euclidian distance is used according to:

$$d(x, y) = \sqrt{\sum_{i=1}^n (x_i - y_i)^2} \quad (1)$$

where n is the number of dimensions of the data, x the missing value, y the neighbor and using $i = 1, \dots, n$. Depending on the number of neighbors selected, it finds this number of neighbors which have the smallest distance. The missing

value is then imputed based on the weighted mean of these neighbors as used in [21, 27] following

$$x_i = \frac{\sum_{k=1}^K w_k v_k}{\sum_{k=1}^K w_k} \quad (2)$$

where K is the number of neighbors selected, v_k is the nearest neighbor, $k = 1, \dots, K$ and w_k is the weight factor that can be written as

$$w_k = \frac{1}{d(x, y)}. \quad (3)$$

The chosen value for K influences the imputed value the most. Various studies discuss possible optimal values for K , such as $K = 1$ [15], $K = 2$ [3] or even $K \approx \sqrt{N}$ where N is the total number of neighbors for that measurement [7]. For the tool the value of K can be set by the user.

Outliers are particularly high or low values compared to the rest of the dataset that can be caused by an error. However, they can be just a correct measurement, and thus are not required to be removed. The detection of these outliers follows the IQR method [10] as cited in [25]. This method considers x to be an outlier if

$$x < (Q1 - 1.5 \cdot IQR) \quad (4)$$

or

$$x > (Q3 + 1.5 \cdot IQR) \quad (5)$$

where $Q3$ is the third quartile or 75th percentile, $Q1$ the first quartile or 25th percentile, the factor f is set to 1.5, and IQR is calculated using

$$IQR = Q3 - Q1. \quad (6)$$

The IQR uses a breakdown point of 25% for the range of all values. However this can be set to a different point by the user. When outliers are removed they are then treated as missing values and imputed following one of the methods previously discussed.

2.2 Clustering

Weather variables such as irradiance and temperature affect residential electricity demand, creating different LPs between seasons. Since Dutch weather often varies, clustering is done by month rather than by season. Data are also split into weekdays and weekends, and further clustered to group similar load patterns. This helps identify consumer types and enables targeted modeling and simulations.

The final clustering process follows the K-means algorithm [18]. This method focuses on minimizing the distance between all data points in a cluster and the center point in that cluster, called the centroid. These data points are then assigned to the cluster where this distance to the centroid is the smallest. By first initializing the centroids, the distance for all data points can be calculated.

Then the centroids are recomputed, and the distances are recalculated for these new centroids. This continues until the centroids stop changing. The objective function that the K-means algorithm tries to optimize can be defined as

$$E = \sum_{i=1}^c \sum_{x \in C_i} \|x - u_i\|^2 \quad (7)$$

where c is the number of different clusters, and u_i the centroid of cluster C_i .

To find the optimal number of clusters, many researchers use internal validation criteria which are able to validate the effectiveness of clusters, the compactness of each cluster and the distance between each cluster [5]. In this research the Davies-Bouldin index (DBI) is used which compares each different cluster with another [13]. From this comparison an index value follows and the lower this index value the less similar the clusters are to each other. Computing the DBI for different numbers of clusters results in finding the optimal number of clusters. The DBI can be written as

$$DBI = \frac{1}{N} \sum_{i=1}^N \max_{i \neq j} \frac{S_i + S_j}{M_{i,j}} \quad (8)$$

where N is the number of clusters, S_i and S_j are the dispersion's of clusters i and j respectively and $M_{i,j}$ the distance between vectors which are chosen as characteristic of clusters i and j [6].

2.3 Multivariate Elliptical Copulas

A copula is a multivariate distribution function, where its marginal distribution functions are uniform in the interval $(0, 1)$ [22]. This function was formulated according to Sklar's Theorem [26], where a multivariate joint distribution $F(x_1, \dots, x_d)$ is represented using its marginal distribution functions $F_i(x_i)$ and a copula $C(\cdot)$ for $i = 1, \dots, d$ is illustrated below:

$$F(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d)) \quad (9)$$

Since this function models the correlation patterns between the marginal uniform random variables $U_i = F_i(X_i)$, the previous equation is rewritten as it follows:

$$F(x_1, \dots, x_d) = C(u_1, \dots, u_d) \quad (10)$$

The MEC model proposed in [8] uses a set of elliptical copulas consisting of both the Multivariate Gaussian Copula (MGC) and the Multivariate t-Distribution Copula (MTC), which capture the dependency structure between random variables. Each random variable has a different marginal distribution function. Furthermore, the copulas can assign different dependence values to all pairs of random variables. Both copulas are used to model a LP and using the BIC are compared to find the model that better describes the input dataset.

2.4 Validation

To validate the generated SLP dataset, multiple metrics are employed that compare it with the original smart meter input dataset. To give insight in the mean error both the RMSE and MAPE are used. However, RMSE is generally considered to be more sensitive to outliers than MAPE [4]. The RMSE is defined as

$$RMSE = \sqrt{\frac{\sum_{t=1}^T (y_t - \hat{y}_t)^2}{T}} \quad (11)$$

where y_t is the measured active power consumption and \hat{y}_t is the synthetically generated active power consumption. Based on the same variables the MAPE can be written as

$$MAPE = 100 \frac{1}{T} \sum_{t=1}^T \left| \frac{y_t - \hat{y}_t}{y_t} \right| \quad (12)$$

To give insight in the statistical parameters the Jensen-Shannon divergence (JS-D) is used, which compares the probability distributions and is based on the Kullback-Leibler divergence (KL-D) but will always have a finite value [17,23]. It can be written as

$$D_{JS}(P||Q) = \frac{D_{KL}(P||\frac{P+Q}{2}) + D_{KL}(Q||\frac{P+Q}{2})}{2}. \quad (13)$$

where D_{KL} is the KL-D which is defined as

$$D_{KL}(P||Q) = \sum_{x \in X} P(x) \log \left(\frac{P(x)}{Q(x)} \right) \quad (14)$$

where $P(x)$ is the probability distribution of the measured data and $Q(x)$ is the probability distribution of the synthetically generated data [16].

3 Tool Design

The tool, whose code is available in [24], is a Python 3.11.4 script with a GUI, compatible with Vision Network Analysis and OpenDSS. Its main purpose is to generate and implement SLPs, with additional features such as preprocessing, clustering, and validation. Each feature can run independently, except for generation, which requires prior clustering and is followed by validation to ensure quality. All features, except validation, can handle two inputs at once.

In Fig. 1 the core algorithm can be seen where each row represents a certain environment. From top to bottom, these are: the user, Python, Excel, and one of the power system analysis software previously mentioned.

To begin the process, an initialization is started in which the features required by the user are selected. When the implementation feature is selected, the tool also requires which software is used for implementation and some network parameters. These parameters consist of a list of loads including both their name and

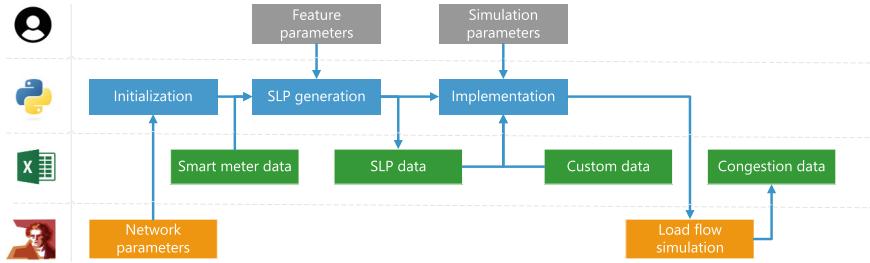


Fig. 1. Flowchart of core algorithm

type (normal or transformer load). Based on this list, the tool knows how many SLPs to generate. When multiple types are present, it is possible to assign different SLPs to each type.

Next, we have the SLP generation process that consists of preprocessing, clustering, generation, and validation. For this step, an Excel file is required which contains the active power consumption measurements. In addition, some feature parameters are required from the user.

When this process is successful, an SLP dataset should be generated, which acts as the input for the implementation process. Instead, it is also possible to use a custom dataset for this part.

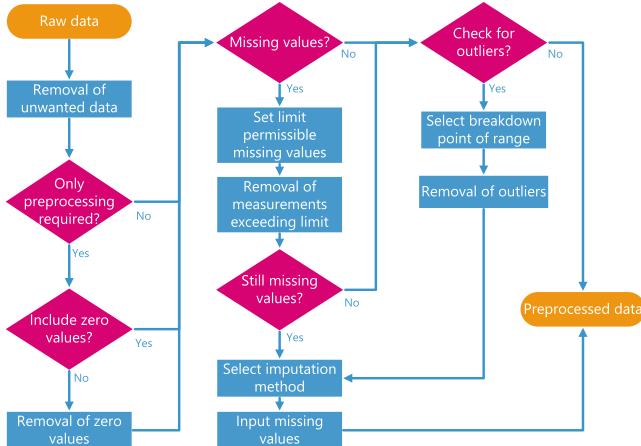


Fig. 2. Flowchart of preprocessing stage

Figure 2 outlines the preprocessing steps applied to the raw measurement data. First, empty entries and duplicate timestamps are removed. If preprocessing is selected, zero-only measurements can be included or removed. A threshold for missing data (as a percentage of the total) is set, beyond which measurements

are discarded; the rest are imputed using a method from Sect. 2.1. Outliers can then be detected, removed, and imputed again. The final data can be visualized as a daily load profile with standard deviation or a histogram with kernel density estimation, and exported to CSV or Excel.

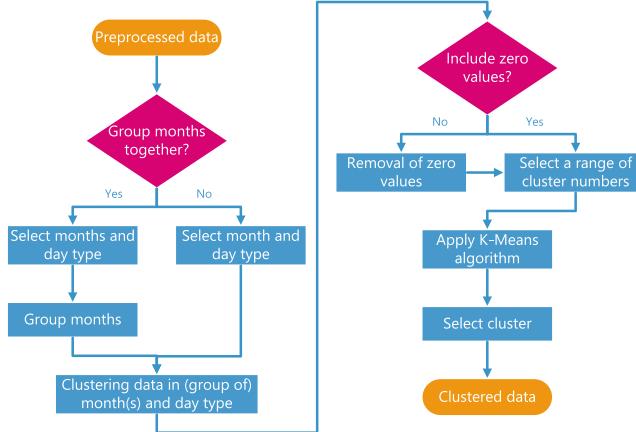


Fig. 3. Flowchart of clustering stage

In Fig. 3 the steps for the clustering process are presented. Clustering is done per month and day type (weekday or weekend), with the extra option to group months together. When these clustering steps are performed, the following optional step is the removal of zero measurements. The next clustering step is applying the K-means algorithm and is done by defining a minimum and maximum number of clusters. This minimum should be at least 2 to work properly. The tool then tries to cluster the data in this range of number of clusters set by the user. When the algorithm is successful, it results in a list containing the optimal number of clusters. For each of these clusters, the DBI and the number of measurements are presented. The final clustering step is to choose one of these clusters to be used for generation. Like the preprocessing feature, visualization and the possibility to export data is provided.

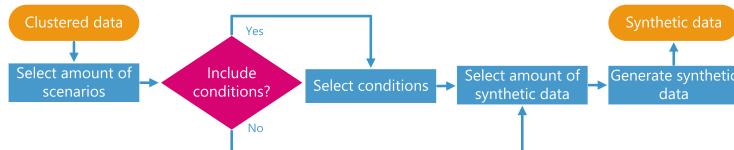


Fig. 4. Flowchart of synthesizing stage

Figure 4 depicts the steps taken in the generation process. Only in this feature it is possible to generate two separate scenarios from the same input data. This provides the option to perceive the influence the conditioned MEC model can have on the data, by comparing a conditioned scenario with an unconditioned scenario. The different conditions that can be selected are annual consumption, mean load, and peak load. For each condition a range is provided, based on the minimum and maximum values of this condition coming from the input data, and from this a value can be set for it by the user. The last step before generation is to set the desired number of generated SLPs. The MEC model is incorporated into the script by using the EllipticalCopula module from the multicopula package developed in [8].

If the generation is successful, the tool will provide the success rate (SR), MAPE, RMSE and JS-D for the generated SLPs. This SR comes from the fact that the EllipticalCopula module sometimes generates negative infinite values at certain time indices. It is calculated using

$$SR = 100 - \left(\left(\frac{D_{-\infty}}{D_{tot}} \right) \cdot 100 \right) \quad (15)$$

where $D_{-\infty}$ is the number of negative infinite values in the synthetic dataset and D_{tot} is the total number of data points in the synthetic dataset. All SLPs that contain exclusively these infinite negative values are subsequently removed from the generated dataset. Any remaining negative infinite values are imputed using linear interpolation as previously discussed in Sect. 2.1. Section 6 further discusses why such values might be generated. Finally, there is the option to go back and choose a different cluster, in case the resulting SR and validation are not satisfactory. However, when the generated data are deemed sufficient, the generation process is concluded by providing visualization where the resulting SLPs can be compared with the cluster input LPs. Both plots showing the mean value or the peak value at each time instance can be compared.

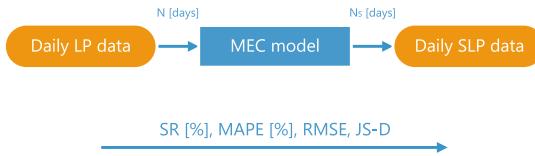


Fig. 5. Flowchart of validating stage

Both MAPE and RMSE are applied to the mean daily LP and the mean daily SLP, as seen in Fig. 5. For the JS-D, the probability distribution of both the input dataset and the generated synthetic dataset are required. Furthermore, they should be for the same random variable; however, this is not the case when dealing with two different datasets containing a wide variety of differing values. By constructing a histogram for both datasets that enclose the same range and

use the same number of bins, both distributions can still be compared with each other. In this way, the number of values in a certain bin, and thus the probability, can still be compared between distributions.

For the generated SLPs to be implemented in Vision or OpenDSS, a number of steps are required. First, from the network file of the network chosen for implementation a list of all the loads needs to be extracted. This list should contain all the exact names for each load and when using Vision it should also contain which type of load it is, either normal or transformer load. When a network consists of two types of load in Vision, it is possible to assign different SLPs to each type by making use of the option to process two inputs simultaneously throughout the tools algorithm. Based on the list of loads, the tool knows how many SLPs are required and will sample these randomly from the generated SLP dataset. Each SLP is then assigned to one of the loads, and when this is done, a load flow simulation is performed. As a final step, the congestion results are extracted from the load flow results which entail the load rate for all time steps for each line, cable, or transformer that has a load rate above 100% for at least one of the time steps. More details of the tool can be found in Appendix A.

4 Case Study

To assess the impact of user defined parameters, multiple case studies were conducted using a base case for comparison. This approach helps to analyze how each parameter influences modeling and validation. Adjustable settings include the imputation method, number of neighbors for kNN, outlier breakdown point (OBP), clustering configuration, treatment of zero measurements, number of generated SLPs, condition modeling, and the IQR factor f . Section 6 provides recommendations based on these results.

The data used in this work comes from an industrial area in the Dutch city of Dordrecht. It contains a large number of logistics centers, but also numerous residential buildings. In total, active power measurements were obtained for 99 industrial and 96 residential consumers, recorded over an entire year with one hour resolution. Both datasets were used as input for the tool. The network of this area was also used for the implementation process and subsequent load flow simulation. This was done to see the effects conditioning can have on a network and on congestion.

For the base case, the parameter inputs were set as follows: no imputation method was selected (as both datasets do not contain missing values), outlier detection was not utilized, zero measurements were not removed, no K-means clustering was applied, the desired number of SLPs was set to the same number of LPs used as input for the generation, and no conditioning was applied. For all cases, the modeling was done for all four seasons separately and for both weekdays and weekends, meaning clustering by season and day type.

All parameters used in the tool are part of the preprocessing, clustering, or generation process. For preprocessing, these include the OBP, the IQR factor f , and the imputation method, which varied between linear interpolation and

kNN with different k values. Clustering parameters include whether zero measurements were included and the range of clusters tested. The generation process involved setting the desired number of SLPs. For each case study, one of these parameters was modified to test its impact compared to the base case. OBP values of 1%, 10%, and the default 25% were examined. The IQR factor f was tested at values 1.5, 3, and 10. Imputation methods included linear interpolation and kNN with K values of 1, 2, 5, 10, and 94, which the latter is based on the \sqrt{N} rule, where $N = 8760$ (24×365). The number of SLPs tested ranged from 100 to 10000 and included a case where it matched the input size. Some clustering cases included zero measurements, and the cluster range tested went from 2 to 20.

The final adjustable parameters in the generation process are the conditioning settings, for which two setpoints per condition were tested: one representing the mean of the valid range and the other representing the maximum. For the AEC condition, the industrial setpoints were 1278.77 GW and 2550.00 GW, and the residential setpoints were 667.09 GW and 1330.00 GW, respectively. For the mean load condition, values of 0.26 GW and 0.50 GW were tested for the industrial case, and 0.10 GW and 0.20 GW for the residential case. Similarly, the peak load condition used 0.49 GW and 0.90 GW for industrial, and 0.17 GW and 0.30 GW for residential. In all scenarios, the same value level, either mean or maximum, was applied to both industrial and residential sectors simultaneously.

5 Results

Figures 6 and 7 show the input datasets: mean LPs and PDFs, respectively. Tables 1 and 2 present base case modeling and validation results for weekdays and weekends. The value N refers to the total number of input days (see Fig. 5). Table 3 reports the number and percentage of outliers detected at each breakdown point. Figure 8 displays the DBI for the tested cluster range, while Fig. 9 compares mean LPs of clustered vs. original inputs. Table 4 summarizes the con-

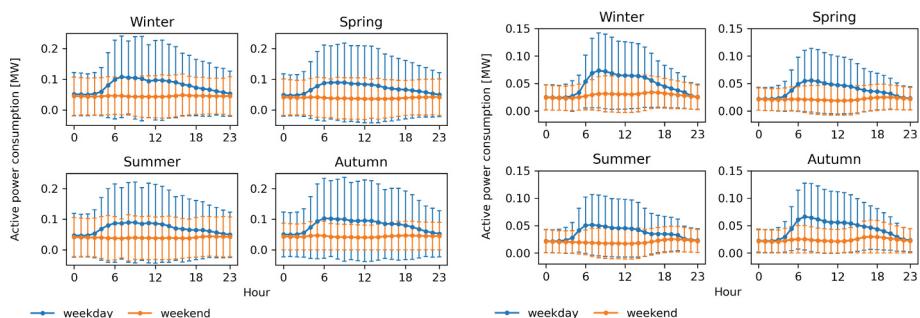


Fig. 6. Mean LPs for both weekdays and weekends for each season of industrial (left) and aggregated residential (right) data, depicting standard deviation as error bar

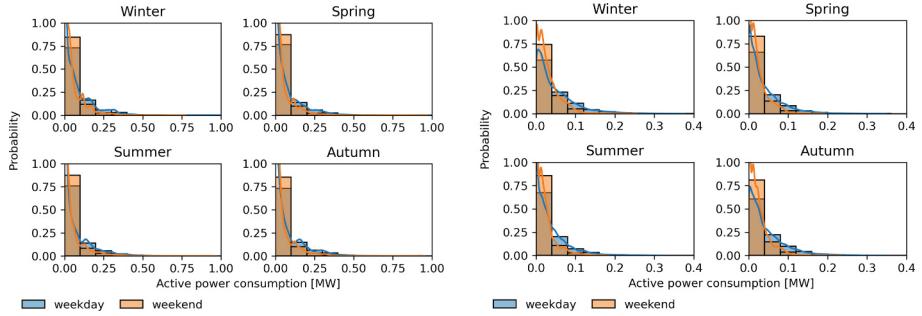


Fig. 7. Histogram for both weekdays and weekends for each season of industrial (left) and aggregated residential (right) data, including the kernel density estimation

Table 1. Base case generation and validation results for both industrial and residential for the entire year regarding weekdays (W = winter, SP = spring, SU = summer, F = fall)

Season	Industrial					Residential				
	N [days]	SR [%]	MAPE [%]	RMSE	JS-D	N [days]	SR [%]	MAPE [%]	RMSE	JS-D
W	6336	81.26	11.21	9.0×10^{-3}	4.3×10^{-3}	6144	98.43	0.88	6.0×10^{-4}	1.0×10^{-4}
SP	6534	78.56	7.42	5.1×10^{-3}	2.8×10^{-3}	6336	93.62	0.70	3.0×10^{-4}	1.0×10^{-4}
SU	6534	77.30	6.00	4.2×10^{-3}	1.5×10^{-3}	6336	92.93	1.20	4.0×10^{-4}	2.0×10^{-4}
F	6435	83.95	5.76	5.0×10^{-3}	1.2×10^{-3}	6240	97.46	0.62	3.0×10^{-4}	1.0×10^{-4}

Table 2. Base case generation and validation results for both industrial and residential for the entire year regarding weekends (W = winter, SP = spring, SU = summer, F = fall)

Season	Industrial					Residential				
	N [days]	SR [%]	MAPE [%]	RMSE	JS-D	N [days]	SR [%]	MAPE [%]	RMSE	JS-D
W	2574	80.99	11.11	5.0×10^{-3}	2.7×10^{-3}	2496	97.49	1.67	5.0×10^{-4}	4.0×10^{-4}
SP	2574	77.22	9.52	3.7×10^{-3}	1.9×10^{-3}	2496	89.69	3.23	7.0×10^{-4}	5.0×10^{-4}
SU	2574	75.75	10.61	4.3×10^{-3}	1.9×10^{-3}	2496	87.26	3.00	7.0×10^{-4}	5.0×10^{-4}
F	2574	81.45	8.92	4.0×10^{-3}	1.3×10^{-3}	2496	94.27	0.60	2.0×10^{-4}	1.0×10^{-4}

Table 3. Preprocessing outlier detection results

OBP	Industrial		Residential	
	Total outliers	Total rate [%]	Total outliers	Total rate [%]
1%	110	1.3×10^{-2}	7	8.0×10^{-4}
10%	4363	0.50	1519	0.18
25%	32393	3.74	20813	2.47
1% with moving window	70	8.0×10^{-3}	11	1.3×10^{-3}
10% with moving window	3061	0.36	775	0.09
25% with moving window	24477	2.91	13910	1.65

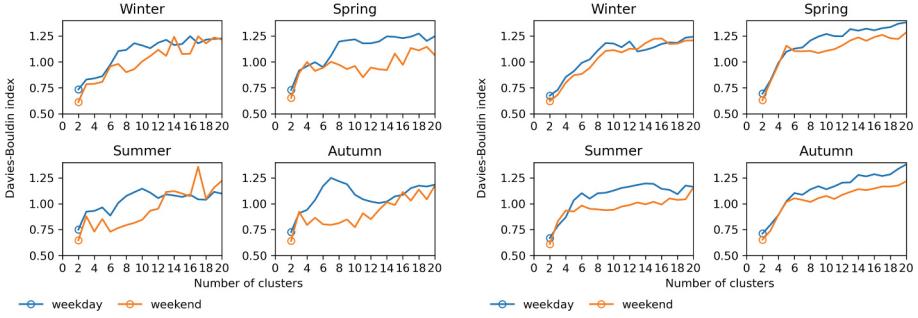


Fig. 8. DBI results of K-means algorithm dividing data in range of number of clusters for industrial (left) and aggregated residential (right) LPs

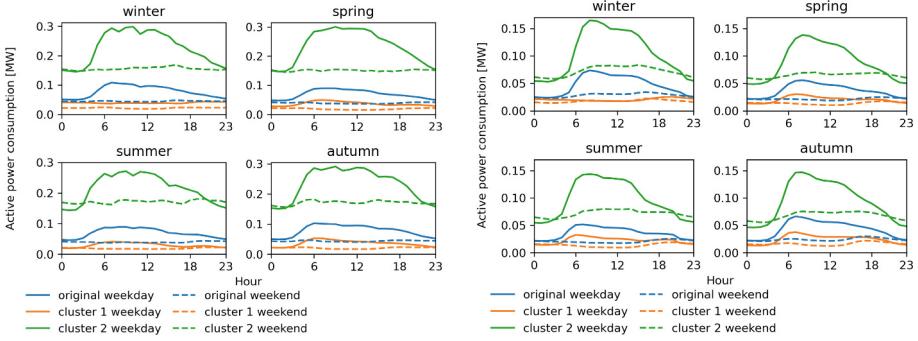


Fig. 9. Original and clusters mean LPs for both weekdays and weekends for each season for industrial (left) and aggregated residential (right) LPs

Table 4. Congestion results

	Cables	Load rate [%]	Transformers	Load rate [%]
Base case	0	—	2	120.00
AEC lower value	2	111.00	3	135.00
AEC higher value	24	223.00	11	155.00
Mean load lower value	6	142.00	4	111.00
Mean load higher value	52	380.00	14	257.00
Peak load lower value	8	158.00	7	123.00
Peak load higher value	35	269.00	13	197.00

gestion results, with overloaded cables and transformers and their maximum load rates. A deeper interpretation of these results is presented in Sect. 6.

6 Discussion

Figure 6 shows that in some time steps, the error bars drop below zero, indicating high variability, especially in industrial data. Figure 7 reveals that most of the data points fall within a low active power range, close to the mean values. Despite this, a large standard deviation is observed, likely due to the mix of many small consumers and a few large ones in the dataset.

Considering that the goal is to maximize SR (thereby generating more data) and minimize MAPE, RMSE, and JS-D (to ensure that the generated data are more statistically similar to the original), it can be deduced from Tables 1 and 2 that the generated SLPs for the residential dataset resemble the input LPs more closely than those for the industrial dataset. A possible reason is that residential data have a standard deviation closer to the mean compared to industrial data, meaning that data with fewer variations may be easier to model. This is reflected in the SR values, which are consistently higher for the residential dataset across all seasons and day types, as the MEC model tends to produce negative infinite values more frequently when the distributions of the profiles differ too greatly. Additionally, a lower (or closer to 0) JS-D further confirms a greater similarity between the synthetic and original PDFs in the residential data.

Data from the fall season were also observed to be slightly easier to replicate, as synthetic scenarios from this season achieved the lowest MAPE among all (Tables 1 and 2). This trend is further supported by RMSE values, which are also the lowest in most fall cases, along with SR values, which are the highest in many instances. Similarly, JS-D values are lowest during the same period, indicating a strong match in the statistical distribution of the data. These results suggest not only high numerical accuracy in the generated SLPs, but also strong statistical fidelity, due to the closer alignment of the respective PDFs.

Since the idea is to recommend parameter settings based on the metrics mentioned above, the tests showed that the best preprocessing settings involve using the OBP of 1% (Table 3) and imputing missing values with kNN, where $k = 2$ (Fig. 8). Zero measurements should be removed, as they offer little insight into load behavior individually. However, their complete removal can be debated, as they may reflect realistic load patterns within a larger dataset.

The lowest DBI value was consistently achieved by K-means with two clusters across all seasons and day types for both datasets (Fig. 8). For industrial data, only the second cluster improved all validation metrics. In the residential case, it reduced MAPE but increased RMSE and JS-D, probably due to RMSE's sensitivity to outliers and their impact on probability distributions, affecting JS-D. Although selecting the second cluster can improve accuracy, it limited the amount of data used. Figure 9 shows that the second cluster more closely resembles the original LP shape. In addition, using more SLPs than input LPs improves accuracy, with 10000 SLPs yielding the best results.

In Table 4 the AEC condition has the least overloaded cables and transformers compared to the others. This might be due to the mean and peak conditions depending only on the daily input LPs, whereas the AEC condition depends on all daily LPs for the whole year. This contributes to the mean and peak condi-

tions having more influence on the modeling compared to the AEC condition. The higher set value for the mean load also has more influence than the higher set value for the peak load. This is to be expected, as when conditioning, the mean encompasses all time steps, while the peak only encompasses a few.

Note that this tool can be applied to larger and more complex networks from different locations, as long as the associated load data from these topologies are also available. The tool is also expected to perform well with low voltage scenarios instead of MV ones. However, there are some limitations. Currently, the tool cannot handle more than two cases simultaneously. Additionally, the occurrence of negative infinite values can become a more significant issue if the LPs differ too much from each other. This issue could be mitigated by further clustering to ensure that only similar LPs are used as input. Furthermore, the MEC model may become computationally intensive when generating large volumes of SLPs, which poses a challenge to scalability. To address this, dimensionality reduction algorithms could be employed to select only the most relevant features as input.

7 Conclusion

The main goal of this research was to improve the utilization of SLPs in distribution network operations, with two specific sub-goals: improving the use of probabilistic generation techniques and enhancing the implementation of SLPs in power system software.

The first sub-goal was achieved through the development of an open-source tool capable of generating SLPs using the MEC approach, based solely on active power measurements as input. Additionally, the inclusion of preprocessing and clustering significantly improved the accuracy of the generated SLPs, achieving an average MAPE of 0.56% for industrial consumers and 0.51% for aggregated residential consumers over an entire year.

The second sub-goal, enhancing the use of SLPs in power system software, was met by enabling the implementation of generated SLPs in both Vision and OpenDSS. Moreover, load flow calculations were performed using these SLPs, and the resulting congestion metrics were also extracted.

Acknowledgments. The authors would like to acknowledge that this work is financially supported by RVO, project no. MOOI622001.

A Appendix: Step by Step Guide

A.1 Preface

For the tool to operate properly, all files used should be inside the same folder. These include: provided Python tool modules, provided Vision macro file (vmf), input dataset(s), and Vision network file (vnf) or OpenDSS network file (dss file).

The tool makes use of the following packages, which should be previously installed: calendar, matplotlib, multicopula, numpy, pillow, pandas, scikit-learn, scipy, seaborn and tk.

To start the tool, run the master module. When implementation in Vision is required, the tool should instead be started up within Vision in the macro editor by opening and executing the vmf named *vision_implementation* (Fig. 10).



Fig. 10. Start screen

Press start to initiate the tool, and exit to terminate it.

A.2 Feature selection

Step 1: Select one or more features you wish to employ (Fig. 11).

Keep in mind that not all combinations of features work, e.g. preprocessing and validation cannot be combined. Furthermore, when selecting MEC generation, clustering and validation are also automatically selected.

Step 2: When the selection is performed correctly, the continue button at the bottom is enabled and should be pressed.

The following appendix subsections show the order of screens when all features are selected.

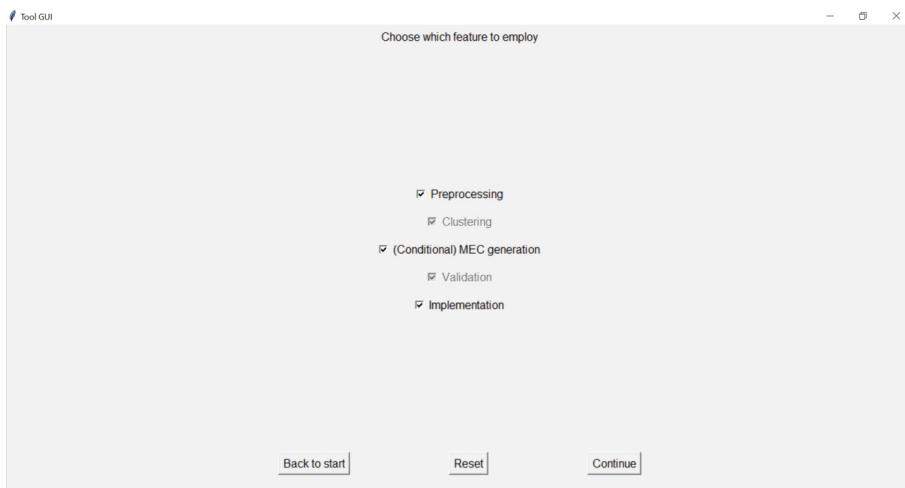


Fig. 11. Feature screen

A.3 Initialization

Step 1: Select the software for implementation; when OpenDSS is selected, insert the network name as well (Fig. 12).

Step 2: Press select. If done correctly, the console shows information on the input type selected and network parameters (number of normal/transformer loads). The continue button is enabled and should be pressed.



Fig. 12. Implementation initialization screen

A.4 Importing data

Step 1: Determine the number of inputs that should be processed by checking or unchecking input 2 (Fig. 13).

The tool is able to process two inputs simultaneously. This is especially useful when implementing into Vision, since different SLPs can be assigned to normal loads and transformer loads. If an implementation in OpenDSS is chosen, only one input can be processed, as it does not differentiate between normal and transformer loads. This also holds when the network in Vision only contains normal loads or transformer loads.

Step 2: Fill in each input, selecting its type, file name, file type, and unit.

The input type acts as a label in all the following screens to help distinguish between inputs. The unit should be the same for both inputs when two inputs are selected.)

Step 3: Press import. If done correctly, the console shows information on how many files are imported, and the continue button is enabled and should be pressed.



Fig. 13. Importing data screen

A.5 Preprocessing

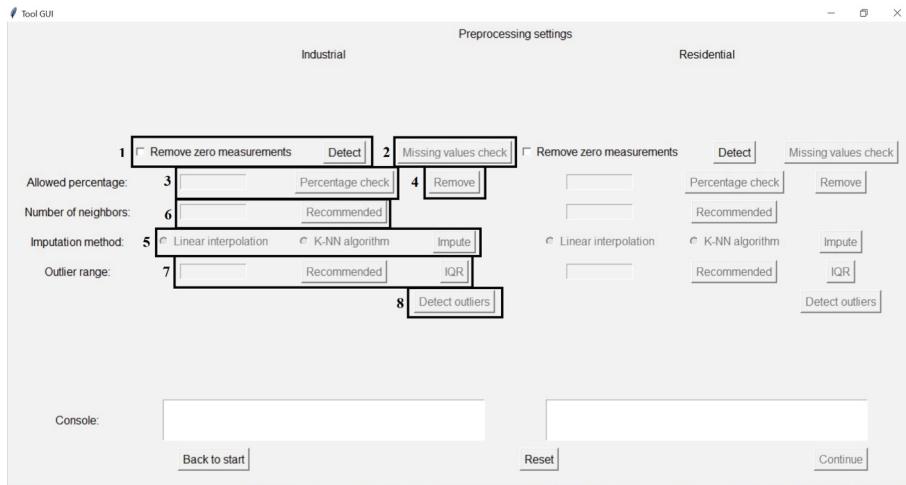


Fig. 14. Preprocessing settings screen

Step 1: Check or uncheck the removal of zero measurements and press the detect button as seen in box 1. The console shows information on the number of duplicates, empty, and zero measurements (Fig. 14).

The option for removing zero measurements is only available when only the preprocessing feature is selected. Otherwise, this option can be selected in the clustering stage.

Step 2: Press the missing values check button in box 2. The console shows information on the number of missing values and its percentage in the dataset.

If the input dataset contains missing values, the percentage check button is enabled in box 3. Otherwise skip to step 6.

Step 3: Fill in a value for the allowed percentage and press the percentage check button in box 3. The console shows information on the number of measurements that exceed this limit.

The value must be between 0 and 100. If measurements contain more missing values than the allowed percentage, the remove button in box 4 is enabled. In addition, the buttons in box 5 are also enabled.

Step 4: The remove button can be pressed to remove these measurements from the dataset.

If the dataset does not contain any more missing values, skip to Step 6.

Step 5: Choose an imputation method and press the impute button. When choosing K-NN, the number of neighbors (custom or recommended) should first be set in box 6 before pressing the impute button.

The custom number of neighbors should be an integer greater than 0 and less than the total data points of one measurement. The recommended number is 2.

Step 6: Choose to continue or search for outliers by setting the outlier range breakdown point (custom or recommended) in the now enabled box 7 and pressing the detect outliers button in box 8. The console shows the number of outliers and its percentage over the dataset.

The custom outlier range breakdown point should be a percentage expressed in decimal form and between 0 and 100. For example, 10% should be inserted as 0.1. The recommended value is 1% or 0.01.

Step 7: Pressing the reset button will reset the settings and requires steps 1 to 6 to be followed again (Fig. 15).

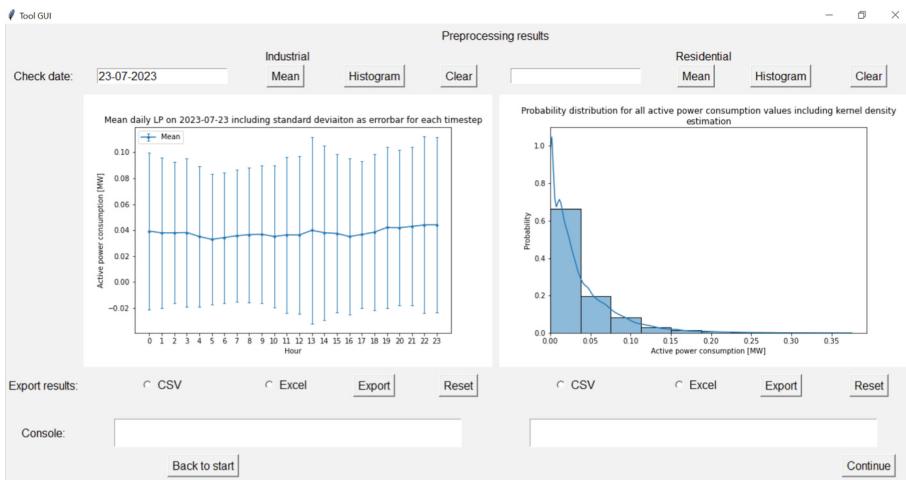


Fig. 15. Preprocessing results screen

Step 1: Press either the mean button or the histogram button.

This is optional. Furthermore, a specific date can be inserted to show the mean of a certain day contained in the dataset. Otherwise, a random date will be chosen.

Step 2: Select CSV or Excel and press the export button.

This is again optional. The exported file is named label_preprocessed_results, where the label refers to the set of input labels on the importing screen.

Step 3: Press continue to finish the preprocessing stage.

A.6 Clustering

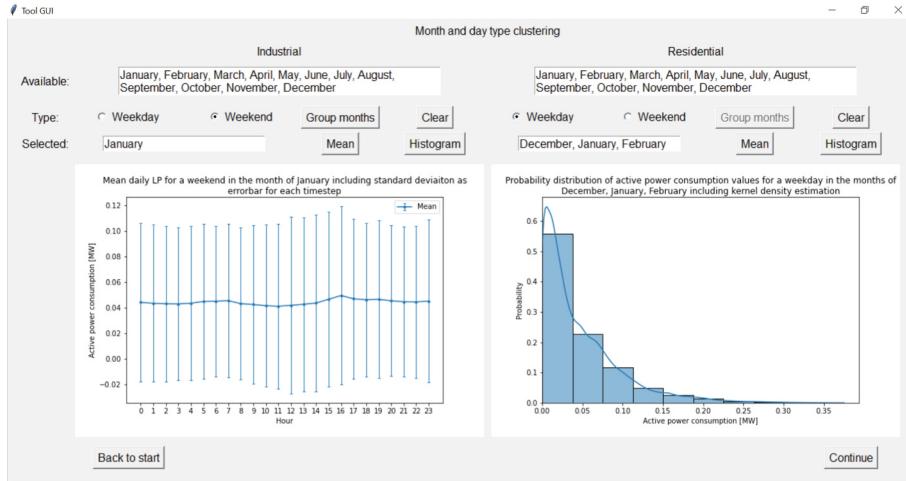


Fig. 16. Clustering selection screen

Step 1: Insert one or more month names from the available months into the selected text box (Fig. 16).

Make sure the name is correctly spelled and starts with capital letters. When inserting multiple names, make sure that they are separated with a comma, space after the comma is not required.

Step 2: When multiple names are inserted, first press the group months button. Now, select the type of day available and press the mean or histogram button.

Step 3: Pressing the clear button will reset this clustering process and requires Steps 1 and 2 to be performed again. Otherwise, the continue button should be pressed.

Step 1: Check or uncheck the removal of zero measurements and fill in the minimum and maximum number of clusters (Fig. 17).

Step 2: Press the Apply K-means button. The console shows information on number of removed zero measurements, the optimum number of clusters with corresponding DBI and the number of days each cluster consists of.

The number of days refers to the number of daily LPs.

Step 3: Select CSV or Excel and press the export button.

This is optional. The exported file is named label_cluster_clusternumber, where the label refers to the set of input labels on the importing screen and the cluster number to which cluster it is.

Step 4: Pressing the clear button will reset this clustering process and requires Steps 1 and 2 to be performed again. Otherwise, the continue button should be pressed (Fig. 18).

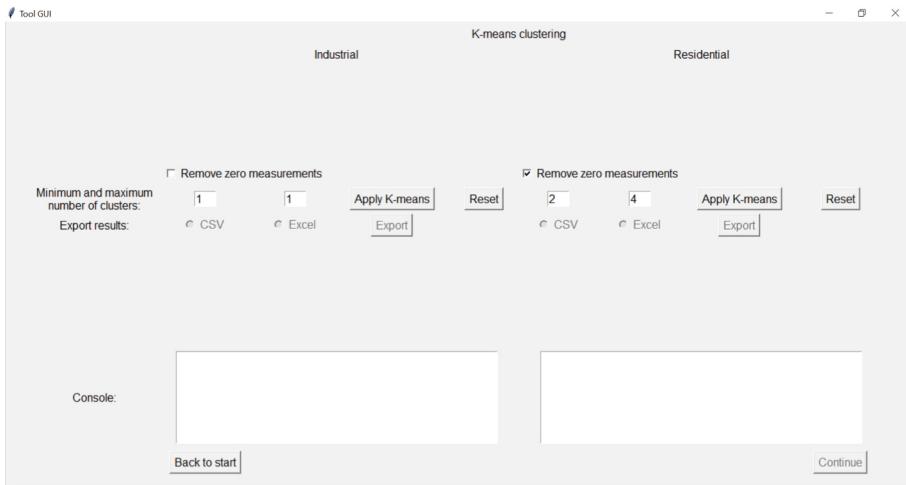


Fig. 17. Clustering process screen

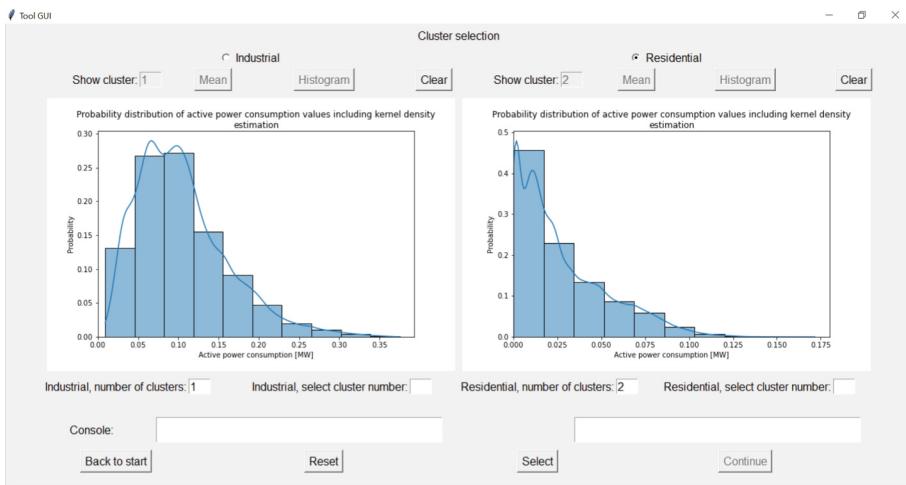


Fig. 18. Clustering results screen

Step 1: Select either of the input labels at the top when two inputs are processed.

If only 1 input is processed this is not required.

Step 2: Fill in which cluster should be shown in the text box above the plot and on the left side of the screen. Followed by pressing either the mean or histogram button next to it.

Below the plot, the number of clusters in which each input is divided in is shown. For example, if this number is 2, then clusters 1 or 2 can be shown.

Step 3: Repeat step 2 but now for the right side of the screen.

Step 4: Pressing the clear button on either side of the screen will reset the plot and inserted cluster on that side of the screen.

Step 5: Fill in which cluster number should be selected to be used in the model in the text box select cluster number below the plot.

Again, if for example the number of clusters in which the input is divided is 2, then cluster 1 or 2 can be selected.

Step 6: Repeat step 5 for the right side of the screen if 2 inputs are processed.

Step 7: Press the select button. The console shows information on which cluster is selected and the number of days it consists of.

Step 8: Pressing the reset button will reset this selection process and requires Steps 5, 6 and 7 to be performed again. Otherwise, the continue button should be pressed.

A.7 Generation

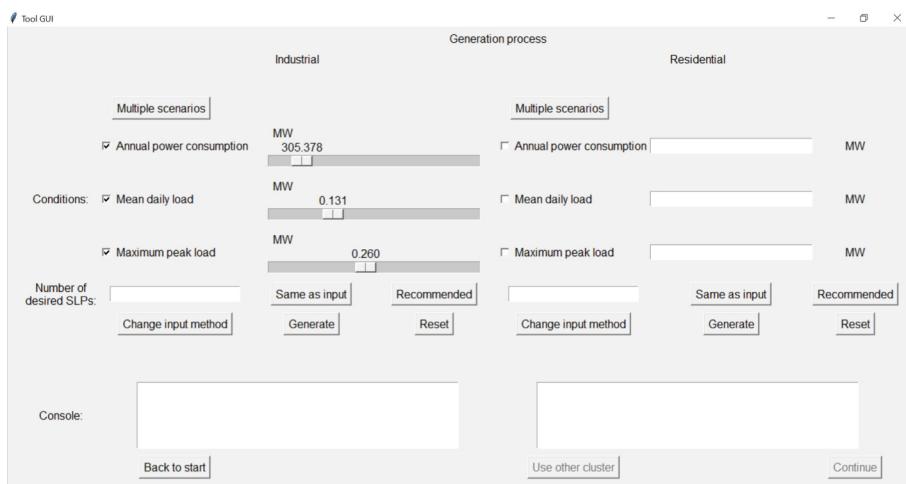


Fig. 19. Generation settings screen

Step 1: Determine if multiple scenarios should be generated, if so, press the multiple scenarios button (Fig. 19).

Step 2: Select the required conditions by either checking or unchecking the relevant condition.

Step 3: Set the value for each selected condition by moving the slider.

It is also possible to insert exact values by pressing the change input method button, which will change the sliders in text boxes where an exact value can be inserted. Make sure that this exact value is within the range allowed provided by the slider.

Step 4: Set the number of desired SLPs (custom, recommended, or same as input).

The custom number must be an integer greater than 0. The recommended number is 10000.

Step 5: Press the generate button. The console shows information on the number of successfully generated SLPs, SR, MAPE, RMSE, and JS-D.

Step 6: If multiple scenarios were selected, repeat Steps 2 to 5.

The console will show the same information for the second scenario. It is possible to still view the information of the first scenario by scrolling in the console.

Step 7: Pressing the reset button will reset the generation process and requires Steps 1 to 6 to be performed again.

Step 8: Either press the continue button or the use other cluster button.

Pressing the use other cluster button results in going back to the previous clustering results screen where the selection of cluster number has to be performed again. Furthermore, Steps 1 to 6 on the generation settings screen have to be performed again as well.

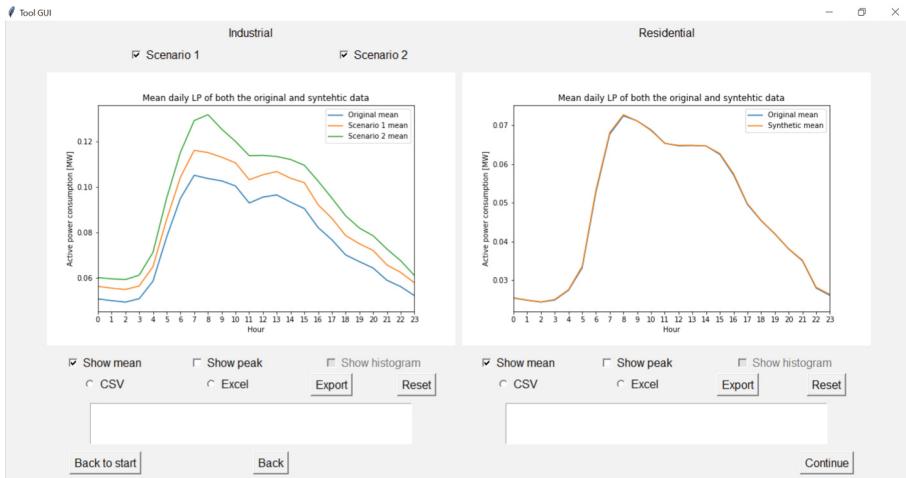


Fig. 20. Generation results screen

Step 1: If multiple scenarios were generated, select above the plot which scenarios should be shown in the plot (Fig. 20).

Step 2: Select to show either the mean, peak, or histogram.

It is possible to show both the mean and the peak in the same plot. When pressing show mean or show peak, the show histogram button is disabled and vice versa.

Step 3: Select CSV or Excel and press the export button.

This is optional. The exported file is named label_generated_SLPs_scenario_x, where the label refers to the set of input labels on the importing screen and scenario_x to which scenario if multiple scenarios were selected.

A.8 Implementation

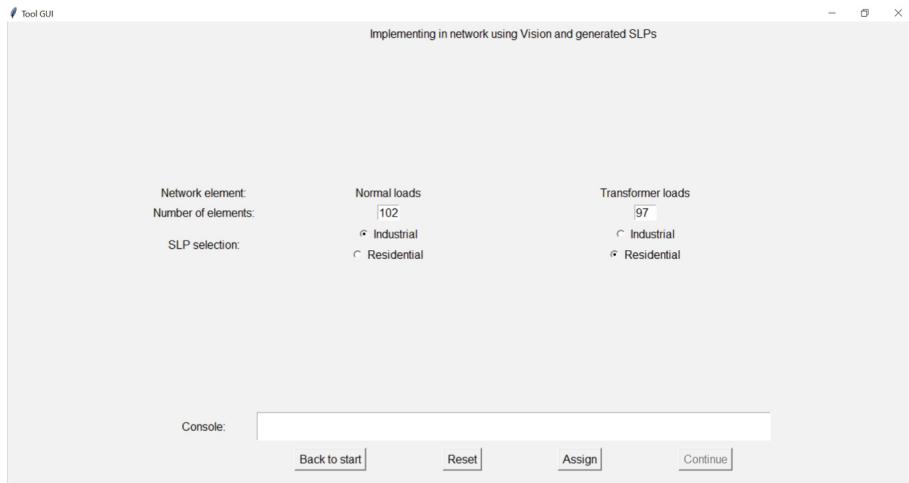


Fig. 21. Implementation process screen

Step 1: Select which SLPs generated should be assigned to which type of loads (Fig. 21).

Only when using Vision for implementation can the network contain both normal loads and transformer loads. Furthermore, if the number of generated SLPs is less than the number of loads, the SLPs are resampled.

Step 2: Press the assign button. The console shows information on whether the assignment was successful.

Step 3: Press the continue button.

If Vision is used, this will terminate the tool script as the load flow calculation is initiated in the macro language. For OpenDSS this will perform the load flow calculation in Python and goes to the next screen.

A.9 Congestion results

Step 1: Press export and finish (Fig. 22).

This will export the congestion results to Excel and terminate the tool script.

A.10 Validation

If only the validation feature were to be chosen, the following steps should be taken (Fig. 23).

Step 1: Press either the mean button or the histogram button.

This is optional and pressing the clear button will reset the plot.

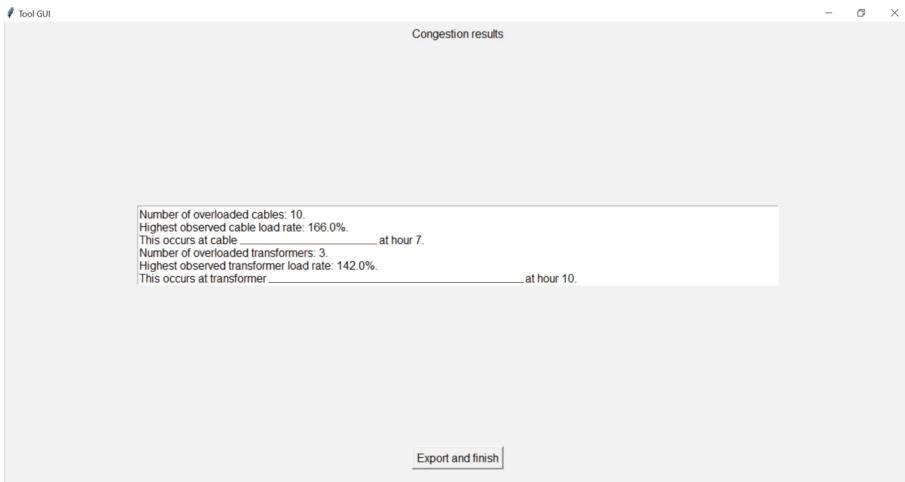


Fig. 22. Congestion results screen

Step 2: Press the validate button. The console shows information on all the validation results.

Step 3: Press the continue button.



Fig. 23. Validation screen

References

1. Arruda Moura, R.A., Khalid Mehmood, K., Van Der Molen, A., Van Der Wielen, P., Nguyen, P.H.: Generating industrial and residential load profile scenarios with a multivariate elliptical copula approach. In: 2024 59th International Universities Power Engineering Conference (UPEC), pp. 1–6 (2024). <https://doi.org/10.1109/UPEC61344.2024.10892609>
2. Arruda Moura, R.A., Khalid Mehmood, K., Van Der Molen, A., Van Der Wielen, P., Nguyen, P.H.: Generation of load profile scenarios using multivariate elliptical copulas. In: IEEE Young Researchers Symposium 2024 (2024). <https://research.tue.nl/en/publications/generation-of-load-profile-scenarios-using-multivariate-elliptica>
3. Cartwright, M.H., Shepperd, M.J., Song, Q.: Dealing with missing software project data. In: Proceedings - International Software Metrics Symposium. vol. 2003-Janua, pp. 154–165 (2003). <https://doi.org/10.1109/METRIC.2003.1232464>
4. Chicco, D., Warrens, M.J., Jurman, G.: The coefficient of determination R-squared is more informative than SMAPE, MAE, MAPE, MSE and RMSE in regression analysis evaluation. PeerJ Comput. Sci. **7**, 1–24 (7 2021). <https://doi.org/10.7717/PEERJ-CS.623>, <https://research.rug.nl/en/publications/the-coefficient-of-determination-r-squared-is-more-informative-th>
5. Dahunsi, F.M., Olawumi, A.E., Ale, D.T., Sarumi, O.A.: A systematic review of data pre-processing methods and unsupervised mining methods used in profiling smart meter data. AIMS Electron. Electr. Eng. **5**(4), 284–314 (2021). <https://doi.org/10.3934/electreng.2021015>
6. Davies, D.L., Bouldin, D.W.: A cluster separation measure. IEEE Trans. Pattern Anal. Mach. Intell. **PAMI-1**(2), 224–227 (1979). <https://doi.org/10.1109/TPAMI.1979.4766909>
7. Duda, R., Hart, P.: Pattern classification and scene analysis. In: A Wiley-Interscience publication (1974)
8. Duque, E.M.S., Vergara, P.P., Nguyen, P.H., van der Molen, A., Slootweg, J.G.: Conditional multivariate elliptical copulas to model residential load profiles from smart meter data. IEEE Trans. Smart Grid **12**(5), 4280–4294 (2021). <https://doi.org/10.1109/TSG.2021.3078394>
9. Ebrahimi, S.R., Rahimiyan, M., Assili, M., Hajizadeh, A.: Home energy management under correlated uncertainties: a statistical analysis through copula. Appl. Energy **305**, 117753 (2022). <https://doi.org/10.1016/j.apenergy.2021.117753>
10. Fávero, L.P., Beifiore, P.: Manual de Análise de Dados - Estatística e Modelagem Multivariada com Excel®, SPSS® e Stata®. Elsevier Brasil (2017)
11. Gebhard, T., Brucherseifer, E., Steinke, F.: Monitoring electricity demand synchronization using copulas. In: 2022 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe), pp. 1–6 (2022). <https://doi.org/10.1109/ISGT-Europe54678.2022.9960369>
12. Giorgi, L., Obushevs, A., Korba, P.: Open source module for the investigation of the impact of electric vehicles in a low voltage grid. In: 2021 IEEE 62nd International Scientific Conference on Power and Electrical Engineering of Riga Technical University (RTUCON), pp. 1–10 (2021). <https://doi.org/10.1109/RTUCON53541.2021.9711593>
13. Halkidi, M., Batistakis, Y., Vazirgiannis, M.: On clustering validation techniques. J. Intell. Inf. Syst. **17**(2–3), 107–145 (2001). <https://doi.org/10.1023/A:1012801612483>

14. International Energy Agency: Electricity Grids and Secure Energy Transitions: Enhancing the Foundations of Resilient, Sustainable and Affordable Power Systems. OECD Publishing, Paris (2023). <https://doi.org/10.1787/455dd4fb-en>
15. Jain, A.K., Duin, R.P., Mao, J.: Statistical pattern recognition: a review. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(1), 4–37 (2000). <https://doi.org/10.1109/34.824819>
16. Kullback, S.: *Information Theory and Statistics*. Courier Corporation (1997)
17. Lin, J.: Divergence measures based on the Shannon entropy. *IEEE Trans. Inf. Theory* **37**(1), 145–151 (1991). <https://doi.org/10.1109/18.61115>
18. MacQueen, J., et al.: Some methods for classification and analysis of multivariate observations. In: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, no. 14, pp. 281–297 (1967). <https://projecteuclid.org/ebooks/berkeley-symposium-on-mathematical-statistics-and-probability/Proceedings-of-the-Fifth-Berkeley-Symposium-on-Mathematical-Statistics-and/chapter/Some-methods-for-classification-and-analysis-of-multivariate-observations/bsm>
19. Mehmood, K.K., Arruda Moura, R.A., Van Der Molen, A., Tonkoski, R., Van Der Wielen, P., Nguyen, P.H.: Efficient congestion management: optimizing non-firm capacities for maximized grid utilization. In: *2024 IEEE International Conference on Environment and Electrical Engineering and 2024 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe)*, pp. 1–6 (2024). <https://doi.org/10.1109/EEEIC/ICPSEurope61470.2024.10751051>
20. Minakshi Vohra, R.G.: Missing value imputation in multi attribute data set. *Int. J. Comput. Sci. Inf. Technol.* **5**(4), 5315–5321 (2014). www.ijcsit.com
21. Murti, D.M.P., Pujianto, U., Wibawa, A.P., Akbar, M.I.: K-Nearest Neighbor (K-NN) based missing data imputation. In: *Proceeding - 2019 5th International Conference on Science in Information Technology: Embracing Industry 4.0: Towards Innovation in Cyber Physical System, ICSITech 2019*, pp. 83–88. Institute of Electrical and Electronics Engineers Inc. (2019). <https://doi.org/10.1109/ICSI Tech46713.2019.8987530>
22. Nelsen, R.B.: *An Introduction to Copulas*. Springer Series in Statistics, 2 edn. Springer, New York (2006). <https://doi.org/10.1007/0-387-28678-0>
23. Nielsen, F.: On the Jensen-Shannon symmetrization of distances relying on abstract means. *Entropy* **21**(5) (2019). <https://doi.org/10.3390/e21050485>
24. Plomp, H.: Synthetic Medium Voltage Load Profile Generator (2025). <https://github.com/hendrikplomp/Synthetic-medium-voltage-load-profile-generator>, GitHub repository
25. Silva, W.N., Bandória, L.H., Dias, B.H., de Almeida, M.C., de Oliveira, L.W.: Generating realistic load profiles in smart grids: an approach based on nonlinear independent component estimation (NICE) and convolutional layers. *Appl. Energy* **351**, 121902 (2023). <https://doi.org/10.1016/j.apenergy.2023.121902>
26. Sklar, M.: Fonctions de répartition à N dimensions et leurs marges. *Annales de l'ISUP* **VIII**(3), 229–231 (1959). <https://hal.science/hal-04094463>
27. Wang, L., Fu, D.M.: Estimation of missing values using a weighted k-nearest neighbors algorithm. In: *Proceedings - 2009 International Conference on Environmental Science and Information Application Technology, ESIAT 2009*, vol. 3, pp. 660–663 (2009). <https://doi.org/10.1109/ESIAT.2009.206>



Decentralized Reinforcement Learning for Adaptive Power Sharing in Hybrid DC Microgrids

Abd Alelah Derbas¹✉, Chiara Bordin¹, Sambeet Mishra², and Frede Blaabjerg³

¹ The Arctic University of Norway (UiT), Tromsø, Norway
{abd.a.derbas,chiara.bordin}@uit.no

² University of South-Eastern Norway (USN), Porsgrunn, Norway
sambeet.mishra@usn.no

³ Aalborg University (AAU), Aalborg East, Aalborg, Denmark
fbl@energy.aau.dk

Abstract. This paper proposes a decentralized voltage control strategy for islanded DC microgrids that replaces conventional droop control with a reinforcement learning (RL)-based approach. Using a Deep Deterministic Policy Gradient (DDPG) agent, the controller learns to generate real-time voltage references based solely on local measurements, eliminating the need for inter-unit communication. Compared to droop control, the proposed method reduces power sharing error from +30% to +8% and halves bus voltage deviation under high line impedance scenarios. The framework adapts to dynamic load and network conditions, offering a scalable and resilient control solution for next-generation microgrids.

Keywords: Reinforcement learning · data-driven control · energy management · decentralized control · DC microgrids · hybrid distributed generation

1 Introduction

Islanded DC microgrids (DC-MGs) have emerged as a promising solution for resilient and sustainable power supply in remote, rural, and off-grid regions due to their efficient integration of photovoltaic (PV) and battery energy storage systems (BESS). Ensuring voltage stability and accurate power sharing in these networks is critical, especially under variable load and generation conditions. Traditional decentralized control strategies, such as droop-based methods, offer simplicity and plug-and-play capability but suffer from inherent limitations in the presence of line impedance mismatch, dynamic disturbances, and non-uniform load profiles [1, 2]. To address these challenges, advanced control techniques, such as distributed optimization [3] and hierarchical control frameworks [4, 5], are

This work is funded by the Aurora Research Center.

being explored. These methods aim to enhance stability, improve power sharing accuracy, and mitigate the effects of mismatched line impedances and dynamic disturbances in islanded DC-MGs.

Classic droop control relies on static voltage-current relationships, which are highly sensitive to physical network asymmetries and can lead to inaccurate power distribution and voltage deviations [6, 7]. Moreover, methods that attempt to compensate for these limitations often require centralized coordination or inter-unit communication, which compromises scalability and system resilience. These challenges are particularly pronounced in DC systems, where resistive voltage drops dominate and fast dynamic response is critical [8, 9].

Recent advances in data-driven control and reinforcement learning (RL) provide a transformative alternative to traditional rule-based or model-driven approaches in microgrid control [10]. Complementary to RL strategies, adaptive fuzzy logic controllers (AFLCs) have also been applied to maximum power point tracking (MPPT) in PV systems, demonstrating improved resilience and tracking accuracy under fluctuating solar conditions [11]. These methods collectively reframe the voltage regulation problem as a policy optimization task, where the objective is to learn a mapping from system observations to control actions that optimize long-term performance [12]. Unlike classical techniques that rely on predefined models, RL-based controllers learn policies through direct interaction with the environment, enabling them to adapt and improve over time [13].

In particular, deep reinforcement learning combines neural function approximation with trial-based learning, allowing the controller to handle nonlinearities, parameter uncertainties, and unmodeled dynamics that are difficult to capture analytically [14]. This is especially valuable in DC microgrids, where fluctuating load conditions, converter dynamics, and line impedance mismatches introduce complex system behavior [15]. Moreover, RL enables purely local decision-making—each distributed generation (DG) unit can operate autonomously without requiring explicit communication with others or access to centralized coordination, significantly enhancing system modularity and fault tolerance [16].

By optimizing control policies in a simulated environment that mirrors real-world conditions, RL frameworks can be trained offline and deployed with lightweight inference logic on embedded controllers [17]. This makes them well-suited for scalable, real-time voltage regulation in modern microgrids that demand adaptability, robustness, and plug-and-play compatibility [18, 19].

This paper introduces a fully decentralized voltage control strategy for islanded DC microgrids based on deep reinforcement learning. A Deep Deterministic Policy Gradient (DDPG) agent is trained to generate real-time voltage references using only local state measurements. The proposed controller operates above conventional PI-regulated buck converters and replaces classical droop control with a neural network-based policy that adapts to local conditions in real time. The remainder of this paper is organized as follows: Sect. 2 describes the system architecture, Sect. 3 presents the dynamic modeling of the microgrid, Sect. 4.2 introduces the reinforcement learning-based control strategy, Sect. 5 details the simulation results, and Sect. 6 concludes the paper.

2 System Architecture

The proposed islanded DC microgrid consists of n hybrid DG units, each comprising a PV array and a BESS. These sources are connected via independent buck-type DC/DC converters and collectively supply power to a shared DC bus that supports aggregated residential, commercial, and industrial loads.

2.1 Hybrid DG Unit Configuration

Each hybrid DG unit is composed of two power electronic subsystems:

- A PV array connected through a unidirectional buck converter, regulated by a maximum power point tracking (MPPT) algorithm.
- A BESS interfaced through a bidirectional buck converter, enabling controlled charge and discharge operations based on load demand and state-of-charge (SoC) constraints.

The outputs of the PV and battery converters are merged at a local DC node and connected to the main DC bus through a line impedance modeled by a series resistance r_{l_i} and inductance L_{l_i} . This configuration allows flexible energy dispatch from both renewable and storage sources.

2.2 DC Bus and Load Interface

All DG units inject current into a centralized DC bus stabilized by a capacitor C_{dc} , which buffers power imbalances and maintains voltage stability. The DC bus supplies heterogeneous loads, including residential appliances, data centers, and commercial facilities, typically interfaced via downstream DC/DC converters.

2.3 Decentralized Control Architecture

Each DG unit is governed by a three-layer local control architecture:

- **Measurement Layer:** Acquires local state information including output current i_{o_i} , converter voltage V_{c_i} , SoC, PV power, and estimated local load.
- **Decision Layer:** A data-driven policy $\pi(\mathbf{s}_i; \theta)$ computes the optimal converter voltage reference V_i^* based solely on local observations.
- **Execution Layer:** Inner-loop PI controllers regulate converter dynamics to track V_i^* by generating the appropriate PWM duty cycle d_i .

This decentralized architecture eliminates the need for inter-unit communication and enhances system scalability and fault tolerance (Figs. 1 and 2).

3 Dynamic Modeling for Decentralized Control

This section presents the nonlinear dynamic model of the microgrid components as a foundation for the design and training of decentralized control strategies. The modeling focuses on converter dynamics, distribution line behavior, and DC bus voltage regulation, based on the system architecture described in Sect. 2.

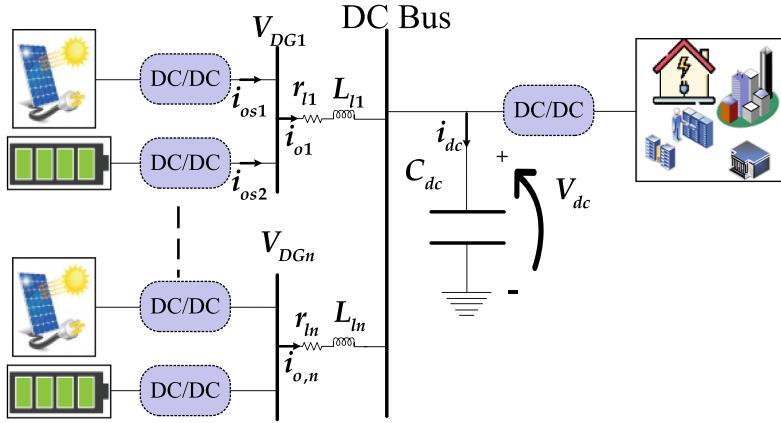


Fig. 1. High-level architecture of the islanded DC microgrid showing hybrid PV-battery DG units interfaced to a common DC bus, with decentralized converter-level control.

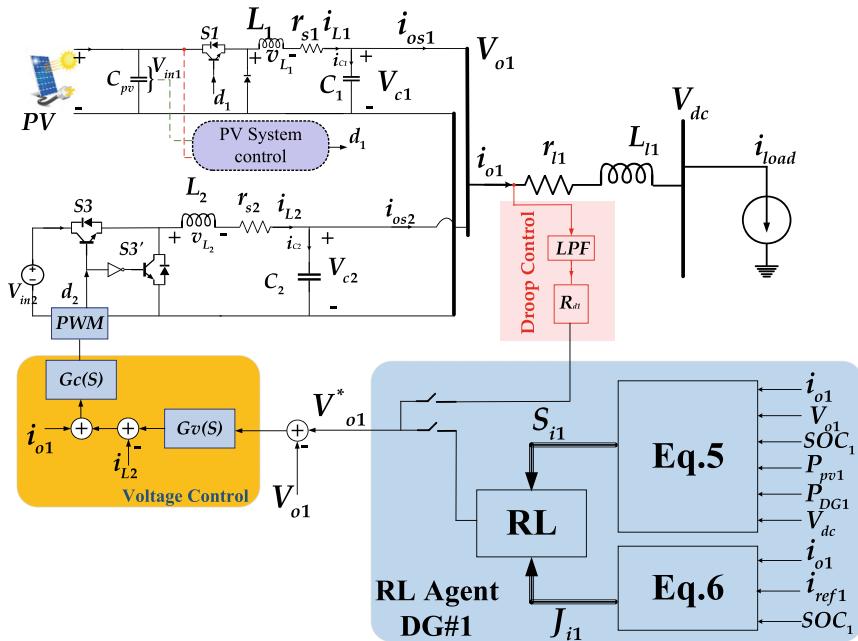


Fig. 2. Block diagram of the proposed RL-based decentralized control strategy for a hybrid PV-battery DG unit interfaced via a buck converter in an islanded DC-MG.

3.1 Integrated Modeling of Hybrid DG Units with Buck Converter Interface

1) Converter-Level Dynamics. The electrical behavior of each buck converter is described by the following state-space equations:

$$\begin{aligned} L_i \frac{di_{L_i}}{dt} &= V_{\text{src},i} \cdot d_i - V_{c_i} - r_{s,i} i_{L_i}, \\ C_i \frac{dV_{c_i}}{dt} &= i_{L_i} - i_{o_i}, \end{aligned} \quad (1)$$

where L_i and C_i denote the output filter inductance and capacitance, respectively; $r_{s,i}$ models the inductor's series resistance; $V_{\text{src},i}$ is the source-side voltage from either PV or BESS; d_i is the duty cycle from the PWM controller; i_{L_i} is the inductor current; V_{c_i} is the converter output voltage; and i_{o_i} is the output current injected into the microgrid.

2) Distribution Line Model. The interface between each DG converter and the main DC bus is represented by an RL line impedance, which introduces voltage drop and dynamic lag. The associated current dynamics are:

$$L_{l_i} \frac{di_{o_i}}{dt} = V_{c_i} - V_{\text{dc}} - r_{l_i} i_{o_i}, \quad (2)$$

where L_{l_i} and r_{l_i} represent the line inductance and resistance, respectively, and V_{dc} is the instantaneous voltage at the DC bus.

3) DC Bus Voltage Dynamics. The global voltage stability of the microgrid is governed by the dynamics of the shared DC bus capacitor. Applying Kirchhoff's current law yields:

$$C_{\text{dc}} \frac{dV_{\text{dc}}}{dt} = \sum_{i=1}^n i_{o_i} - i_{\text{load}}, \quad (3)$$

where C_{dc} is the DC bus capacitance and i_{load} is the total current drawn by the aggregate DC loads.

4) Modeling Implications for Control Design. Equations (1)–(3) collectively form a nonlinear, coupled dynamic model of the microgrid that captures both local converter dynamics and inter-unit interactions through the shared DC bus. RL is employed to develop adaptive control strategies that map observed local states to optimal voltage references. Depending on system nonlinearity and the desired balance between exploration and stability, algorithms such as Deep Deterministic Policy Gradient (DDPG), Proximal Policy Optimization (PPO), or Twin Delayed DDPG (TD3) can be utilized for policy training. This dynamic model serves as the simulation environment for training decentralized RL agents, enabling accurate emulation of converter dynamics, distribution line behavior, and DC bus voltage fluctuations.

Furthermore, the structured state-space formulation facilitates closed-loop validation of the proposed data-driven control framework. It allows performance benchmarking against baseline droop-based strategies under varying operational conditions, including PV intermittency, load transients, and battery SoC fluctuations.

5) Role in Reinforcement Learning-Based Control. The complete model derived in Eqs. (1)–(3) serves as the core environment for training a RL agent, which is used to design a decentralized voltage control policy for each DG unit. The state-space representation enables accurate simulation of real-time dynamics, making it suitable for reward-driven policy learning and closed-loop performance evaluation. Further details on the RL formulation and training process are presented in Sect. 4.2.

4 Reinforcement Learning-Based Voltage Control Framework

This section presents a novel, fully decentralized control framework for energy management and power sharing in islanded DC-MGs with hybrid DG units. The control strategy replaces conventional droop-based methods with a data-driven voltage reference policy, trained via reinforcement or supervised learning. Each DG unit operates autonomously based on local measurements, enabling real-time adaptability and plug-and-play scalability.

4.1 Data-Driven Voltage Reference Generation

Traditional droop control adjusts the voltage reference of each DG unit based on static coefficients linked to the output current, which is sensitive to impedance mismatch and lacks adaptability. To overcome these limitations, we propose a real-time, data-driven voltage reference policy.

Each DG unit computes its local voltage reference V_i^* as a function of its current state:

$$V_i^* = \pi(\mathbf{s}_i; \theta), \quad (4)$$

where $\pi(\cdot)$ is a parameterized policy model (e.g., a neural network), θ denotes the policy parameters, and \mathbf{s}_i is the local state vector defined as:

$$\mathbf{s}_i(t) = [i_{oi}(t), V_{ci}(t), \text{SoC}_i(t), P_{\text{pv},i}(t), P_{\text{load},i}(t)]. \quad (5)$$

This formulation enables each DG to respond to local variations in real time, facilitating intelligent voltage regulation and autonomous participation in system-wide power balancing.

4.2 Policy Learning and Deployment

The policy $\pi(\cdot)$ is trained to minimize a multi-objective cost function that reflects power sharing fairness, energy efficiency, and battery health. The local cost function for DG unit i is:

$$J_i = \alpha_1 (i_{oi} - i_i^{\text{ref}})^2 + \alpha_2 R_i i_{oi}^2 + \alpha_3 (\text{SoC}_i - \text{SoC}_{\text{ref}})^2, \quad (6)$$

where i_i^{ref} is the ideal proportional current, R_i is the local line resistance, SoC_i is the measured state of charge, and SoC_{ref} is the desired battery energy level. The weights $\alpha_1, \alpha_2, \alpha_3$ define the relative priorities of the objectives.

The optimal control policy is obtained by minimizing the expected cumulative cost over time:

$$\pi^* = \arg \min_{\pi} \mathbb{E} \left[\sum_{t=0}^T \gamma^t J_i(t) \right], \quad (7)$$

where $\gamma \in (0, 1]$ is the discount factor. The optimization is conducted using either RL or supervised learning (SL) methods.

1) Reinforcement Learning-Based Training. In the RL approach, the control policy is trained through repeated interactions with the simulation environment modeled in Sect. 3. At each time step, the agent observes the local system state, selects a voltage reference action, and receives a reward based on the negative instantaneous cost $-J_i(t)$. The standard reinforcement learning setup is defined by:

- **State:** $\mathbf{s}_i(t) \in \mathbb{R}^n$ —the local observation vector.
- **Action:** $a_i(t) = V_i^* \in [44, 54]$ —the computed voltage reference.
- **Reward:** $r_i(t) = -J_i(t)$, where $J_i(t)$ is defined in Eq. (6).
- **Policy Update:** Actor-critic methods such as DDPG, PPO, or TD3.

The policy $\pi(\cdot; \theta)$ is parameterized by θ and can be implemented using deep neural networks. Actor-critic frameworks like DDPG are particularly effective in continuous action spaces and enable real-time adaptive control. This setup allows the agent to autonomously discover robust and generalizable control policies across varying operational scenarios.

The complete policy training process for decentralized voltage reference generation is summarized in Algorithm 1.

2) Supervised Learning Alternative. When optimal voltage references are available from centralized solvers such as Model Predictive Control (MPC) or Mixed-Integer Linear Programming (MILP), the control policy can be trained offline using supervised learning. In this setting, a labeled dataset of input-output pairs $\{(\mathbf{s}_i, V_i^*)\}$ is generated, where each input state vector \mathbf{s}_i corresponds to an optimal voltage command V_i^* . The policy network is trained to approximate this mapping, minimizing a regression loss between predicted and target voltage references.

Algorithm 1: Policy training using DDPG for voltage reference control.

Input: Environment model \mathcal{E} (see Sect. 3), initial policy $\pi(\cdot; \theta)$
Output: Optimized policy parameters θ^*

1 Initialization: Randomly initialize actor parameters θ and critic parameters ϕ

2 foreach episode $e = 1$ to N_{episodes} **do**

3 Reset environment \mathcal{E} , obtain initial state $\mathbf{s}_i(0)$

4 foreach time step $t = 0$ to T **do**

5 Select action: $V_i^*(t) = \pi(\mathbf{s}_i(t); \theta) + \mathcal{N}_t$

6 Apply $V_i^*(t)$ to environment, observe: next state $\mathbf{s}_i(t+1)$, reward $r_i(t) = -J_i(t)$

7 Store experience $(\mathbf{s}_i(t), V_i^*(t), r_i(t), \mathbf{s}_i(t+1))$ in replay buffer \mathcal{D}

8 Sample mini-batch from \mathcal{D}

9 Update critic parameters ϕ by minimizing TD error:

$$y_i = r_i + \gamma Q'(\mathbf{s}_i(t+1), \pi'(\mathbf{s}_i(t+1); \theta'); \phi')$$

$$L = \frac{1}{N} \sum_i (y_i - Q(\mathbf{s}_i(t), V_i^*(t); \phi))^2$$

10 Update actor parameters θ using deterministic policy gradient:

$$\nabla_{\theta} J \approx \frac{1}{N} \sum_i \nabla_{V_i^*} Q(\mathbf{s}_i, V_i^*) \cdot \nabla_{\theta} \pi(\mathbf{s}_i; \theta)$$

11 Soft update target networks:

$$\theta' \leftarrow \tau \theta + (1 - \tau) \theta', \quad \phi' \leftarrow \tau \phi + (1 - \tau) \phi'$$

12 **end**

13 **end**

3) Policy Execution and Integration. RL is applied at the decision-making layer to learn a data-driven voltage reference policy that maps local system states to optimal control actions. The trained policy operates above the converter-level PI controllers, enabling intelligent and adaptive voltage regulation in a fully decentralized manner.

During online operation, the trained policy is deployed as a lightweight evaluator within each DG unit. At every control interval:

1. The local state vector \mathbf{s}_i is measured or estimated.
2. The policy evaluates $V_i^* = \pi(\mathbf{s}_i; \theta)$, producing the optimal voltage reference.
3. The inner-loop PI controllers track V_i^* by adjusting the PWM duty cycle d_i of the buck converter.

This structure ensures fast and reliable real-time execution with minimal computational overhead. The complete inference and control routine is summarized in Algorithm 2.

Algorithm 2: Online execution of the learned RL policy using local state \mathbf{s}_i and PI-based voltage tracking.

Input: Local state vector $\mathbf{s}_i(t) = [i_{oi}(t), V_{ci}(t), \text{SoC}_i(t), P_{\text{pv},i}(t), P_{\text{DG},i}(t), V_{\text{dc}}(t)]$
Output: Voltage reference $V_i^*(t)$ for converter control

- 1 **Initialization:** Load pre-trained policy $\pi(\cdot; \theta)$ for DG unit i
- 2 **foreach** control time step t **do**
- 3 Measure local observations $\mathbf{s}_i(t)$
- 4 Compute optimal voltage reference:
- 5 $V_i^*(t) \leftarrow \pi(\mathbf{s}_i(t); \theta)$
- 6 Track $V_i^*(t)$ using inner-loop PI voltage and current controllers
- 7 Generate PWM duty cycle $d_i(t)$ for buck converter switching
- 8 **end**

4.3 Decentralized and Scalable Operation

The proposed control framework is fully decentralized. Each DG unit independently computes its voltage reference based solely on local measurements, without requiring communication or coordination with other units. This architecture inherently supports:

- **Resilience to Communication Failures:** Control remains stable under loss of connectivity or isolated operation.
- **Modularity:** Additional DG units can be added or removed without reconfiguring the global control logic.
- **Scalability:** The control scheme scales linearly with system size, suitable for large-scale or heterogeneous microgrids.

By coupling intelligent voltage reference generation with converter-level autonomy, the framework enables adaptive power sharing, enhanced bus voltage stability, and efficient energy management under renewable variability and dynamic load profiles. The next section presents simulation results validating the proposed method's effectiveness under multiple operating conditions.

5 Simulation and Results

To validate the performance of the proposed RL-based control strategy, a detailed MATLAB/Simulink simulation model is developed using the system configuration shown in Fig. 3. The testbed replicates an islanded DC-MG comprising two hybrid DG units with rated powers of 2.2 kW and 1.1 kW, respectively. Each DG interfaces a PV source and a (BESS) system through buck-type DC/DC converters. Local and common loads are included to emulate realistic demand-side variations.

All DGs are connected to a central DC bus through distribution lines modeled as resistive-inductive rl elements. The control system operates in a decentralized

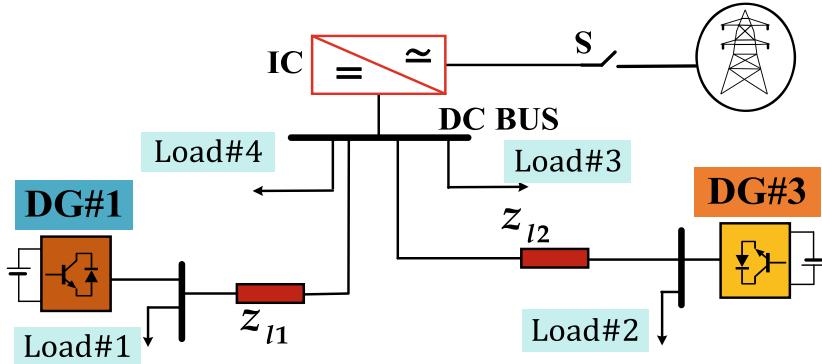


Fig. 3. Configuration of the implemented islanded DC-MG.

Table 1. Simulation Parameters of the Islanded DC Microgrid

Category	Parameter	Value
DG Units	Rated Power (DG1, DG2)	2.2 kW, 1.1 kW
Local Loads	P_{L1}, P_{L2}	0.5 kW each
PCC Loads	P_{L3}, P_{L4}	0.7 kW, 1.5 kW
Line Impedance	$L_{\text{line}}, R_{\text{line}}$	6 μH , 4.6 $\text{m}\Omega$
Converter Filter	L_f, R_f	2.5 mH, 0.01 Ω
	C_f	500 μF
DC Bus	Bus Capacitance C_{dc}	500 μF
PI Controllers	Voltage ($K_{p,v}, K_{i,v}$)	4, 300
	Current ($K_{p,i}, K_{i,i}$)	8, 200
Battery Model	Nominal Capacity (C_{nom})	40 Ah
Reinforcement Learning	Agents	2 \times DDPG (DG1, DG2)
	Observation Dimension	6
	Action Bounds	44 V to 54 V
	Policy Sample Time	2 ms
Simulation	Sampling Time T_s	20 μs

fashion, with inner-loop PI controllers for voltage and current regulation, and an RL agent in the outer loop responsible for generating adaptive voltage references.

Simulation parameters, including converter filter components, line impedances, and controller gains, are summarized in Table 1. The simulations are executed in MATLAB/Simulink (R2024b) with a sampling time of 20 μs and a total simulation duration of 10 s, ensuring adequate capture of both steady-state behavior and dynamic transients. Two voltage control strategies are implemented and compared:

1. **Baseline:** Conventional classic droop control.

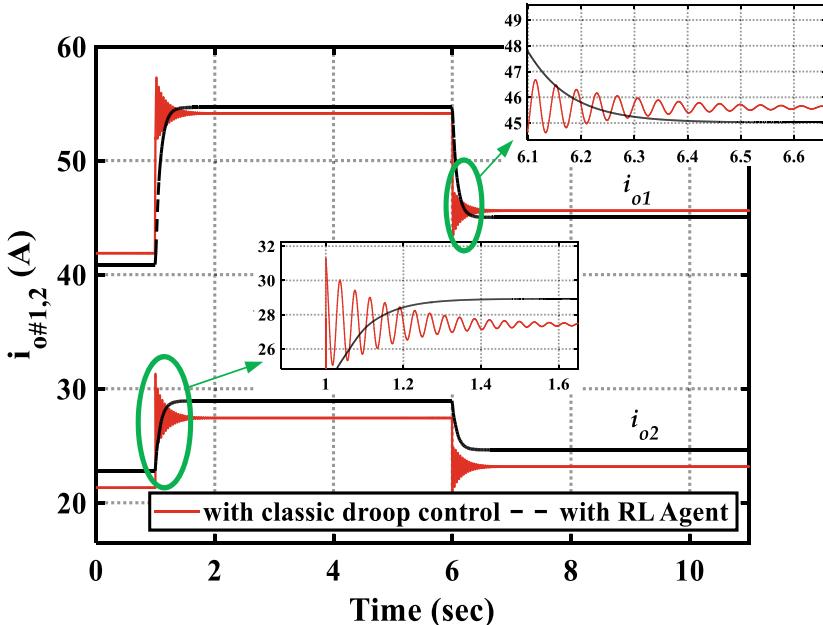


Fig. 4. DG output currents i_{o1} and i_{o2} during load variation scenario. RL-based control eliminates oscillations compared to droop control.

2. **Proposed:** A decentralized reinforcement learning-based voltage control framework employing two DDPG agents, each associated with a hybrid DG unit.

The RL agents are trained and deployed using MATLAB's Reinforcement Learning Toolbox, interacting with the Simulink-based microgrid model in real time. The reward function is designed to optimize power sharing accuracy, battery SoC regulation, and energy efficiency under varying operating conditions.

5.1 Scenario 1: Load Variation

To evaluate the adaptability of the proposed RL-based voltage control strategy under dynamic load conditions, Scenario 1 considers two successive events at the PCC. At $t = 1$ s, Load#4 is increased from 1.5 kW to 2.5 kW, followed by the disconnection of Load#3 at $t = 6$ s. These changes impose significant power redistribution demands on the microgrid.

Figure 4 shows the output currents i_{o1} and i_{o2} under both classical droop control and the proposed RL-based controller. The droop controller induces oscillations with dominant frequencies around 26–33 Hz following both load transients. In contrast, the RL controller effectively suppresses these oscillations and stabilizes current trajectories faster, enhancing dynamic response and power sharing stability.

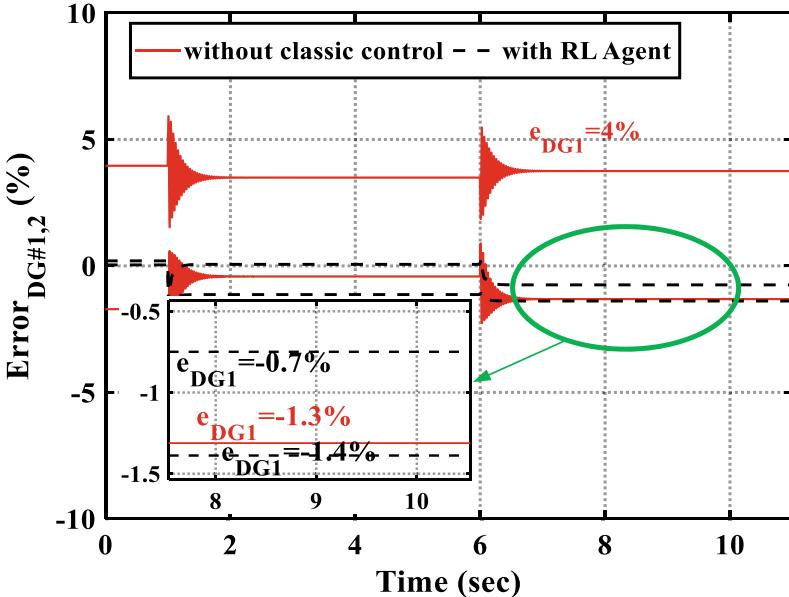


Fig. 5. Power sharing error comparison between RL-based and droop control under load transients.

The power sharing accuracy is assessed in Fig. 5. Under classical droop control, DG#1 exhibits a sharing error of approximately +4%, whereas the RL-based controller limits the error to around -1.3%, yielding a significantly smaller deviation. This improved performance demonstrates the RL agent's ability to maintain accurate and stable current sharing while responding adaptively to sudden load variations.

5.2 Scenario 2: Impact of High Line Impedance

This scenario evaluates the impact of elevated distribution line resistance R_{line} on voltage regulation and power sharing performance in the DC microgrid. While inductive effects dominate in AC systems, resistive line drops are more critical in DC networks. Accordingly, the line resistance is artificially increased tenfold from the nominal value in Table 1, i.e., $R_{\text{line}} = 46 \text{ m}\Omega$, to emulate extended cable runs or degraded conductor conditions, while keeping L_{line} unchanged. The same disturbance conditions applied in Scenario 1 are also used here: Load #4 is increased from 1.5 kW to 2.5 kW at $t = 1 \text{ s}$, and Load 3 is disconnected from the PCC at $t = 6 \text{ s}$. These events are used to trigger transient power redistribution in the microgrid and evaluate controller performance.

Figure 6 shows the power sharing error for both the conventional droop-based controller and the proposed RL-based controller. The droop control, which relies on static $V - I$ characteristics, suffers from severe power imbalance due to the

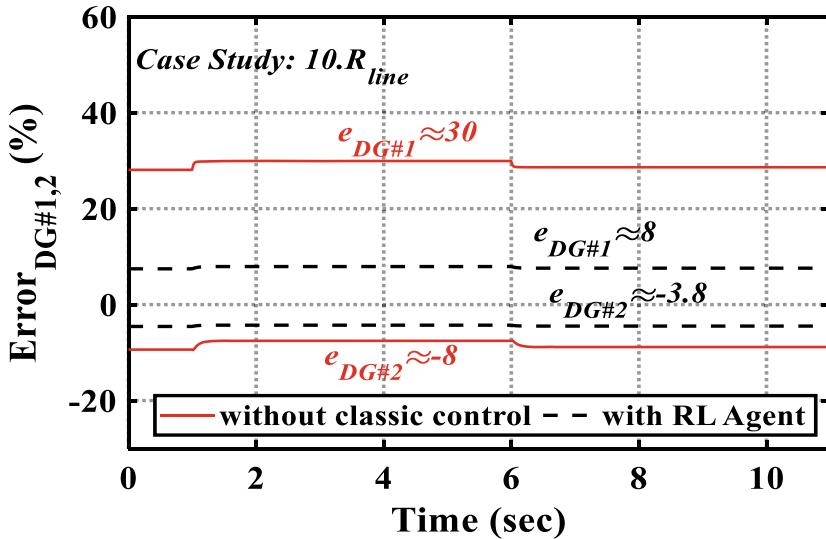


Fig. 6. Power sharing error of DG units under high line impedance ($10 \times Z_{line}$).

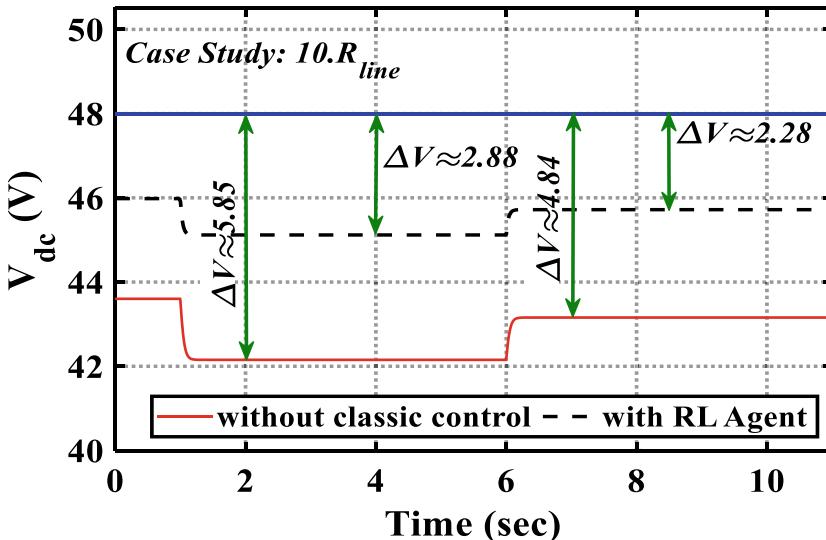


Fig. 7. DC bus voltage profile under high impedance conditions using classic control and RL agent.

asymmetric voltage drops introduced by the resistive lines. This results in significant deviation from ideal current sharing, with DG#1 showing an overshoot of approximately +30%, and DG#2 undershooting by around -8%. In contrast, the RL-based policy limits these deviations to within +8% and -3.8%, respectively.

In Fig. 7, the bus voltage V_{dc} is also shown to degrade considerably under the droop controller, with transient voltage dips reaching up to 5.85 V below nominal. The RL agent is able to constrain the deviation to less than 2.88 V, reflecting improved robustness to resistive voltage drops.

Despite the absence of low-frequency oscillations under high impedance conditions for both methods, the RL strategy significantly outperforms the droop controller in terms of voltage regulation accuracy and current sharing fairness, validating its robustness in high-impedance microgrid applications.

Across both scenarios, the RL-based controller outperforms classical droop in terms of dynamic response, power sharing fairness, and voltage regulation. Particularly in high-impedance conditions, RL mitigates current imbalance and avoids excessive bus voltage dips. These results underscore the framework's robustness to real-world variability.

6 Conclusion

A detailed nonlinear dynamic model of the DC-MG, including converter dynamic, distribution lines, and shared bus behavior, was developed to serve as the RL training environment. The control strategy was evaluated through MATLAB/Simulink simulations across representative scenarios involving sudden load variations and high line resistance condition. The RL-based controller demonstrated superiority in enhancing the dynamic voltage response, and reducing power sharing errors during load disturbances. It also maintained the accuracy of power sharing and reduced the dc bus voltage deviations under high line resistance conditions. These results underscore the RL strategy's adaptability, resilience, and scalability, making it a strong candidate for decentralized control in multi DG microgrid.

While this work benchmarked performance against conventional droop control, future studies will include comparisons with more advanced strategies such as model predictive control and hierarchical schemes to offer a broader evaluation.

Acknowledgements. This research is supported by the UiT Aurora project MAS-COT - Mathematical Structures in Computation. Additionally, participation in the conference is supported by the Nordic Energy Informatics Academy Network.

References

1. Al-Ismail, F.S.: A critical review on DC microgrids voltage control and power management. *IEEE Access* **12**, 30345–30361 (2024)
2. Nabatirad, M., Razzaghi, R., Bahrani, B.: Decentralized energy management and voltage regulation in islanded DC microgrids. *IEEE Syst. J.* **16**(4), 5835–5844 (2022)
3. Dissanayake, A.M., Ekneligoda, N.C.: Multiobjective optimization of droop-controlled distributed generators in DC microgrids. *IEEE Trans. Industr. Inf.* **16**(4), 2423–2435 (2020)

4. Nahata, P., La Bella, A., Scattolini, R., Ferrari-Trecate, G.: Hierarchical control in islanded DC microgrids with flexible structures. *IEEE Trans. Control Syst. Technol.* **29**(6), 2379–2392 (2021)
5. Derbas, A.A., Bordin, C., Mishra, S., Hamzeh, M., Blaabjerg, F.: AC microgrid modeling and adaptive control using biomimetic valence learning: an AI-based approach. In: 2024 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), pp. 117–122 (2024)
6. Derbas, A.A., Hamzeh, M.: A new power sharing method for improving power management in DC microgrid with power electronic interfaced distributed generations. In: 2019 27th Iranian Conference on Electrical Engineering (ICEE), pp. 624–629 (2019)
7. Derbas, A.A., Oshnoei, A., Azzouz, M.A., Awad, A.S.A., Blaabjerg, F., Anvari-Moghaddam, A.: Adaptive damping control to enhance small-signal stability of dc microgrids. *IEEE J. Emerg. Sel. Top. Power Electron.* **11**(3), 2963–2978 (2023)
8. Tah, A., Das, D.: An enhanced droop control method for accurate load sharing and voltage improvement of isolated and interconnected dc microgrids. *IEEE Trans. Sustain. Energy* **7**(3), 1194–1204 (2016)
9. Davari, M., Mohamed, Y.A.-R.I.: Robust droop and DC-bus voltage control for effective stabilization and power sharing in VSC multiterminal DC grids. *IEEE Trans. Power Electron.* **33**(5), 4373–4395 (2018)
10. Liu, Y., Yang, T., Hao, X., Qi, Y., Li, W.: Data-driven power sharing control within multisource paralleled DAB converters. *IEEE Trans. Transport. Electrif.*, 1 (2025)
11. Ali, M., Ahmad, M., Koondhar, M.A., Akram, M.S., Verma, A., Khan, B.: Maximum power point tracking for grid-connected photovoltaic system using adaptive fuzzy logic controller. *Comput. Electr. Eng.* **110**, 108879 (2023)
12. Arwa, E.O., Folly, K.A.: Reinforcement learning techniques for optimal power control in grid-connected microgrids: a comprehensive review. *IEEE Access* **8**, 208992–209007 (2020)
13. Glavic, M.: (deep) reinforcement learning for electric power system control and related problems: a short review and perspectives. *Annu. Rev. Control.* **48**, 22–35 (2019)
14. Cui, Y., Xu, Y., Li, Y., Wang, Y., Zou, X.: Deep reinforcement learning based optimal energy management of multi-energy microgrids with uncertainties. *CSEE J. Power Energy Syst.*, 1–12 (2024)
15. Mi, Y., et al.: A power sharing strategy for islanded DC microgrid with unmatched line impedance and local load. *Electric Power Syst. Res.* **192**, 106983 (2021)
16. Mahmoud, M.S., Abouheaf, M., Sharaf, A.: Reinforcement learning control approach for autonomous microgrids. *Int. J. Model. Simul.* **41**(1), 1–10 (2021)
17. Guo, C., Wang, X., Zheng, Y., Zhang, F.: Real-time optimal energy management of microgrid with uncertainties based on deep reinforcement learning. *Energy* **238**, 121873 (2022)
18. Wang, C., et al.: Multiagent deep reinforcement learning-based cooperative optimal operation with strong scalability for residential microgrid clusters. *Energy* **314**, 134165 (2025)
19. Derbas, A.A., Bordin, C., Mishra, S., Blaabjerg, F.: ANN-based real-time optimal voltage control in islanded AC microgrids. In: 2024 IEEE 15th International Symposium on Power Electronics for Distributed Generation Systems (PEDG), pp. 1–5 (2024)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.



District Heating, Thermal Systems, and Retrofit Strategies



Digitalization in District Heating: Comparative Insights from Denmark and Sweden on Adoption, Barriers, and Value Creation

Zheng Ma¹  Joy Dalmacio Billanes¹  and Kristina Lygnerud² 

¹ SDU Center for Energy Informatics, The Maersk Mc-Kinney Moller Institute, The Faculty of Engineering, University of Southern Denmark, Odense, Denmark

 {zma, joydbi}@mmsmi.sdu.dk

² Energy, LTH, Lund University, P.O. Box 118, SE-221 00 Lund, Sweden

Abstract. Digitalization holds significant potential to improve the efficiency, flexibility, and sustainability of District Heating (DH) systems. However, the transition toward digitalized DH is often constrained by challenges beyond technology—such as regulatory uncertainty, unclear business models, fragmented organizational capacity, and misaligned stakeholder incentives. This paper offers a comparative, qualitative analysis of how digitalization is adopted, perceived, and implemented in DH systems in Denmark and Sweden—two of the most mature DH markets globally. Based on expert survey data collected under the IEA DHC Annex TS9 initiative, the study explores adoption trends across operational, customer-facing, and strategic domains. Findings reveal that while digital tools like smart meters, Supervisory Control and Data Acquisition (SCADA) systems, and AI-based analytics are being introduced, systemic barriers, such as General Data Protection Regulation (GDPR) related constraints, lack of clear return on investment, and digital skills gaps, continue to slow progress. Denmark's strengths in infrastructure and coordination contrast with Sweden's more diverse but fragmented innovation landscape. The study identifies key enablers: tariff reform, regulatory sandboxes, and capacity building, and offers context-sensitive policy and strategic recommendations. These insights contribute to advancing digitalization in the DH sector and to shaping future research and policymaking in smart energy system transformation.

Keywords: Digitalization · District heating · Denmark · Sweden · Smart energy systems · Energy policy · Barriers · Business model innovation

1 Introduction

DH systems form a critical part of urban energy infrastructure in many countries, particularly in Europe, where they play a key role in providing centralized space heating and domestic hot water in a cost-effective and low-carbon manner [1]. As the energy sector undergoes rapid decarbonization and digital transformation, DH systems are also evolving to meet new expectations for efficiency, flexibility, and sustainability. This evolution is embodied in the concept of Fourth Generation District Heating (4GDH), which

emphasizes low-temperature operation, integration of renewable energy sources, and system-level flexibility [2, 3].

Digitalization is increasingly seen as a strategic enabler in this transition. It promises to enhance system performance through predictive maintenance, real-time optimization, customer-side demand response, and data-driven energy management [4, 5]. Smart meters, sensors, and intelligent control systems offer the potential to reduce thermal losses, optimize grid operations, and improve consumer engagement. Moreover, as urban areas strive to become “smart” and climate-neutral, digitally enhanced DH systems are expected to play a foundational role in integrated energy planning.

However, the digitalization of DH systems remains in its early stages in most countries. While some frontrunners—such as Denmark and Sweden—have introduced advanced digital tools, widespread adoption remains uneven and fragmented. Existing literature and policy reports often focus on individual technical solutions or national case studies, but there remains a lack of comparative, cross-disciplinary analysis that explores how digitalization is perceived, implemented, and challenged in DH systems across different institutional, regulatory, and organizational settings. Moreover, few studies address the interplay between digital tools, business models, and policy environments, or the barriers that prevent digitalization from scaling beyond pilot projects.

This paper aims to fill this gap by providing a comparative, qualitative analysis of the digitalization of DH in Denmark and Sweden—two of the most advanced DH markets globally. While the empirical focus is on these two countries, the broader objective is to generate insights that are relevant to other regions seeking to digitalize their DH systems, including countries with emerging DH markets or those undergoing modernization. The study identifies key enablers and obstacles related to technological integration, organizational capacity, customer engagement, and regulatory frameworks, contributing to the growing field of energy informatics and digital energy system transformation.

The central research questions guiding this study are:

1. What are the current trends in digital tool adoption within DH systems in Denmark and Sweden?
2. What benefits, barriers, and strategic considerations are perceived by DH experts regarding digitalization?
3. How can lessons from these two countries inform broader strategies for digital transformation in the heat sector?

To investigate these questions, the study draws on qualitative data gathered from academic experts in Denmark and Sweden through a structured survey conducted under the IEA DHC Annex TS9 initiative. The analysis focuses on the adoption of digital tools across operational, customer-facing, and strategic domains, while also examining the broader institutional and regulatory conditions shaping digitalization in each national context. By comparing these two advanced DH markets, the study aims to provide transferable insights for other regions considering or advancing digital transitions in the heat sector.

The remainder of the paper is structured as follows: Sect. 2 reviews the state of the art in digitalization of district heating, including key thematic areas; Sect. 3 outlines the methodology and data sources; Sect. 4 presents national-level findings on adoption

trends, perceived benefits, and barriers; Sect. 5 offers a cross-country synthesis and identifies systemic challenges and enabling factors; Sect. 6 discusses the policy, strategic, and organizational implications; and Sect. 7 concludes with recommendations and directions for future research.

2 Literature Review

Digitalization in DH is emerging as a key research and policy topic, particularly in the context of achieving low-carbon, efficient, and consumer-responsive energy systems. This section reviews the state of knowledge across the main dimensions explored in this study: operational performance, customer engagement, data analytics, organizational readiness, and the policy and regulatory environment. The review also identifies research gaps that this paper seeks to address.

2.1 Digitalization and Operational Performance in District Heating

Early digitalization in DH focused on automation and remote monitoring, notably through SCADA systems and Fault Detection and Diagnostics (FDD). These tools improve reliability and thermal efficiency by identifying system anomalies such as improper return temperatures and hidden faults [6, 7]. For instance, smart meter and sensor data in Denmark has enabled detection of customer-level inefficiencies, contributing to operational cost savings and emission reductions [8].

Pilot projects like EnergyLab Nordhavn in Copenhagen have demonstrated how digital control of heat storage and dynamic source switching (e.g., between combined heat and power and electric boilers) can increase system flexibility and align heat production with electricity market signals [9]. Despite promising results, real-time optimization and AI-driven control remain rare in mainstream DH operations.

2.2 Customer Engagement and Demand-Side Digitalization

Digital technologies also offer significant potential on the demand side, including smart thermostats, consumption visualization tools, and feedback-based energy behavior interventions. These tools empower end-users and enable prosumer participation, such as feeding excess heat into the network [10].

Studies have shown that better user control and real-time feedback can enable buildings to operate effectively at reduced supply temperatures, facilitating the transition to 4GDH [11]. However, most implementations are limited to pilot projects, and customer-facing services are often underdeveloped, especially in conservative or municipally owned utilities.

2.3 Data Analytics and Artificial Intelligence in DH

Big data and AI offer new opportunities to optimize DH systems through predictive maintenance, demand forecasting, and network performance analytics. For example, AI models can identify performance issues at the substation level or optimize distribution

temperatures based on weather and user profiles [12]. According to Statista, in 2023, nearly all respondents reported using AI tools, with 77% using Generative AI, 71% using Large Language Models, and only 2.86% not using any of these technologies [13] (as shown in Figure 1).

Moreover, the IEA DHC Annex TS4 project guidebook provides an overview of how digitalization can enhance District Heating and Cooling (DHC) systems [14]. However, adoption is uneven. Some utilities are cautious due to policy restrictions, internal skill gaps, or data privacy concerns. The report [14] notes that even when data is collected, it is not always effectively used, highlighting a gap between data availability and analytical maturity [4].

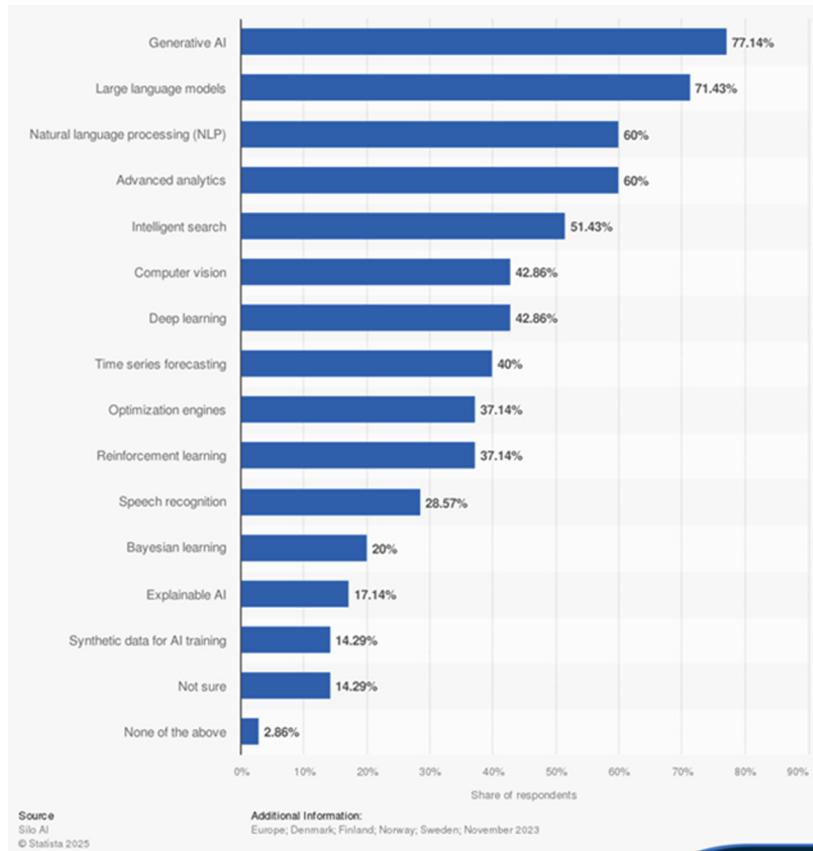


Fig. 1. Artificial Intelligence and the Technologies Used in Nordic Organizations [13].

2.4 Organizational Readiness and Business Model Innovation

Organizational capacity and business model adaptability are critical for realizing the benefits of digitalization. Literature emphasizes the importance of IT/OT integration,

in-house digital skills, and leadership commitment [15]. Yet, many DH companies—particularly smaller or municipally owned ones—lack dedicated digital teams or strategic digitalization roadmaps.

Business model innovation, such as transitioning from heat sales to “heat-as-a-service (HaaS),” is gaining attention but remains at an early stage. HaaS refers to a model where an energy company supplies heat to customers as a service, without selling the equipment or energy directly [16]. A Swedish study of municipal DH firms found that traditional cost-plus pricing models often disincentivize efficiency improvements, while uncertainty about return on investment (ROI) hinders digital experimentation [17].

2.5 Policy, Regulation, and Data Governance

The policy and regulatory environment significantly shape digitalization efforts. While supportive energy and climate policies can accelerate digital innovation, restrictive or unclear regulations—particularly around data privacy—often act as barriers. The EU’s GDPR, while vital for safeguarding user rights, has introduced administrative burdens and ambiguity regarding the permissible use of granular consumption data [18].

Literature also highlights the absence of clear regulatory frameworks for emerging services such as demand response participation or prosumer heat integration, leaving utilities uncertain about investing in enabling technologies.

3 Methodology

This study applied qualitative data from a structured questionnaire developed for the IEA DHC Annex TS9 project on “Digitalization in DH and Cooling”. The survey explored how digital tools create value in DH systems, covering efficiency, flexibility, customer service, risk management, and innovation. It included mostly closed-ended questions, with some open-ended responses. Respondents are consisting of four academic experts (two each from Denmark and Sweden) who are engaged in DH digitalization research.

The analysis maintains full confidentiality, using only anonymized and aggregated data. Selected quotes, edited for clarity and without attribution, illustrate key points. These insights are examined in the discussion to inform future digitalization strategies in DH policy and implementation.

To present the findings, a narrative comparison in the text and summary tables are presented. Table 1 provides an overview of the extent of digitalization in key areas of DH operations as indicated by the respondents, comparing Denmark and Sweden. The legend is also provided to provide a concise, standardized label (e.g., “Moderate Adoption”) for the longer, descriptive phrases found within the table.

Table 1. Perceptions on Digitalization in Key Areas of DH.

Key Areas	Denmark	Sweden
Operations & Maintenance	Moderate Adoption: Some advanced fault detection in place, but not widespread	Limited to Moderate Adoption: Focus on efficiency and return temps growing, but many processes still manual
Energy Management	Limited Adoption: Some pilots on price-based source switching and load shifting, but real-time optimization is rare.	Limited Adoption: Some forecasting tools used, but grid flexibility and dynamic optimization are still early-stage
Customer & Demand Side	Partial Implementation: Some utilities offer apps; prosumer heat feed-in and smart thermostat pilots exist but remain rare	Partial Implementation: Online consumption data is common, but interactive services and customer heat integration are still limited to pilots
Data & Analytics	Mixed Adoption: Some utilities use AI and predictive models; others remain cautious. Data collection is widespread, with GDPR managed internally	Highly Mixed Adoption: Some avoid AI due to policy concerns; others fully adopt it. Data use is moderate overall, with privacy fears limiting use, especially in smaller firms
Strategy & Business	Emerging Impact: Digitalization appears in strategies, but new business models are limited. Some early steps toward “heat-as-a-service”	Mixed Impact: Some use digital tools for new services; others focus on maintaining existing operations

Maturity Level Legend:

- Moderate Adoption: Digital tools are used, but widespread advanced implementation may still be lacking.
- Limited to Moderate Adoption: Some digital tools are implemented, but adoption is not extensive, and many processes may remain manual.
- Limited Adoption: Basic digital tools may be present, but advanced applications, such as real-time optimization or dynamic control, are rare or in early stages.
- Partial Implementation: Certain aspects of digitalization are in place, but comprehensive integration or advanced features are still underdeveloped.

- Mixed Adoption: Varies significantly among different entities within the sector; some adopt extensively, while others remain cautious or avoid certain tools.
- Highly Mixed Adoption: Extreme variability in adoption, with some leading in digital integration and others facing significant constraints.
- Emerging Impact: Digitalization is beginning to appear in strategic plans, but its transformative effect on business models is still limited.
- Mixed Impact: Digital tools are used, but their influence varies from supporting existing operations to exploring new services.

4 National Perspectives on Digitalization in DH: Insights from Denmark and Sweden

This section presents national level insights into the digitalization of DH in Denmark and Sweden, based on expert survey responses. Each sub-section outlines the national context and policy environment, the current state of digital tool adoption, perceived benefits, barriers to implementation, and illustrative examples from practice. Moreover, the National policy timelines for DH Digitalization in Denmark and Sweden is shown in Table 2.

4.1 Denmark

National Context and Policy Landscape. Denmark has one of the most advanced DH systems globally, with over 60% of households connected and a strong commitment to decarbonization. By 2030, 90% of DH energy is expected to come from renewable sources [3]. The Danish energy strategy emphasizes 4GDH, focusing on low-temperature operation, improved efficiency, and integration with renewable electricity [2]. This provides a favorable environment for digital innovation, further supported by national research and demonstration projects such as EnergyLab Nordhavn.

Digital Adoption Across Key Domains. Danish experts reported moderate adoption of digital tools across most operational areas. Fault detection, predictive maintenance, and SCADA systems are used, especially in larger utilities. However, real-time optimization and dynamic source switching (e.g., based on electricity prices) remain limited to pilots.

On the customer side, apps and digital interfaces are partially implemented, and prosumer models—where customers can feed heat back into the grid—exist mostly as experimental projects. Smart thermostat integration and end-user comfort controls have been seen limited to moderate implementation.

In terms of data and analytics, adoption is mixed. While some utilities extensively use AI and predictive models, others remain cautious. Smart meter data is widely available, and GDPR compliance is generally handled internally.

Strategically, digitalization appears in long-term planning documents but has yet to fully influence business model transformation. Concepts such as “heat-as-a-service” are emerging, but traditional heat delivery remains dominant.

Perceived Benefits and Value Creation. Experts identified several benefits of digitalization. Improved data use has enhanced system monitoring, control, and fault detection. For example, the use of “heat boosters” (small, decentralized heat pumps) in

Viborg allowed for lower supply temperatures and reduced return temperatures to 30°C, enabling demand-side flexibility and consumer savings through the “motivation tariff” scheme. The motivation tariff refers to the pricing program in DH that encourage consumers to lower their return temperature [19]. Digital tools also support energy efficiency improvements by identifying substation inefficiencies and reducing heat losses.

Barriers to Adoption and Implementation. Barriers in Denmark are predominantly economic and organizational rather than technical. Experts noted unclear business cases and insufficient alignment of financial incentives across stakeholders. For instance, building owners may not directly benefit from digital upgrades that reduce heat loss, creating a “split incentive” issue. Without clear ROI, many DH companies, especially smaller or non-profit ones, are hesitant to adopt advanced digital technologies. Technical concerns, such as data integration or GDPR compliance, were mentioned but not considered significant impediments due to Denmark’s strong infrastructure.

Notable examples include EnergyLab Nordhavn’s demonstration of smart control in DH networks and the motivation tariff scheme, which encourages consumers to reduce return temperatures. The Viborg case illustrates how combining digital control with hardware innovation can yield financial and environmental benefits.

4.2 Sweden

National Context and Policy Landscape. Sweden’s DH sector expanded rapidly in the 1970s, driven by the oil crisis, environmental concerns, and state support [20]. Today, DH accounts for more than 50% of space heating, especially in urban areas. The Swedish Energy Agency has funded multiple projects focused on smart energy systems, and several utilities have pioneered digital initiatives (e.g., Växjö Energi’s use of AI for demand forecasting [18]). However, the sector includes both large, innovative companies and smaller, risk-averse municipal utilities, leading to uneven digital progress.

Digital Adoption Across Key Domains. Swedish experts reported a similarly moderate but varied level of digital adoption. Fault detection, predictive maintenance, and smart control tools are used to some extent, though implementation is often partial. Supply-demand forecasting and grid flexibility services are at an early stage, and peak load management is largely unimplemented.

Customer-facing platforms are limited. While online consumption dashboards exist, interactive services and prosumer integration are mostly confined to pilots. Smart thermostats have seen some implementation, particularly in trials by Stockholm Exergi.

Data analytics adoption varies significantly. One expert noted full AI integration, while another cited internal restriction preventing its use. These differences reflect divergent digital strategies and capabilities within Swedish utilities.

Perceived Benefits and Value Creation. Swedish experts agreed on the operational value of digitalization. Substation-level analytics have enabled detection of faulty equipment contributing to high return temperatures. Digitalization has also helped standardize operations, improve temperature stability, and support internal coordination. Some utilities are exploring business model innovations, such as building energy management

services and integrated cooling offerings, indicating potential for diversification beyond heat supply.

Barriers to Adoption and Implementation. Barriers in Sweden are primarily regulatory and organizational. GDPR compliance and rigid procurement standards were seen as significant obstacles, especially for smaller firms. One expert criticized the excessive focus on data governance and procedural compliance before any services or use cases are clearly defined. This “process over innovation” culture slows digital progress and deters experimentation.

Organizational conservatism, particularly in municipally owned utilities, also hinders adoption. Without dedicated digitalization teams or strong internal mandates, many companies outsource innovation or wait until technologies are fully proven before adoption. Skills shortages, especially in data science and IT/OT integration further limit progress.

Examples include Växjö Energi’s AI-driven forecasting, Stockholm Exergi’s smart thermostat pilots, and utilities exploring service innovations such as cooling and comfort management. These efforts highlight the potential of digital tools, but also the systemic obstacles that need to be overcome.

Table 2. National Policy Timelines for DH Digitalization.

Country	Policy/Development Area	Description	Relevant Date/Period
Denmark	Renewable Energy Target	90% of DH energy is expected to be renewable	By 2030
	Tariff Innovation	Motivation tariff for behavioral and technological change	Ongoing (existing)
	Demonstration Projects	Projects like EnergyLab Nordhavn demonstrate smart DH control	Ongoing
Sweden	DH Sector Expansion	DH sector rapidly grew due to several factors	1970s
	Digital Initiatives	Swedish Energy Agency has funded projects (e.g., Växjö Energi and Stockholm Exergi)	Ongoing

5 Cross-Country Insights on Digitalization Progress, Challenges, and Enablers

This section synthesizes and compares insights from Denmark and Sweden to provide a broader understanding of how digitalization is progressing in DH, what benefits are being realized, which barriers persist, and what enabling factors can support broader adoption.

5.1 Patterns of Digital Adoption Across Domains

Digitalization across the DH sectors in both countries is advancing incrementally. Experts consistently described implementation as “moderate” in key operational areas, including fault detection, predictive maintenance, and supply-demand forecasting. Smart meters and SCADA systems are in regular use, particularly in Denmark, where infrastructure and data access are well developed. Sweden shows a more fragmented picture, with some utilities adopting AI and advanced analytics, while others lag due to strategic hesitancy or policy constraints.

Notably, real-time optimization, AI-driven automation, and consumer-facing services remain limited to pilot projects or early-stage deployment. Sweden demonstrates a wider disparity among utilities, with some digital frontrunners contrasted by others showing minimal innovation. In Denmark, uptake is more uniformly moderate—even in smaller systems—thanks to strong infrastructure and a centralized policy push.

The demand-side digitalization is still emerging in both countries. While customer dashboards and smart thermostats are being piloted, interactive features and prosumer integration are not widely adopted. This highlights an untapped potential for user-driven efficiency and flexibility services.

5.2 Perceived Benefits and Value Creation

Experts from both countries agree that digitalization brings clear benefits, especially in operational efficiency and cost savings. Benefits highlighted include:

- Improved operational control through fault detection, return temperature management, and smart metering.
- Energy savings and emission reductions via optimization of temperature levels and more dynamic load shifting.
- Data-driven planning and diagnostics that improve asset maintenance and investment decisions.
- Emerging service innovation in Sweden, where some utilities are developing integrated comfort, cooling, or energy management services.

The Danish perspective emphasizes internal efficiency gains—reducing waste, losses, and supply temperatures—within a relatively stable, often non-profit system. In contrast, the Swedish case illustrates how digitalization can support competitive differentiation, particularly in markets facing pressure from heat pumps or private heating providers. These contrasting priorities suggest that digitalization serves both cost-effectiveness and business model diversification, depending on context.

5.3 Barriers to Digital Transformation

While technical tools are available and maturing, both countries face persistent systemic barriers. Table 3 summarizes key barriers, their descriptions, and how prominently they appear in each national context.

Table 3. Key Barriers to Digitalization in DH.

Barrier/Challenge	Description and Context
Unclear Use Cases / Business Case	Difficulty in identifying and quantifying the tangible benefits of digital investments. Many DH companies are unsure how specific digital tools translate to revenue or cost savings.
Misaligned Incentives	The payer for digital improvements may not be the beneficiary, and without mechanisms to share savings, stakeholders lack motivation to engage in digitalization.
Regulatory & Data Privacy Hurdles	Strict regulations (e.g., GDPR) and unclear guidelines on data sharing and flexibility remuneration create uncertainty and slow digital initiatives.
Organizational Conservatism	Many DH organizations resist digitalization due to a preference for proven methods, risk aversion, or outsourcing innovation, delaying adoption.
Digital Skills Gap	Shortage of expertise in data science, IT, and modern control systems within DH companies.
Integration Challenges	Technical barriers in integrating new digital tools with legacy infrastructure, (e.g., lack of standardized data formats, older equipment)

These findings suggest that many of the barriers are non-technical. Economic and institutional factors—such as weak incentive structures or regulatory ambiguity—exert more influence than technology maturity. This is consistent with literature on digital energy transitions, which highlights that governance and organizational culture often determine innovation outcomes more than technological availability.

5.4 Enablers for Scaling Digitalization

Despite the barriers, both countries have several levers that can accelerate digital transformation in DH systems:

- **Strong policy signals:** National climate strategies and international collaborations (e.g., IEA DHC, EU projects) help anchor digitalization in broader policy frameworks.
- **Tariff innovation:** Denmark’s “motivation tariff” is an example of how pricing structures can incentivize behavioral and technological change. Similar mechanisms could reward flexibility or data-driven performance.
- **Demonstration projects:** Real-world pilots, like EnergyLab Nordhavn in Denmark or Växjö Energi in Sweden, illustrate the value of digital tools. Expanding such projects and sharing best practices across the sector can reduce uncertainty and build confidence.
- **Organizational capacity building:** Utilities with internal digital teams or academic partnerships are better positioned to adapt. Support for upskilling and digital workforce development will be critical, particularly for smaller utilities.

- **Regulatory flexibility:** “Sandbox” environments that temporarily relax data-sharing or procurement rules can facilitate innovation without full-scale reform. These could be particularly valuable in overcoming GDPR-related hesitations.
- **Customer-driven momentum:** As customer expectations rise, especially in urban and environmentally conscious markets, utilities may be compelled to offer smarter, more transparent services.

While Denmark and Sweden are both digitalization leaders in the DH sector, their experiences diverge in key respects. Denmark’s progress is rooted in system-level planning, energy efficiency, and strong infrastructure. Sweden’s progress is more uneven but shows promise in service innovation and market responsiveness. These differences underscore the importance of context-sensitive strategies that account for local policy, ownership models, and market dynamics.

6 Policy and Strategic Recommendations for Scaling Digitalization in DH

This section draws on cross-country findings and relevant literature to propose strategic and policy measures for advancing the digital transformation of DH. The analysis emphasizes the need for differentiated but coordinated interventions across regulation, governance, business models, and workforce development. While Denmark and Sweden serve as empirical anchors, the lessons are relevant to a broader international audience navigating the transition toward smarter, low-carbon thermal systems.

6.1 Priorities Emerging from Cross-Country Lessons

The comparative findings from Denmark and Sweden reveal that digitalization in DH is not constrained by technology availability but by a range of systemic, institutional, and economic factors. Both countries demonstrate moderate and uneven adoption across key digital domains, yet the underlying drivers differ.

Denmark’s digital progress is rooted in strong policy alignment and technical infrastructure but limited by organizational conservatism and a lack of clear business cases for many digital investments. In contrast, Sweden shows a more fragmented digital landscape, with some utilities adopting innovative AI-based services while others remain constrained by regulatory caution and resource gaps.

These differences point to the need for context-sensitive strategies. In Denmark, enhancing digitalization requires better alignment of incentives and clearer economic rationales for utilities—especially those operating under non-profit models. In Sweden, overcoming policy rigidity and building internal digital capacity are more urgent. For both countries, and the sector globally, the path forward must prioritize integration of digitalization into long-term decarbonization, customer engagement, and resilience strategies, consistent with the broader principles of 4GDH as defined by [2].

6.2 Policy Measures and Implementation Actions

Public policy must play an enabling role in overcoming the structural barriers identified. A priority is reforming tariff structures. Existing cost-recovery pricing models often disincentivize energy efficiency or flexible heat delivery. In Denmark, the motivation tariff already rewards lower return temperatures; such outcome-based pricing could be expanded to support real-time control, demand response, and prosumer integration. Sweden, where competitive utilities face pressure from heat pumps, would benefit from more dynamic pricing schemes that reflect value-added digital services.

Another critical lever is data governance. While the EU's GDPR provides necessary protections, its strict interpretation in energy contexts—particularly in Sweden—has created procedural burdens that hinder innovation. Governments must clarify how energy utilities can legally and ethically process, share, and aggregate consumption data. Establishing sector-specific guidelines, trusted data platforms, or regulatory sandboxes would help balance privacy with innovation [17].

Continued public investment in digitalization R&D and demonstration projects is essential. Both countries have benefitted from national and EU-funded pilots—such as EnergyLab Nordhavn and Västerås's flexible DH grid—but policy frameworks must now move beyond experimentation to incentivize scale-up. This includes support for replication, dissemination of successful use cases, and the removal of legal or procedural barriers to implementation.

Policy must explicitly support business model innovation. Municipal regulations and ownership structures should be updated to allow utilities to transition from heat-only providers to broader service integrators—offering energy management, smart control, or even cooling. In Sweden, this may also require revisiting competition law to allow constructive cooperation between utilities and technology providers in service innovation, as discussed in the context of sustainability transitions [7].

Moreover, governments and regulatory bodies in both Denmark and Sweden are exploring ways to reduce ambiguity in data governance. Pilot projects and regulatory sandbox approaches, for example, are being tested to allow utilities to trial innovative digital services involving personal consumption data under relaxed compliance conditions. These developments represent practical steps toward reconciling GDPR obligations with the innovation demands of a digitalized DH sector.

6.3 Strengthening Sectoral Readiness and Innovation Capacity

The success of any digitalization strategy depends not only on technology or policy, but on the internal capabilities of DH operators. A recurring theme across expert responses is the shortage of in-house digital skills, particularly in smaller or municipally owned utilities. These organizations often lack staff with expertise in data science, IT/OT integration, cybersecurity, or digital business development.

To address this, governments, universities, and industry associations should jointly support sector-wide capacity building. National training initiatives on digital energy analytics, AI for thermal networks, and cybersecurity compliance can help upskill existing staff. In Denmark, this effort could be aligned with existing energy informatics programs.

In Sweden, partnerships between utilities and technical universities can help overcome resource limitations and encourage in-house learning.

Regulators also have a role to play in supporting innovative culture. A more agile regulatory environment—such as allowing time-bound pilot projects to operate outside full compliance—is needed. Sandbox approaches, in which rules are temporarily relaxed under strict oversight, have been effective in the electricity sector and could be adapted to DH. Such frameworks allow experimentation without risking legal or reputational harm, encouraging utilities to test new services, data-sharing models, or AI tools, as supported by the IEA’s digitalization guidance for the heating sector [4].

Finally, engaging customers as co-creators of value will be vital. As digital tools enable more responsive, transparent, and user-centered heating services, utilities must be prepared to explain, market, and adapt these offerings. Proactive customer education and digital service design—especially in response to heat pump competition—will help ensure long-term engagement and differentiation.

6.4 Broader Applicability and International Relevance

While this study focused on Denmark and Sweden, its findings are applicable to other countries seeking to modernize or expand their DH systems. Many European countries—such as Finland, Germany, Poland, and the UK—face similar challenges in balancing legacy infrastructure with new digital imperatives. In emerging markets like China, where DH expansion is ongoing, integrating digital tools from the outset could avoid future retrofit costs and embed flexibility into system design.

Pan-European instruments such as Horizon Europe and EU cohesion funds can be leveraged to promote cross-border learning, fund technical interoperability standards, and build regional digital infrastructure for thermal energy. The IEA DHC Annex TS9 and TS4 initiatives already provide transnational platforms for exchange; these efforts should be expanded and aligned with policy developments in smart grids, building automation, and urban resilience. As highlighted in recent IEA reports, digitalization must be treated as a foundational pillar for energy system transformation across all vectors, including heat [8].

In general, countries with high DH penetration must recognize that digitalization is not an optional upgrade but a foundational element of future-proofing the heat sector. Without it, low-carbon targets, energy efficiency gains, and consumer-centric services will be increasingly difficult to deliver.

7 Conclusion

This study examined the digitalization of DH systems through a comparative qualitative analysis of Denmark and Sweden, two leading DH markets. Drawing on expert responses and literature, the study assessed the current state of digital adoption, identified perceived benefits, and analyzed structural barriers and enabling conditions. By exploring operational, organizational, and regulatory dimensions, the paper contributes to a more holistic understanding of the systemic factors shaping digital transitions in the heat sector.

Findings show that while digital tools such as smart meters, SCADA systems, and AI-based analytics are gaining traction, implementation remains moderate and uneven. Denmark benefits from centralized coordination and infrastructure, but adoption is hindered by limited business case clarity and incentive misalignment. Sweden shows a broader range of innovation, especially in service offerings, but is constrained by regulatory rigidity and organizational conservatism. Both countries illustrate how non-technical barriers—such as governance structures, data privacy rules, and workforce readiness—often outweigh technological limitations.

These insights suggest that policy frameworks must go beyond supporting technology development to address institutional and market design. Regulators and municipalities can enable progress by modernizing tariff systems, clarifying GDPR-compliant data usage, and creating safe spaces for experimentation through regulatory sandboxes. Strategic investments in workforce development, cross-sectoral partnerships, and customer engagement are also essential to build long-term digital capacity in DH systems.

The study was conducted on small samples with only four academic experts (two from each country). The approach significantly limits the representativeness of findings and may not capture the full diversity of perspectives within the sector. Future study should consider a broader range of stakeholders (e.g., industry practitioners, utility operators, and regulatory stakeholders) to triangulate insights beyond academic perspective. Future research should also examine end-user behavior and perception analysis, particularly concerning digital feedback and smart tariff schemes, addressing consumer response, uptake behavior. Comparative studies across more varied policies and market contexts (e.g., Central and Eastern Europe or East Asia) would also enhance generalizability.

Overall, digitalization is emerging as a critical enabler of smarter, more sustainable, and customer-responsive district heating. While Denmark and Sweden are relatively advanced, their experiences reflect challenges that many DH systems will face globally. Coordinated policy reform, innovative business models, and sustained knowledge exchange will be key to accelerating the transition. With the right conditions, the coming decade could mark a turning point for digital DH as a pillar of the low-carbon urban energy future.

Acknowledgments. This paper is part of the project titled “Danish participation in IEA DHC Annex TS9 – Digitalization of District Heating and Cooling: Improving Efficiency and Performance Through Data Integration”, funded by EUDP (project number: 95-41006-2410289).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Lauersen, B., Moesgaard, R.: District Heating Sustainability in Denmark and its Scandinavian Neighbors. *District Energy Magazine* (2020)

2. Lund, H., et al.: 4th Generation District Heating (4GDH): integrating smart thermal grids into future sustainable energy systems. *Energy* **68**, 1–11 (2014)
3. Leiria, D., et al.: Using data from smart energy meters to gain knowledge about households connected to the district heating network: a Danish case. *Smart Energy* **3**, 100035 (2021)
4. Tunzi, M., et al.: Digitalization of the Demand Side: How Does the End-User's Behavior Influence the Operating Temperatures in DHN? (2024)
5. Magnusson, D.: Swedish district heating – a system in stagnation: current and future trends in the district heating sector. *Energy Policy* **48**, 449–459 (2012)
6. Williamsson, J.: Business model innovation for digitalization in the Swedish district heating sector. *Energies* **16**(21), 7457 (2023)
7. Bolton, R., Hannon, M.: Governing sustainability transitions through business model innovation: towards a systems understanding. *Res. Policy* **45**(9), 1731–1742 (2016)
8. IEA. Digitalisation and Energy (2023). <https://www.iea.org/reports/digitalisation-and-energy>
9. Ribeiro, A.: Achieving IT-OT Integration Emerges as Critical Step for Industrial Efficiency and Security (2024). <https://industrialcyber.co/features/achieving-it-ot-integration-emerges-as-critical-step-for-industrial-efficiency-and-security/>
10. Tunzi, M., et al.: Demand side digitalisation: a methodology using heat cost allocators and energy meters to secure low-temperature operations in existing buildings connected to district heating networks. *Energy* **264**, 126272 (2023)
11. IEA. Annex TS9: Digitalisation of District Heating and Cooling (2024). <https://www.iea-dhc.org/the-research/annexes/2024-2028-annex-ts9>
12. Yang, Q., et al.: Identifying untraced faults associated with high return temperatures from heating systems in buildings connected to district heating networks. *Energy* **309**, 133097 (2024)
13. Statista. Artificial Intelligence and the Technologies Used in 2023 by Nordic Organizations (2025)
14. IEA. Annex TS4: Guidebook for the Digitalisation of District Heating: Transforming Heat Networks for a Sustainable Future (2023)
15. Salite, D., Miao, Y., Turner, E.: A comparative analysis of policies and strategies supporting district heating expansion and decarbonisation in Denmark, Sweden, the Netherlands and the United Kingdom – lessons for slow adopters of district heating. *Environ. Sci. Policy* **161**, 103897 (2024)
16. Galis, D.: Heat-as-a-Service in Action: Insights from Early Renewable Heat Projects (2024). <https://www.wbcsd.org/news/heat-as-a-service-in-action-insights-from-early-renewable-heat-projects/>
17. The General Data Protection Regulation (GDPR). Regulation (EU) 2016/679 2016. https://commission.europa.eu/law/law-topic/data-protection/legal-framework-eu-data-protectio_n_en
18. IEA. Swedish Energy Agency. <https://www.iea.org/articles/swedish-energy-agency>
19. Lund, H., et al.: Fourth-generation district heating and motivation tariffs. *ASME Open J. Eng.* **1** (2022)
20. Fälting, L., Forssell, A., Åberg, M.: Drivers of district heating's dominance in Sweden's urban areas: a historical perspective. *Util. Policy* **92**, 101860 (2025)



Evaluating Retrofit Strategies and Decentralized Systems for the Transition to Low-Temperature District Heating: A Simulation-Based Case Study in Borlänge, Sweden

Vignesh Pechiappan Ayyathurai^(✉)  and Abdelmomen Najmadin

Dalarna University, Borlänge, Sweden

h24vigpe@du.se

Abstract. Transitioning to low-temperature district heating (LTDH) is an imperative approach for decarbonizing energy supply systems in the Nordic region. This study introduces a novel evaluation of scalable renovation strategies and explores the integration of renewable technologies in a real Swedish residential area. Through aggregated urban building energy modeling and statistical historic level inputs, two sets of energy efficiency retrofit scenarios together with the projected climate conditions for 2020, 2050, and 2080 are systematically analyzed to determine their impact on energy demand, peak loads, emissions, and system optimization. Both operational and embodied carbon emissions are quantified, and the techno-economic feasibility of ground source heat pumps and photovoltaic-thermal (PVT) systems is assessed. The results indicate that deep retrofit measures can achieve up to a 35% reduction in peak loads, alongside significant improvements in energy efficiency and solar energy utilization. These findings highlight the value of coupling demand-side renovations with decentralized renewable generation to facilitate the transition towards LTDH implementation and enhance urban climate resilience. This research supports district-level renovation strategies, aligning with Sweden's pathway toward net-zero emissions.

Keywords: Low-Temperature District Heating (LTDH) · Building Stock Renovation · Peak Load Reduction · Decentralized Energy Systems · Operational and Embodied Carbon Emissions · Urban Energy Simulation

1 Introduction

1.1 Context and Motivation

The European building sector is responsible for approximately 40% of final energy consumption and 36% of CO₂ emissions, with residential buildings comprising a significant share of this footprint [1]. In Sweden, district heating (DH) networks dominate thermal energy supply, servicing over 90% of residential units, primarily due to the historical availability of affordable biomass resources, municipal energy planning strategies, and established centralized infrastructure [2]. Historically, these systems have been designed

for high-temperature operation (70–90°C), significantly limiting their compatibility with lower-temperature renewable energy sources and waste heat recovery technologies, thus hindering broader decarbonization and energy efficiency initiatives [3].

As Sweden advances toward its net-zero emissions target by 2045, modernizing existing DH infrastructure is imperative. The transition to LTDH, operating below 55 °C, offers substantial benefits in terms of system efficiency, reduced distribution losses, and enhanced integration of decentralized renewable energy sources, such as solar thermal, shallow geothermal, and sewage heat recovery [4, 5]. Despite substantial energy-saving opportunities, the renovation rate of the existing European building stock remains critically low, frequently below 1% annually. This sluggish pace is predominantly due to demand-side barriers, including limited occupant engagement, split incentives, and insufficient public awareness, alongside supply-side limitations such as constrained contractor capacity and prohibitive upfront investments, collectively impeding critical national emissions reduction targets [6, 7]. To address these challenges, the European Union's Renovation Wave Strategy underscores the need for data-driven, scalable approaches that can accelerate renovation efforts and inform strategic resource planning [8].

1.2 Problem Statement and Research Gap

Current renovation practices predominantly adopt a “bottom-up,” single-building-at-a-time approach, which is resource-intensive and lacks scalability [9]. While detailed, these models fall short in supporting strategic, district-level renovation planning. Conversely, “top-down” simulation tools, capable of aggregating building data and modelling urban-scale energy dynamics [10], present a viable pathway to upscale renovation strategies and optimize resources across building stocks, while taking advantage of statistical historic level construction database. However, despite the recognized potential, there remains limited research explicitly integrating urban-scale, top-down simulation frameworks with decentralized energy optimization strategies such as building-level heat pumps, PVT systems, and shallow geothermal in a cohesive model. Addressing this gap is crucial to enable scalable, data-driven solutions capable of significantly accelerating renovation rates and informing strategic district-level planning decisions. Moreover, the coupling study of energy demand reduction, resource allocation, and climate impact at the district level is underexplored, especially in the context of Nordic residential areas.

1.3 Research Objectives

This research specifically focuses on the Rymdgatan residential district in Borlänge, Sweden, consisting of ten multi-family buildings constructed in the 1990s. These buildings typify Swedish residential stock of this era, characterized by suboptimal thermal envelopes and high dependence on legacy high-temperature DH systems. The district serves as a strategically representative case, providing valuable insights and scalable examples for similarly constructed residential areas across Nordic climates. Using statistical historic level inputs, the research simulates a range of retrofit strategies

including light and deep renovations and evaluates the performance of decentralized optimization measures such as heat pumps, PVT systems, and geothermal energy, within

the context of a LTDH transition. By addressing existing methodological and practical gaps, this study develops a calibrated, aggregated urban building energy model for this representative Swedish residential area to evaluate district-scale renovation strategies concerning the following objectives:

1. Quantify seasonal and annual energy demand profiles across multiple retrofit scenarios, assessing variations in hourly load patterns, peak demand magnitudes, and their implications for LTDH system compatibility.
2. Evaluate emissions impacts by comparing operational CO₂ reductions achieved through retrofit measures and decentralized supply integration, alongside embodied carbon associated with material use and technology deployment.
3. Assess solar energy utilization potential at the building and district level, examining PVT contributions to self-sufficiency, self-consumption, and temporal load balancing under different technological and climatic scenarios.
4. Analyse the performance and feasibility of decentralized energy systems, including ground-source heat pumps, PVT, and shallow geothermal technologies, focusing on their role in peak load mitigation, supply-demand matching, and long-term techno-economic viability in support of LTDH transitions.

2 Literature Review and Theoretical Context

District heating systems have long been central to thermal energy supply in Nordic countries, with over 90% of Swedish multi-family residences connected to such networks [2]. Traditionally, these systems operate at high supply temperatures (70–90°C), which restricts the integration of low-grade renewable sources such as waste heat and solar thermal energy [3]. To address these limitations and improve efficiency, the sector is transitioning to 4th and 5th Generation District Heating and Cooling (4GDH/5GDHC). These next-generation systems operate at supply temperatures below 55°C, allowing for thermal energy exchange through decentralized, bidirectional flows and enhancing compatibility with solar, geothermal, and heat recovery technologies [4, 5, 11]. A key principle underlying LTDH is the creation of energy synergies, wherein the interplay between supply-side systems (e.g., ground-source heat pumps, PVT collectors, and solar thermal) and demand-side efficiencies (e.g., envelope upgrades, mechanical ventilation) enhances performance, resilience, and carbon mitigation [12, 13]. However, dynamic modelling of these synergies at the district scale remains methodologically challenging due to thermal inertia, spatial heterogeneity, and fluctuating user behaviour, requiring advanced multi-objective optimization frameworks [14].

In the Swedish context, the relevance of these challenges is underscored by the characteristics of the residential building stock, particularly post-war developments like the Million Programme housing, which are marked by poor insulation and high thermal loads [6]. Although retrofit potential is high, the renovation rate remains under 1% annually insufficient to meet national climate goals [8]. Conventional renovation approaches largely follow bottom-up methods focused on individual buildings, which lack scalability, are resource-intensive, and fail to support system-level planning [9]. In response, the European Renovation Wave strategy calls for top-down, data-driven methodologies capable of evaluating impacts across entire building stocks [15].

While tools like SimStadt and UrbanOpt also support urban-scale modeling, CEA offers integrated LCA modules, hourly resolution, and GIS coupling, making it more adaptable for LTDH scenario planning in Nordic contexts [16]. Urban Building Energy Models (UBEMs), such as the City Energy Analyst (CEA), have emerged as valuable tools to bridge this methodological gap. These platforms simulate energy flows, retrofit effects, and supply system configurations at urban scales [10, 17–19]. By enabling scenario analysis, capital investment estimation, and the development of spatially informed decarbonization pathways, UBEMs provide a practical pathway application, particularly in medium-density districts typical of Sweden. As it remains limited in the current literature, it is highly relevant for advancing both research and practice in this context. Building on these insights, this study employs a calibrated top-down UBEM approach to investigate the energy, emissions, and economic implications of retrofit and decentralized supply scenarios for a representative cold-climate Swedish residential district.

3 Methodology

3.1 Case Study Area and Simulation Platform

This study employs a simulation-based approach to assess retrofit strategies and decentralized energy configurations for 10 building blocks in Rymdgatan residential district in Borlänge, Sweden. The representative multi-family buildings constructed in the early 1990s that being characterized by moderate insulation levels and reliance on high-temperature district heating (HTDH). Building energy performance was modelled in a top-down approach. The CEA tool as one of the most well recognized top-down UBEM is employed. Due to the characteristics of aggregated data and socio-economic drivers to simplify complexity, reduce data intensity, and enhance scalability, it allows for efficient analysis and planning across large urban areas without the need for detailed, building-specific data collection, making it more practical for large-scale renovation projects [10, 19]. Moreover, CEA's hourly resolution, integrated life-cycle carbon modelling, and GIS compatibility make it suitable for district-level scenario planning under evolving climatic and policy contexts.

To enhance efficiency and achieve more realistic inputs, the model is fed with statistical historic level inputs that incorporate geometry from cadastral GIS datasets, envelope properties in accordance with TABULA(European typology database used for archetype modelling) [20] recommendations for Swedish multi-family houses (1990– 1995), and occupancy/internal load profiles based on established Swedish residential

schedules. Weather data were obtained from SMHI (Sveriges meteorologiska och hydrologiska institut) using Typical Meteorological Year (TMY) datasets for Borlänge. Two retrofit levels light and deep were applied across time milestones (2020, 2050, 2080), including envelope enhancements, ventilation upgrades, and integration of onsite renewable systems (PVT, solar thermal, and shallow geothermal probes).

The developed CEA model was calibrated using measured energy data for 2018. The actual annual heating demand for the district was 672.71 MWh/year, compared to the simulated baseline value of 692.20 MWh/year. A CV (RMSE) (Coefficient of Variation of the Root Mean Squared Error) of 2.89% was achieved well below the 15% benchmark for annual calibration accuracy demonstrating strong agreement and

ensuring model robustness for scenario-based analysis [7]. No manual tuning beyond standardized parameters was applied, supporting the transferability of the modelling approach to other urban districts.

3.2 Simulation Framework and Scenario Development

The simulation framework consisted of a baseline scenario (HTDH, no retrofit) and six retrofit scenarios: 2020 LR/DR, 2050 LR/DR, and 2080 LR/DR. Each combined envelope interventions with varying degrees of decentralized supply integration, including ground-source heat pumps (GSHPs), PVT collectors, and solar thermal systems. Centralized optimization was excluded to reflect a realistic decentralized LTDH transition. Figure 1 illustrates the methodological workflow. The methodology for this study began with comprehensive data collection and input preparation. Occupancy and schedule data were collected to capture occupant loads, daily and seasonal schedules, appliance usage, lighting, domestic hot water consumption, and internal electricity use. Then, the collected data underwent preprocessing and validation. The next phase involved scenario development, where two sets of energy efficiency retrofit scenarios were formulated to represent both shallow and deep renovation strategies, focusing on building envelope improvements, upgrades to heating and ventilation systems, and the integration limits for photovoltaic-thermal (PVT) systems. In addition, weather and occupancy data for the years 2020, 2050, and 2080 were integrated to assess the impact of projected climate conditions on building energy performance. Overall, Urban building energy modeling simulated district-level energy demand, peak loads, and system behavior. Ground source heat pumps and photovoltaic-thermal systems were assessed for integration. Both operational and embodied carbon emissions were quantified. Results were visualized, reporting energy demand, peak load reduction, emissions, and techno-economic metrics to support low-temperature district heating transition and net-zero goals. Table 1 summarizes retrofit characteristics and technology assumptions per scenario. At the end, the overall performance was evaluated across four dimensions: energy demand (seasonal/annual load profiles and peak demand), emissions (operational CO₂ emissions and embodied carbon based on life-cycle assessment principles), techno-economics (CAPEX, OPEX, and TAC), and renewable integration (self-sufficiency and self-consumption ratios of distributed systems).

Cost and performance projections for PVT and GSHP were based on IEA Technology Roadmap (2023), assuming 20% cost decline and +15% efficiency increase by 2050 compared to 2020 levels. Internal loads were based on standard occupancy profiles from Sveby (2020) for Swedish multi-family houses: 2.4 persons/unit, with 7–9 kWh/m²/year appliance loads and 10 W/m² lighting gains.

4 Results and Discussion

4.1 Seasonal Energy Demand Profiles

Figure 2 illustrates average daily energy demand for typical winter (January) and summer (July) across all scenarios. In winter, the baseline has a peak load of 354.6 kWh, which was reduced by 42% in the 2080 Deep Retrofit (DR) scenario, demonstrating both the

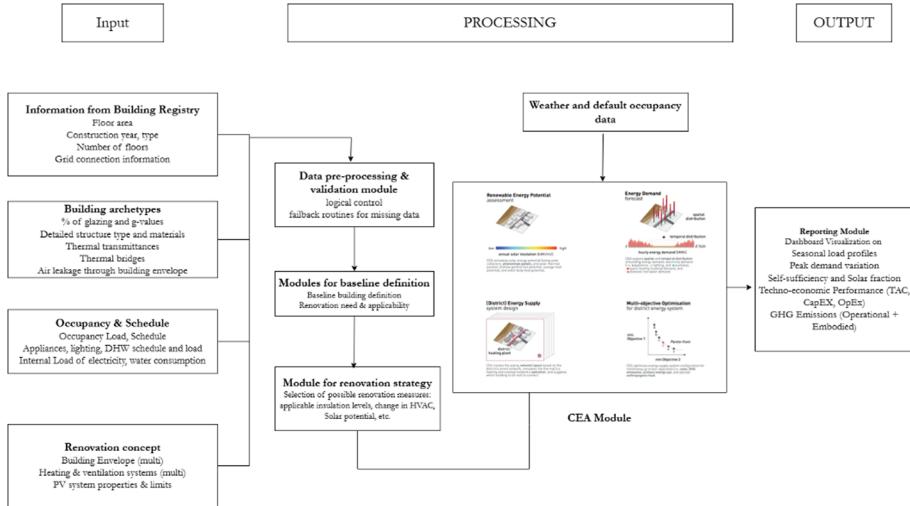


Fig. 1. Methodology Workflow

Table 1. Scenario specifications, including retrofit levels and technology integrations

Scenario	Envelope Improvements	HVAC System Modifications	Technology Integration
2020 LR	Windows: 0.9 W/m ² K Roof: 0.08 W/m ² K Walls: 0.16 W/m ² K Floors: 0.19 W/m ² K	Heating: 70/55°C DHW: 45/10°C	Flat plate solar collectors (SC: FP) PVT: Monocrystalline silicon module
2020 DR	Windows: 0.76 W/m ² K Roof: 0.05 W/m ² K Walls: 0.07 W/m ² K Floors: 0.19 W/m ² K	Heating: 70/55°C DHW: 45/10°C	Flat plate solar collectors (SC: FP) PVT: Monocrystalline silicon module
2050 LR	Windows: 0.9 W/m ² K Roof: 0.08 W/m ² K Walls: 0.16 W/m ² K Floors: 0.19 W/m ² K	Heating: 40/35°C DHW: 35/10°C optimum control	Flat plate solar collectors PVT: Multi-crystalline silicon module
2050 DR	Windows: 0.76 W/m ² K Roof: 0.05 W/m ² K Walls: 0.07 W/m ² K Floors: 0.19 W/m ² K	Heating: 40/35°C DHW: 35/10°C optimum control	Flat plate solar collectors PVT: Multi-crystalline silicon module
2080 LR	Windows: 0.9 W/m ² K Roof: 0.08 W/m ² K Walls: 0.16 W/m ² K Floors: 0.19 W/m ² K	Heating: 40/35°C DHW: 35/10°C Economizer & optimum control	Evacuated tube solar collectors (SC: ET) PVT: Copper indium gallium diselenide solar module
2080 DR	Windows: 0.76 W/m ² K Roof: 0.05 W/m ² K Walls: 0.07 W/m ² K Floors: 0.19 W/m ² K	Heating: 40/35°C DHW: 35/10°C Economizer & optimum control	Evacuated tube solar collectors (SC: ET) PVT: Copper indium gallium diselenide solar module

effectiveness of envelope improvements and system upgrades and the future climate influence. Light retrofits (LR) show moderate reductions, with noticeable performance gaps compared to the DR configurations. In summer, energy demand is largely driven by domestic hot water (DHW) and base loads. The baseline summer peak of 79.2 kWh

drops to 42.1 kWh in 2080 DR. Although the overall load is lower, the DRs still half peak values, confirming their year-round impact. Some early stage retrofit scenarios (e.g., 2020 LR) present minor summer anomalies due to thermal inertia shifts and DHW system lag, reinforcing the need for integrated control strategies during transition phases.

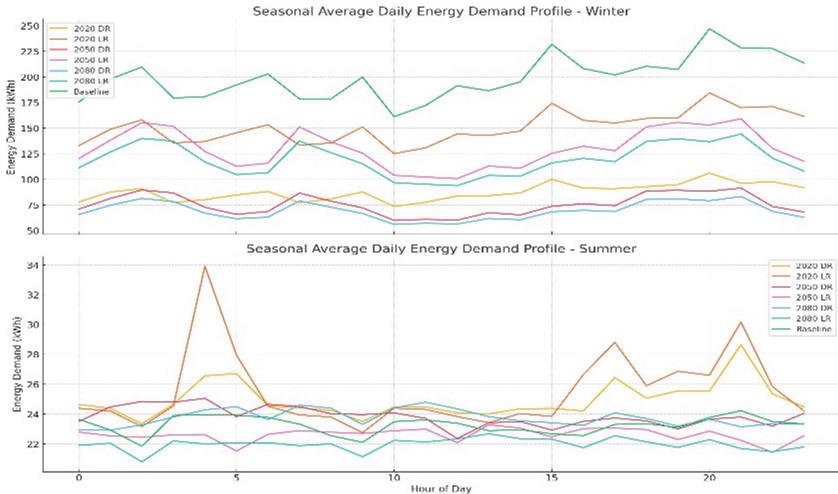


Fig. 2. Average daily energy demands profiles for winter (January) and summer (July) across the Baseline and retrofit scenarios.

4.2 Emissions Impact: Operational and Embodied Carbon

Operational carbon emissions decline by 50% across scenarios from 61.08 tons CO₂ in 2020 LR to 30.74 tons in 2080 DR while per-area emissions fell from 13.8 to 6.95 kg CO₂/m². District heating for DHW remains the largest source of emissions, increasingly dominant as other end uses were decarbonized. Grid electricity emissions remain marginal (2–4.5%), with reductions constrained by energy mix and system operation (Fig. 3).

Embodied emissions, constant across scenarios due to standardized material assumptions, were 57.23 tons CO₂/year (12.99 kg CO₂/m²/year). This reinforces the long-term carbon benefit of operational savings, justifying deep retrofits even with higher upfront material impacts (Fig. 4).

4.3 Solar Energy Utilization Potential

Annual solar energy yield varies substantially with technology and scenario. While the baseline offered 97,850 kWh/year, peak generation was achieved in 2050 DR (240,057 kWh/year) due to expanded solar thermal and PVT deployment. In contrast, 2080 scenarios experienced a reduced yield (\approx 124,000 kWh) due to a different expanded deployment of Copper Indium Gallium Selenide (CIGS)-based PVT modules. Multi-crystalline PVT

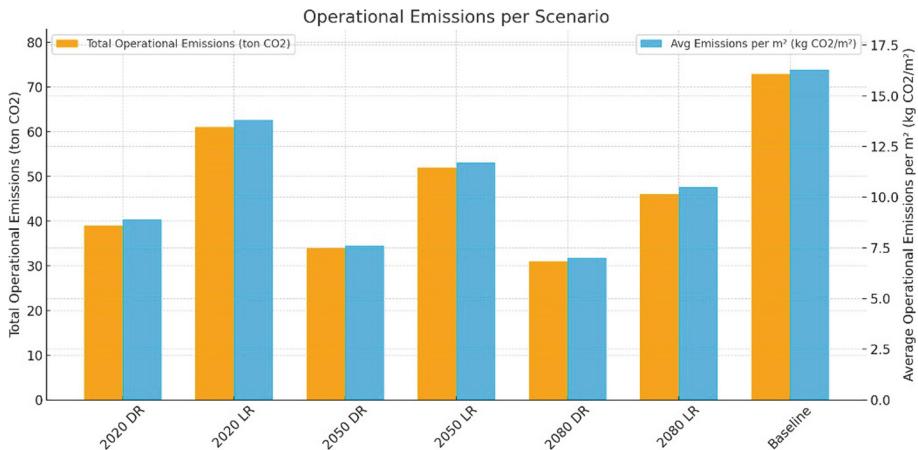


Fig. 3. Total and average operational carbon emissions across Baseline and retrofit scenarios

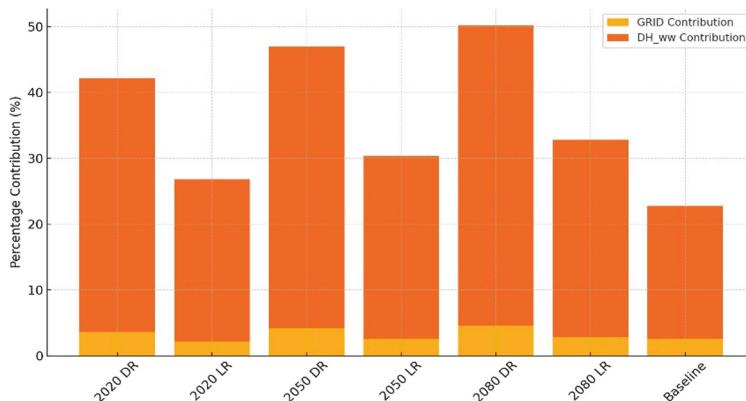


Fig. 4. Contribution of grid electricity and district heating (DHW) to total operational emissions per scenario.

modules show higher efficiency at moderate temperatures and longer lifespan in northern climates compared to CIGS, which are more sensitive to thermal degradation. This aligns with recent findings in Völzel et al., [21]. The selection of technology plays a critical role with one finding that multi-crystalline modules in 2050 scenarios outperform CIGS in 2080 by >30,000 kWh. This highlights the need for strategic selection of future-proof solar technologies aligned with system and climate targets (Fig. 5).

4.4 Decentralized Supply System Performance and Optimization

This analysis investigates how decentralised heating supply systems evolve across energy retrofit scenarios. It focuses on three core metrics: heating system capacity composition, annualized economic performance, and operational greenhouse gas (GHG) emissions.

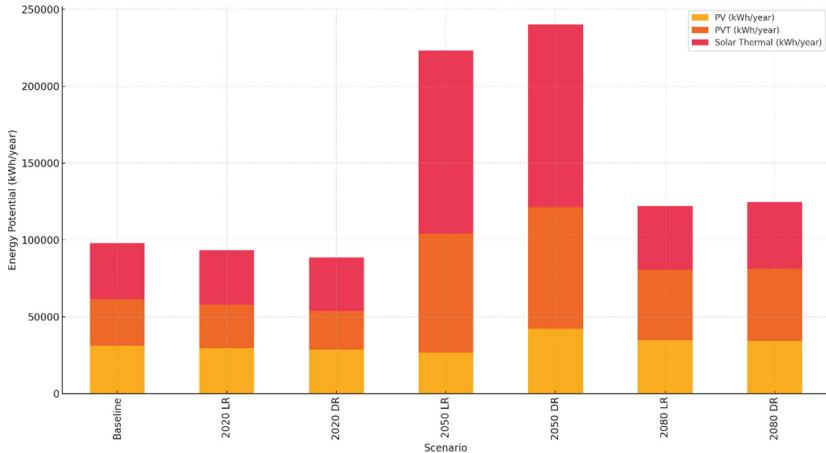


Fig. 5. Solar energy utilization potential across all scenarios from PV, PVT, and solar thermal systems.

All results are derived from optimized configurations for each building, representing the best-performing supply system per scenario (Fig. 6).

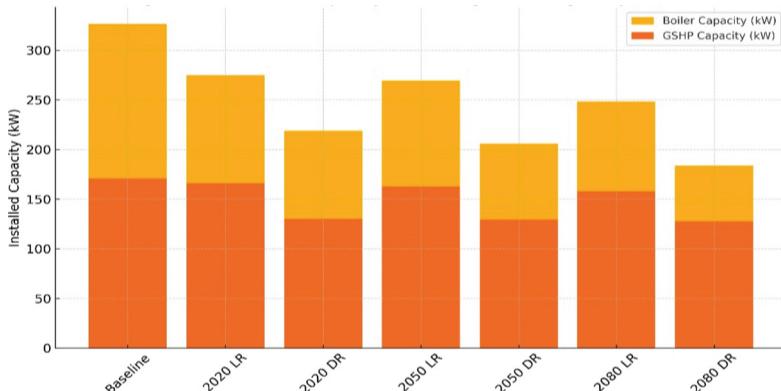


Fig. 6. Installed capacity of boilers and GSHPs by scenario

Regarding technology transition and capacity planning, as shown in Fig. 6, the transition to decentralized systems yields substantial benefits in capacity planning, emissions, and lifecycle cost. GSHPs progressively replaced gas boilers, which declined from 156 kW in the baseline to <100 kW in retrofit scenarios. No fuel cells were selected in any case, due to unfavorable techno-economic parameters.

Regarding peak heating load reductions across scenarios, one of the critical indicators of decentralized system effectiveness in LTDH transitions is the reduction in peak heating loads at the building level. As illustrated in Fig. 7, across all buildings, peak heating loads fell by up to 35% especially in energy-intensive buildings like B1009 and B1010. This

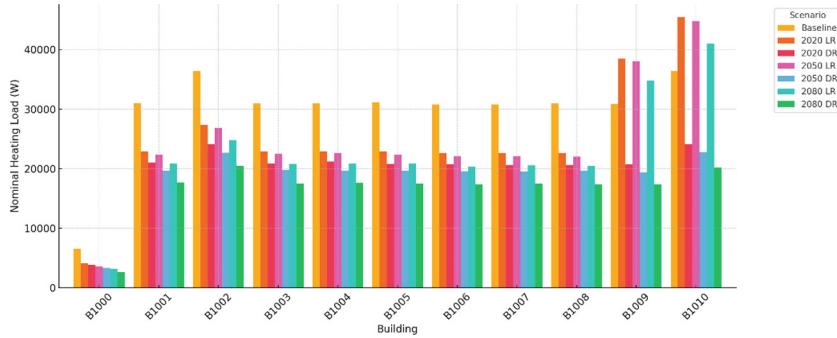


Fig. 7. Peak heating load per building across scenarios, showing significant reductions under retrofit and decentralized supply configurations

supports system downsizing and enhances LTDH compatibility. Such reductions reduce distribution infrastructure needs and seasonal thermal stress.

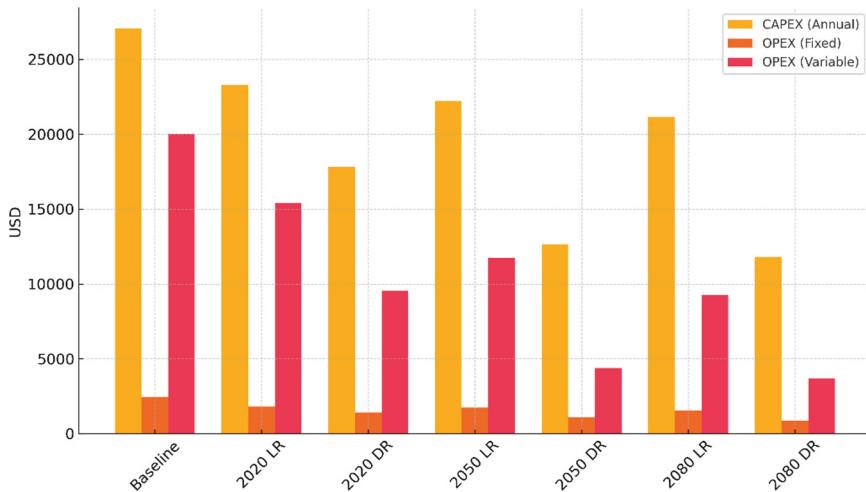


Fig. 8. Annualized CAPEX, fixed and variable OPEX per scenario

Regarding cost efficiency and economic trends, the evolution in system configuration has a direct impact on lifecycle costs Fig. 8. Total Annualized Cost (TAC) fell from USD 50,000 (baseline) to under USD 18,200 in 2050 DR. Cost reductions stemmed from smaller system sizing, reduced variable OPEX, and improved control strategies. The 2080 scenarios retained sub-USD 20,000 TAC, validating future-ready configurations.

Regarding emissions performance, the most dramatic change is observed in operational GHG emissions. As depicted in Fig. 9, emissions drop from 17.2 tons CO₂/year in the Baseline to just 1.8–1.9 tons in the 2050 and 2080 DR scenarios. This nearly 90% reduction is attributed to the significant reduction in boiler usage and increased deployment of GSHPs, which are expected to operate using a decarbonizing electricity

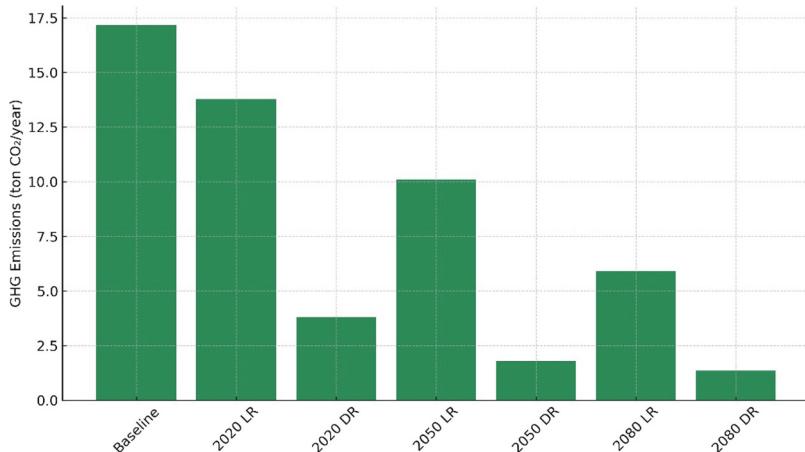


Fig. 9. Operational GHG emissions across all scenarios

grid in future timelines. Importantly, the plateau between 2050 and 2080 scenarios indicates that emissions reductions reach an effective minimum by mid-century, affirming the sufficiency of 2050-level decarbonisation strategies under the modeled conditions.

The findings confirm that decentralised, electrified heating systems can offer substantial environmental and economic benefits when combined with building retrofits. While GSHPs emerge as the technology of choice, natural gas boilers remain useful as backup or peak-load systems. System optimization enables cost-effective solutions that maintain performance while significantly reducing operational carbon. The marginal gains between 2050 and 2080 reflect diminishing returns as systems approach their optimal efficiency and sustainability limits.

5 Conclusion

This study presents a district-scale evaluation of renovation strategies and decentralized energy systems for a representative residential area in Borlänge, Sweden. Using the CEA, seven scenarios were simulated across three timeframes (2020, 2050, 2080), examining their impacts on energy demand, emissions, and techno-economic performance under a LTDH transition pathway. The results demonstrate that deep retrofit strategies, particularly when combined with decentralized supply systems such as GSHPs and PVT modules, significantly reduce both energy consumption and operational carbon emissions. The 2080 DR scenario achieved a 42% reduction in peak winter heating demand and 47% reduction in summer base loads, affirming the effectiveness of advanced envelope measures in shaping seasonal demand profiles.

While this study emphasizes decentralized LTDH systems, centralized LTDH networks using biomass or waste heat may offer economies of scale. However, their rigidity limits responsiveness to future load shifts, making decentralised systems more future-proof in mixed-use districts [11]. Operational CO₂ emissions declined by up to 50%, while peak heating loads were reduced by 35% in high-demand buildings, supporting

LTDH compatibility and infrastructure downsizing. Total annualized system costs are estimated with a drop over 60% in 2050 DR, underscoring the long-term economic viability of electrified, retrofit-integrated solutions. Solar energy contribution is projected to peak in 2050 DR, highlighting the influence of technology choice on renewable yield. While the findings are district-specific, the modeling approach based on statistical archetypes, standardized building physics (TABULA), and climate zones makes the framework transferable to other mid-density districts in cold climates.

Deep retrofits, while beneficial long-term, face social feasibility barriers such as tenant relocation, split incentives in rental housing, and upfront capital constraints. Future work must incorporate stakeholder analysis and co-benefits to improve uptake. By integrating calibrated urban energy simulations with life-cycle emissions and cost assessments, this study provides robust, scalable evidence for policymakers and planners. The findings confirm that mid-century deep renovations, when coupled with distributed energy supply, can meet decarbonization goals without waiting for long-term technological shifts. The calibrated modeling workflow developed herein is transferable to other Nordic cities with similar mid-90s building stock and DH legacy systems.

While this study has primarily focused on energy demand, emissions, and technoeconomic outcomes, it is important to recognize that occupant comfort and resilience will become increasingly significant as district heating systems transition to lower temperatures. Future research should therefore include a detailed assessment of indoor thermal comfort, air quality, and the ability of renovated buildings to maintain healthy and comfortable conditions during extreme weather events or system disruptions. By integrating occupant-centric metrics and resilience indicators into scenario analyses, future work can ensure that renovation and renewable integration strategies not only achieve energy and climate targets but also support the well-being and adaptability of residents in a changing climate. In addition, it should include a comparative analysis of centralized systems alongside decentralized configurations to evaluate their performance under LTDH conditions. Incorporating higher-resolution calibration data, dynamic control strategies, and thermal storage will improve model accuracy and flexibility assessment. Multi-objective optimization considering cost, comfort, and embodied emissions will further enhance the robustness of district-level renovation planning.

6 Appendices

Appendix A. Retrofit Scenario Specifications

Building Element	Baseline U-value / Spec	Light Retrofit	Deep Retrofit
External Wall	0.35 W/m ² K	0.16 W/m ² K	0.07 W/m ² K
Window	2.88 W/m ² K	0.9 W/m ² K	0.16 W/m ² K
Roof	0.20 W/m ² K	0.08 W/m ² K	0.05 W/m ² K
Floor	0.20 W/m ² K	0.19 W/m ² K	0.19 W/m ² K
Air Leakage (n50)	1.0 ach	0.7 ach	0.4 ach

(continued)

(continued)

Building Element	Baseline U-value / Spec	Light Retrofit	Deep Retrofit
Ventilation System	Mech. Vent. (no heat recovery)	Mech. Vent. (no HRV)	Mech. Vent. With HRV

Appendix B. Peak Heating Load per Building Across Scenarios

Building	Peak Load (Baseline) [W]	Peak Load (2020 LR) [W]	Peak Load (2020 DR) [W]	Peak Load (2050 LR) [W]	Peak Load (2050 DR) [W]	Peak Load (2080 LR) [W]	Peak Load (2080 DR) [W]
B1000	6577.87	4133.73	3891.01	3595.62	3390.69	3223.45	2650.23
B1001	31023.27	22892.36	21051.54	22362.95	19642.14	20936.91	17691.16
B1002	36421.2	27387.74	24121.03	26878.14	22715.74	24794.28	20438.61
B1003	30956.61	22890.77	20881.47	22520.79	19829.65	20828.17	17562.41
B1004	30958.37	22946.23	21215.52	22627.92	19697.33	20882.56	17613.96
B1005	31091.66	22886.00	20805.51	22415.08	19669.74	20879.65	17562.41
B1006	30824.98	22652.84	20772.83	22147.43	19590.88	20351.22	17413.87
B1007	30824.98	22599.75	20604.34	22091.34	19563.76	20567.53	17549.87
B1008	30958.37	22656.03	20604.34	22041.45	19669.44	20460.16	17363.26
B1009	30828.51	38513.62	20689.09	38041.85	19460.63	34766.26	17416.29
B1010	36421.20	45413.16	24126.60	44748.22	22767.25	40983.07	20254.12

Appendix C. Cost and Emission Factors

Parameter	Value	Source/Note
CAPEX – Ground Source Heat Pump (GSHP)	1,500 USD/kW	CEA Default + Adjusted from Results
CAPEX – Photovoltaic-Thermal (PVT) System	1,000 USD/m ²	Assumed from Scenario Inputs (2050 vs 2080 tech)
OPEX – Fixed Annual Cost (GSHP + PVT)	5% of CAPEX	Standard industry assumption (5%)
OPEX – Variable Cost per kWh (Grid/DH)	0.10 USD/kWh	Used in OPEX_var_USD column of optimization results
Discount Rate (TAC Calculation)	3% annually	Applied in TAC_USD calculation
Emission Factor – District Heating	0.079 kg CO ₂ /kWh	Used in Sect. 4.2 (Emissions Impact)

(continued)

(continued)

Parameter	Value	Source>Note
Emission Factor – Grid Electricity	0.012 kg CO ₂ /kWh	Low-carbon Swedish grid, as used in model
Embodied Carbon – Retrofit Materials	30–60 kg CO ₂ /m ²	Estimated from retrofit depth and OneClick LCA profiles

Appendix D. Model Calibration Summary

Calibration Parameter	Value	Notes
Measured Annual Heating Demand (2018)	672.71 MWh/year	Actualenergydata from utility records
SimulatedBaseline Heating Demand	692.20 MWh/year	Baseline CEA model output for existing conditions
Calibration Error (Δ kWh/year)	19.49 MWh/year	Absolutedifference between measured and simulated
Coefficient of Variation of RMSE (CV(RMSE))	2.89%	Used to validate model accuracy
Calibration Threshold Benchmark	< 15%(ASHRAE Guideline 14)	Acceptable for annual model calibration

Acknowledgment. The authors would like to appreciate the financial support from the Swedish Energy Agency through the Viable Cities' international program (ENVISION CHANGE project P2023-01559, and PED-ACT project P2022-01000), which are respectively part of the ERA-NET Cofund Urban Transformation Capacities (ENUTC) framework: Building Transformation Capacity Through Arts and Design (BTC-ENUTC), and the Joint Programming Initiative (JPI) Urban Europe framework.

References

1. World Energy Outlook 2023
2. Werner, S.: International review of district heating and cooling. *Energy* **137**, 617–631 (2017). <https://doi.org/10.1016/j.energy.2017.04.045>
3. Averfalk, H., Ingvarsson, P., Persson, U., Gong, M., Werner, S.: Large heat pumps in Swedish district heating systems. *Renew. Sustain. Energy Rev.* **79**, 1275–1284 (2017). <https://doi.org/10.1016/j.rser.2017.05.135>
4. Schmidt, D., Lygnerud, K., Werner, S., Geyer, R., Schrammel, H., Østergaard, D.S., Gudmundsson, O.: Successful implementation of low temperature district heating case studies. *Energy Rep.* **7**, 483–490 (2021). <https://doi.org/10.1016/j.egyr.2021.08.079>
5. Buffa, S., Cozzini, M., D'Antoni, M., Baratieri, M., Fedrizzi, R.: 5th generation district heating and cooling systems: a review of existing cases in Europe. *Renew. Sustain. Energy Rev.* **104**, 504–522 (2019). <https://doi.org/10.1016/j.rser.2018.12.059>

6. Mata, É., Sasic Kalagasisidis, A., Johnsson, F.: Cost-effective retrofitting of Swedish residential buildings: effects of energy price developments and discount rates. *Energ. Effi.* **8**, 223–237 (2015). <https://doi.org/10.1007/s12053-014-9287-1>
7. Ekström, T.: Predicting the energy performance of buildings: a method using probabilistic risk analysis for data-driven decision-support. Division of Building Physics, Department of Building and Environmental Technology, Faculty of Engineering, Lund University, Lund (2011)
8. A Renovation Wave for Europe - greening our buildings, creating jobs, improving lives
9. Johari, F., Lindberg, O., Ramadhan, U.H., Shadram, F., Munkhammar, J., Widén, J.: Analysis of large-scale energy retrofit of residential buildings and their impact on the electricity grid using a validated UBEM. *Appl. Energy* **361**, 122937 (2024). <https://doi.org/10.1016/j.apenergy.2024.122937>
10. Fonseca, J.A., Nguyen, T.-A., Schlueter, A., Marechal, F.: City Energy Analyst (CEA): integrated framework for analysis and optimization of building energy systems in neighborhoods and city districts. *Energy Build.* **113**, 202–226 (2016). <https://doi.org/10.1016/j.enbuild.2015.11.055>
11. Lund, H., et al.: 4th Generation District Heating (4GDH). *Energy* **68**, 1–11 (2014). <https://doi.org/10.1016/j.energy.2014.02.089>
12. Gjoka, K., Rismamchi, B., Crawford, R.H.: Fifth-generation district heating and cooling systems: a review of recent advancements and implementation barriers. *Renewable Sustainable Energy Rev.* **171**, 112997 (2023). <https://doi.org/10.1016/j.rser.2022.112997>
13. Bilardo, M., Sandrone, F., Zanzottera, G., Fabrizio, E.: Modelling a fifth-generation bidirectional low temperature district heating and cooling (5GDHC) network for nearly Zero Energy District (nZED). *Energy Rep.* **7**, 8390–8405 (2021). <https://doi.org/10.1016/j.egyr.2021.04.054>
14. Abugabbara, M.: District heating and cooling systems transition: Evaluation of current challenges and future possibilities. LTH, Lund University, Lund, Division of Building Services (2023)
15. Tsemekidi-Tzeiranaki, S., Paci, D., Clementi, E., Gonzalez Torres, M.: Analysis of the reports on 2020 targets under Article 27 of the Governance Regulation: Energy Efficiency. Publications Office of the European Union, Luxembourg (2022). <https://doi.org/10.2760/27622>
16. Dabirian, S., Panchabikesan, K., Eicker, U.: Occupant-centric urban building energy modeling: approaches, inputs, and data sources A review. *Energy Build.* **257**, 111809 (2022). <https://doi.org/10.1016/j.enbuild.2021.111809>
17. Ali, U., Shamsi, M.H., Hoare, C., Mangina, E., O'Donnell, J.: Review of urban building energy modeling (UBEM) approaches, methods and tools using qualitative and quantitative analysis. *Energy and Buildings* **246**, 111073 (2021). <https://doi.org/10.1016/j.enbuild.2021.111073>
18. Fonseca, J.A., Schlueter, A.: Integrated model for characterization of spatiotemporal building energy consumption patterns in neighborhoods and city districts. *Appl. Energy* **142**, 247–265 (2015). <https://doi.org/10.1016/j.apenergy.2014.12.068>
19. Ferrando, M., Causone, F., Hong, T., Chen, Y.: Urban building energy modeling (UBEM) tools: a state-of-the-art review of bottom-up physics-based approaches. *Sustain. Cities Soc.* **62**, 102408 (2020). <https://doi.org/10.1016/j.scs.2020.102408>
20. Loga, T., Stein, B., Diefenbach, N.: TABULA building typologies in 20 European countries— Making energy-related features of residential building stocks comparable. *Energy Build.* **132**, 4–12 (2016). <https://doi.org/10.1016/j.enbuild.2016.06.094>
21. Völzel, C., Lechner, S.: Grid-coupled geothermal and decentralised heat supply systems in a holistic open-source simulation model for 5GDHC networks. *Sustainability* **16**, 10503 (2024). <https://doi.org/10.3390/su162310503>



Parametric Study Model for Observing Exergy Balance Through the Series of Subsystems in District Heating

Genku Kayo^(✉) 

Hosei University, Tokyo, Japan
genku@hosei.ac.jp

Abstract. District heating (DH) systems have evolved from steam-based heat transport using boilers to combined heat and power (CHP) generation, and more recently to low-temperature heat supply using heat pumps (HP). As DH systems represent a critical energy infrastructure in Nordic regions, it is essential to examine their performance in the context of urban energy transitions. This study investigates exergy consumption in representative archetypes of DH system by developing a parametric model that quantifies the exergy balance across three subsystems: heat generation, heat transfer, and heat exchange within buildings. The results of the parametric analysis indicate that exergy consumption associated with pumping is a more significant contributor to overall exergy balance than the exergy loss through heat dissipation. Specifically, the analysis reveals that the exergy input required for pumps consistently exceeds that required for heat generation. The key finding is that the heat transfer process, particularly the pumping operation, constitutes the primary source of exergy loss in DH systems, underscoring the importance of optimizing this component to enhance overall system efficiency.

Keywords: district heating · exergy consumption · exergy balance equation

1 Introduction

1.1 Energy Theory

The degradation of heat, energy conversion processes, and the qualitative status of energy can be analyzed using exergy theory [1]. A wide range of studies have applied this theoretical framework to evaluate the performance of energy systems, including heat pump systems [2], HVAC systems in buildings [3], and district heating systems [4–6]. Likewise, the energy and exergy analysis on geothermal energy delivery process was conducted to identify the relationship between exergy gain and consumption by applying analytical models of geothermal energy system [7]. In this study, exergy theory is employed to analyze the flow and transformation of exergy across a series of interconnected subsystems using balance equations. Emphasis is placed on evaluating the district heating process by examining the balance between exergy input and exergy consumption.

Exergy consumption in the system arises from the following processes: (1) energy conversion from electromagnetic to mechanical energy, (2) friction losses at the shaft (e.g., in ball bearings) and at the impeller wheel (due to interaction between water and blades), and (3) thermal energy transfer through the piping network. These exergy losses are inherent and unavoidable in the operation of the system. Therefore, evaluating the system from an exergy perspective is essential for assessing its performance and identifying opportunities for improvement. The system's feasibility is assessed based on an exergy balance, which accounts for the exergy input at the pump the exergy consumed throughout the process. If the pump efficiency is denoted by η , this efficiency represents the portion of input energy converted into kinetic power, which is used to generate a pressure differential across the impeller. This kinetic energy is ultimately dissipated as thermal energy through friction. The remaining energy is released as waste heat around the pump housing and into the machine room. These phenomena must be considered in evaluating system performance.

1.2 Objectives

The objective of the research is to analyze the exergy consumption characteristics and input/output balance of different district heating (DH) system archetypes (Boiler based, CHP based, and HP based) by developing a parametric exergy balance model. The model is designed to evaluate exergy flows across subsystems (heat generation, heat transfer, and heat exchange in buildings) under varying operational and environmental conditions. The parametric study systematically varies critical factors such as the distance between the DH plant and buildings, the environmental reference temperature used in exergy calculations, and the volume of water or steam used for heat delivery. Each DH archetype is examined through corresponding case studies to assess and compare performance across system configurations.

2 Methodology

2.1 System Configuration

Figure 1 presents the configuration of the district heating (DH) system, which is composed of four subsystems. Subsystem 1 corresponds to heat generation at the DH plant; Subsystem 2 covers heat distribution from the plant to the buildings; Subsystem 3 involves heat exchange within the buildings; and Subsystem 4 represents the return flow of the heat delivery system. The white arrows in Fig. 1 indicate the directions of exergy input and output.

The primary input factors include: (1) electricity supplied to the pump, which generates a pressure differential across the system, and (2) thermal energy supplied by the heat generator at the DH plant. The output factors include: (3) thermal energy delivered to the buildings and (4) heat losses that occur during heat transfer, which are treated as exergy outputs.

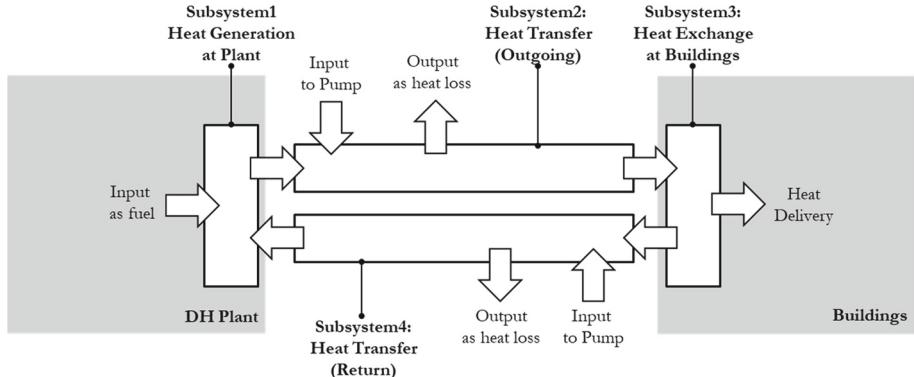


Fig. 1. System configuration

2.2 Exergy Balance Equation at Subsystems

The exergy balance equations for each subsystem are presented in Eqs. (1)–(4). The Eq. (1) describes the exergy balance in subsystem 1 shown in Fig. 1. As well as Eq. (1), the Eqs. (2)–(4) do the exergy balances in subsystems 2, 3, and 4 in Fig. 1. The left-hand side of each equation represents the exergy input to the system, which includes the thermal exergy carried by the water (X_{in}), the kinetic exergy associated with the water flow ($\frac{1}{2}\rho v^2 V$) and the pressure exergy introduced by the pump (X_w). Additionally, the exergy consumed within the system is denoted as X_c . In Eq. (1), represents the exergy gained from the heat generator as fuel (X_{fuel}). The right-hand side of each equation accounts for the exergy output from the system, such as the thermal exergy carried by the water (X_{out}) and the kinetic exergy. Because the system under study is a closed loop with constant velocity and flow rate, the kinetic exergy of the incoming and outgoing streams remains equal. The terms X_{Qloss} in Eq. (2) and Eq. (4) represent exergy outputs as heat losses through heat transferring process, while $X_{delivery}$ in Eq. (3) denotes the exergy delivered to the building side.

$$X_{in1} + \frac{1}{2}\rho v^2 V + X_{fuel} - X_{c1} = \frac{1}{2}\rho v^2 V + X_{out1} \quad (1)$$

$$X_{in2} + \frac{1}{2}\rho v^2 V + X_{w2} - X_{c2} = X_{Qloss2} + \frac{1}{2}\rho v^2 V + X_{out2} \quad (2)$$

$$X_{in3} + \frac{1}{2}\rho v^2 V - X_{c3} = X_{delivery} + \frac{1}{2}\rho v^2 V + X_{out3} \quad (3)$$

$$X_{in4} + \frac{1}{2}\rho v^2 V + X_{w4} - X_{c4} = X_{Qloss4} + \frac{1}{2}\rho v^2 V + X_{out4} \quad (4)$$

The pump locates in subsystem 2 but the power input to the pump is distributed across all subsystems. It is considered as thermal energy dissipated due to friction within each subsystem. The exergy associated with the pump is therefore equivalent to the total mechanical work performed by all subsystems (W_i).

$$X_w = W = Q_f / \eta \quad (5)$$

The friction resulting from the pressure difference between the inlet and outlet of subsystem i is expressed by Eq. (6). The capacity of the circulation pump is determined by the overall pressure level within the piping network, which corresponds to the total pressure drop across all subsystems. The pressure drop in each subsystem is allocated proportionally based on the pipe length associated with that subsystem. The thermal energy generated by friction in subsystem i is calculated using Eq. (7).

$$Q_{fi} = (P_{ini} - P_{outi})V \quad (6)$$

$$Q_{fi} = \frac{\pi}{8} \cdot f_d \cdot D \cdot l_i \cdot \rho \cdot v^3 \quad (7)$$

Here, f_d denotes the dimensionless friction coefficient, D is the pipe diameter [m], l_i is the length of the pipe in subsystem i [m], ρ is the water density [kg/m^3], and v is the flow velocity [m/s]. In the case of cooling (cold water delivery) during summer, the heat generated by friction acts as an undesired thermal gain. Conversely, during the heating in winter, this frictional heat contributes positively to the thermal delivery.

2.3 Model Development

The study model was developed within a visual programming environment, which is widely used in both academic and industrial contexts, particularly in urban environmental and building performance studies. The model can be readily integrated with city or building geometry modeling. This approach offers the potential to transcend disciplinary boundaries and support simulations of energy transitions. Also, the parametric interactions between design variables and objective functions are tangibly understandable.

2.4 Study Case

To represent the main archetypes of district heating networks, three representative cases are defined, as shown in Table 1.

Table 1. Study case

	Case1	Case2	Case3
Heat Generator	Boiler	CHP	HP
Efficiency η	0.9	0.45 (th), 0.35 (el)	COP = 3.0
Heat Carrier	Steam	Water	Water
Main Pipe Diameter [m]	0.250	0.200	0.150
Supply Temp [°C]	180	100	60
Return Temp [°C]	60	60	40

2.5 Parameter Setting

In this model, the tangible input parameters are the distance between the plant and buildings as L [m], environmental temperature as T_o [K], and reference temperature T_{ref} [K] to calculate exergy. T_o simplify the model, the number of branches and layout of pipe are not considered. The volume of steam or water as V [m^3/s] which circulates through the subsystems is determined to keep the velocity in the range between 2.0 and 3.0 m/s. The range of each parameter is listed in Table 2. The material of the pipe is SUS (thickness = 0.014 m) with glass wool insulation (thickness = 0.06 m). The efficiency of the pump (η) is 0.80 constant.

Table 2. The range of parameters

			Min Value	Max Value
Environmental Temperature	T_o	[°C]	-10	20
Reference Temperature	T_{ref}	[°C]	-10	20
Distance	L	[m]	10^3	10^4
Velocity	v	[m/s]	2.0	3.0

3 Results and Discussion

3.1 Exergy Balance of Each Archetype

Figures 2, 3, and 4 illustrate the exergy input–output balance for each subsystem across different district heating (DH) archetypes. In each graph, the upper section represents the exergy input to each subsystem, while the lower section shows the exergy output and the exergy consumed during the processes within the subsystem. Light green bars represent exergy inputs to the system, including fuel supplied to the heat generator and electricity for circulation pumps. Light red bars indicate the exergy carried by the steam or water circulating within the system. Yellow bars represent exergy consumption (X_c) due to irreversibility in the subsystem. Purple bars represent thermal exergy caused by heat losses (XQ_{loss}) to surroundings. And Red bar represents the thermal exergy (XQ_{deli}) delivered to buildings. In the exergy flow, the input exergy (X_{in}) to a given subsystem corresponds to the output exergy (X_{out}) from the preceding subsystem.

Figure 2 is the result in case of boiler setting the values as the volume flow $190 \text{ m}^3/\text{s}$ (in which situation, the velocity is 3.0 m/s). The results indicate that the exergy input required for the pumps is relatively greater than that for the heat generator. A similar pattern is observed in exergy consumption, suggesting that the transfer process, especially pumping, constitutes the most significant source of exergy loss within the system. Figure 3 is the result in case of CHP. The setting values are the volume flow $340 \text{ m}^3/\text{s}$ (in which situation, the velocity is 3.0 m/s), Fig. 4 is the result in case of CHP. The setting values are the volume flow $190 \text{ m}^3/\text{s}$ (in which situation, the velocity is 3.0 m/s), Distance L as 1000 m, environmental temperature T_o as -5°C , and reference temperature T_{ref} as 10°C are set as same in all cases.

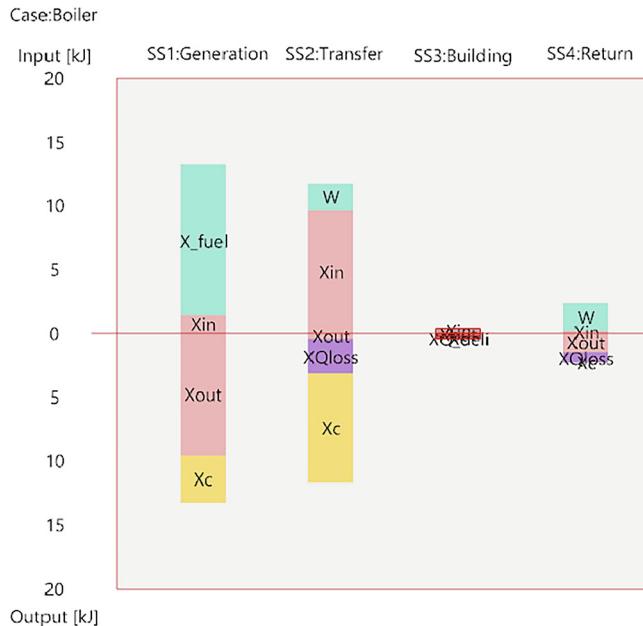


Fig. 2. Exergy balance of each subsystem (Case1: Boiler, $L = 1000$ m, $V = 500$ m³/s)



Fig. 3. Exergy balance of each subsystem (Case2: CHP, $L = 1000$ m, $V = 340$ m³/s)



Fig. 4. Exergy balance of each subsystem (Case3: HP, $L = 1000$ m, $V = 190$ m³/s)

3.2 Parametric Study

Distance Between DH Plant and Buildings

Figure 5 shows the result of CHP case distance L is 3000 m. Compared with the result in Fig. 3, the required exergy input to the pumps also increases as the distance between the district heating (DH) plant and the buildings increases. This is due to the enlargement of the inner pipe surface area, which necessitates higher-capacity pumps. Consequently, exergy consumption caused by friction along the inner pipe surface becomes more significant. The heat generated by this friction surpasses the heat losses through the pipe walls, resulting in a higher exergy output at the end of Subsystem 3.

Volume of Heat Carrier

Figure 6 shows the result of CHP case V is 220 m³/s. Compared with the result in Fig. 3, Similarly, when the volume of steam or water (V) increases, the trend are opposite to that observed with increased distance. A larger volume of transfer medium leads to greater friction losses within the pipe and requires pumps with greater capacity. This parametric study indicates that exergy consumption at the pump is a major contributing factor to the overall exergy balance of the system.

Environmental Temperature

The amount of exergy changes depending on the environmental temperature (T_o) as shown in Eqs. (1)–(4). The parametric study model can demonstrate the behavior of exergy balance by changing the environmental temperature. The model illustrates that

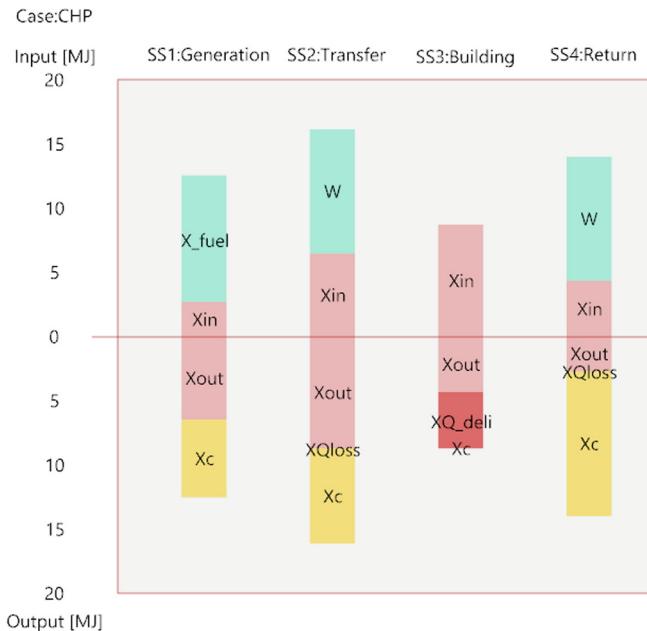


Fig. 5. Exergy balance of each subsystem (Case2: CHP, L = 3,000 m V = 340 m³/s)

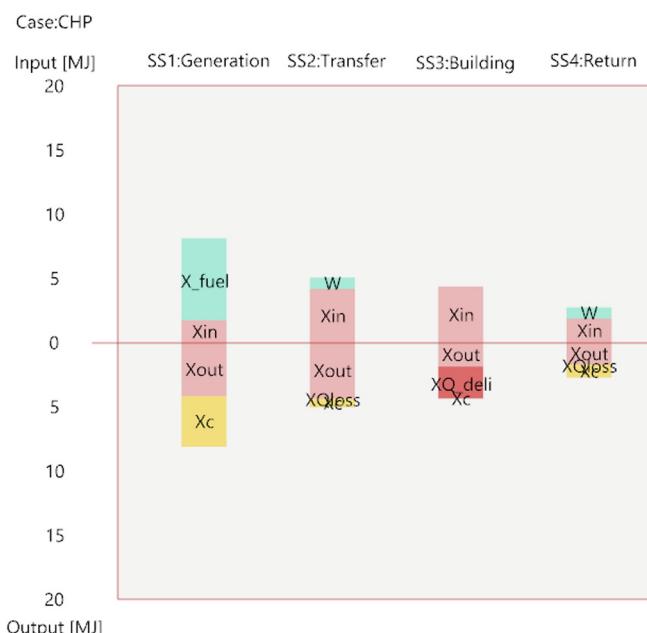


Fig. 6. Exergy balance of each subsystem (Case2: CHP, L = 1,000 m, V = 220 m³/s)

the total amount of exergy input and output are changed but the proportion of exergy balance of each subsystem is not changed.

4 Conclusions

The results of the parametric analysis indicate that exergy consumption associated with pumping is a more significant contributor to overall exergy balance than the exergy loss through heat dissipation. Specifically, the analysis reveals that the exergy input required for pumps consistently exceeds that required for heat generation. The key finding is that the heat transfer process, particularly the pumping operation, constitutes the primary source of exergy loss in DH systems, underscoring the importance of optimizing this component to enhance overall system efficiency.

The aim of the model is to demonstrate the interaction between study parameters and objective variables (exergy). The parameters which affect the quantitative amount of exergy are more than two, such as environmental temperature, volume of water, profile and insulation level of pipes, or the distance between the DH plant and buildings, etc. Therefore, it is hard to illustrate the exergy balance with 2D or 3D plot. The article describes the developed exergy balance model and proposes the way to study the behaviors of exergy tangibly. The key findings are that the exergy input and consumption at the pumps are one of the most influential factors for exergy flow in the whole system. The case study also indicates that the parametric study model makes it easy to capture the exergy behavior even if the system configuration is complicated.

References

1. Shukuya, M.: Bio-Climatology for Built Environment, 1st edn. CRC Press, Boca Raton (2021)
2. Li, R., Ooka, R., Shukuya, M.: Theoretical analysis on ground source heat pump and air source heat pump systems by the concepts of cool and warm exergy. *Energy Build.* **75**, 447–455 (2014)
3. Ferrara, M., Coleman, J., Meggers, F.: Exploring potentialities of energy-connected buildings: performance assessment of an innovative low-exergy design concept for a building heating supply system. *Energy Procedia* **122**, 1075–1080 (2017)
4. Gong, M., Werner, S.: Exergy analysis of network temperature levels in Swedish and Danish district heating systems. *Renewable Energy* **84**, 106–113 (2015)
5. Gong, M., Werner, S.: Mapping energy and exergy flows of district heating in Sweden. *Energy Procedia* **116**, 119–127 (2017)
6. Fitó, J., et al.: Energy- and exergy-based optimal designs of a low-temperature industrial waste heat recovery system in district heating. *Energy Convers. Manage.* **211**, 112753 (2020)
7. Kayo, G., Choi, W., Shukuya, M., Ooka, R.: Exergy analysis on district-scale geothermal energy delivery system. In: IAQVEC2023 Conference, Tokyo, Proceedings, pp. 1–8 (2023)



Demand-Side Frequency Response Based on District Heating System Integrated with Heat Pump

Hui Yan^(✉) and Sara Walker

School of Chemical Engineering, University of Birmingham, Birmingham, UK
h.yan@bham.ac.uk

Abstract. The increasing penetration level of renewable generation in power systems leads to challenges in frequency stability due to system inertia reduction. This study investigates the potential of the district heating system (DHS) integrated with the variable-speed heat pump to provide demand-side frequency response services while meeting end-user heating demands. Dynamic models are developed considering the thermal inertia of the DHS with the heat pump, and the dynamic performance is analyzed under rotational speed disturbances. The demand-side frequency control is proposed which regulates the supply water temperature setpoint to adjust heat pump power consumption during primary and secondary frequency response processes. The effectiveness of the proposed control is validated through a town-level case study. Results show that the proposed control effectively reduces the frequency nadir and accelerates frequency recovery, particularly under large load disturbances. For a step load increase of 0.25 p.u., the frequency nadir improves from -0.0265 p.u. to -0.0241 p.u. and the frequency settling time decreases from 973.2 s to 913.7 s. Additionally, the room temperature variation is less than 0.8 °C taking 2000 s, satisfying the acceptable thermal comfort criteria.

Keywords: District heating system · Heat pump · Frequency response

1 Introduction

The renewable power penetration will continuously increase in power systems to meet various countries' decarbonization targets. The converter-based renewable generation accounted for approximately 35% of electricity in the UK in 2023 and is expected to exceed 60% by 2035[1]. However, the increasing penetration of renewable power poses challenges in ensuring frequency stability, partly due to the reduction of system inertia. System operators are seeking additional sources for frequency regulation beyond the traditional synchronous generators. One promising alternative is the demand-side response, which is widely considered a means of providing ancillary service [2].

Heating systems constitute a growing portion of the total load in power systems. Their inner energy storage capacity and thermal inertia can offer the potential for demand-side flexibility, for individual and district heating systems. Many researchers have conducted

related research on the district heating system (DHS) with the combined heat and power unit (CHP). Xu et al. [3] developed a methodology to quantify the maximum flexibility of DHS with CHP by decomposing them into subsystems and considering thermal inertia and transport delays. Li et al. [4] studied the CHP-coupled DHS and proposed a model to use thermal inertia for cost reduction and wind power integration. Dai et al. [5] modeled the heat network with the consideration of heat transfer limitation.

With the growing trend of heating system electrification, the heat pump becomes a popular and efficient technology to utilize low-temperature heat while reducing the water circulation temperature. Many studies investigated the capability of heat pumps to support the frequency response at different levels. Kim et al. [6] developed a dynamic model of the variable-speed heat pump in a commercial building for secondary frequency regulation by responding to direct load control. Rasmussen et al. [7] indicated that large-scale heat pumps can provide primary frequency support via local droop and inertia control. They also proposed model predictive control coordinated primary frequency control of large-scale and small-scale heat pumps [8]. However, most of the previous studies simplified the heat pump dynamic behaviors into constant transfer functions or empirical fitting curves, which did not fully reflect the nonlinear dynamic behaviors of heat pumps. Song et al. [9] modeled the non-linear power response of the air-source heat pumps and proposed a fast frequency response scheme using them, but its impact on the heating system was not considered.

Heat pump-driven DHS as a demand-side resource has great potential to participate in the frequency response while respecting the end-user heating demand. However, to the best of the authors' knowledge, the aforementioned previous studies mainly focused on heat pump dynamics and DHS operation with CHP, while integrated modelling and analysis of heat pump-driven DHS for frequency response have not been sufficiently addressed. The thermal inertia of the heat-pump-driven DHS requires further assessment, and the control of such systems for primary and secondary frequency response should be further explored. Therefore, this paper develops dynamic models of the DHS integrated with the variable-speed heat pump, describing the interaction between thermodynamic and electrical processes. A demand-side frequency control strategy is proposed which adjusts heat pump power consumption in response to grid frequency deviations by regulating the supply water temperature, thereby providing both primary and secondary frequency responses. The effectiveness of the control is validated through a case study of a town-level integrated energy system.

2 Model Development

2.1 System Description

Figure 1 describes the integration of a DHS and the distribution grid. The DHS contains the centralized heat pump system, the secondary-network (SN) water supply and return pipes and the heat users. In this study, the heat pump system replaces the secondary heat exchange station of the conventional heating system, which directly provides the hot water as the heat source for the heat users.

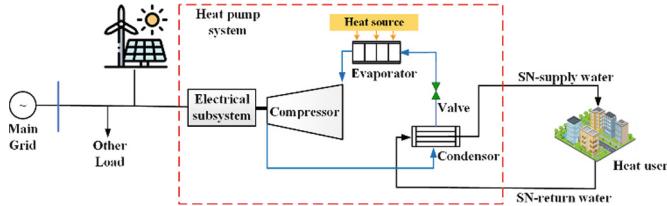


Fig. 1. Diagram of the integration of the electrical and district heating system

2.2 Dynamic Model of the Heat Pump

The variable-speed heat pump supplies heated water to the DHS by extracting heat from low-temperature heat sources. As shown in Fig. 2, the heat pump consists of an evaporator, condenser, compressor and throttle valve. In the evaporator, the refrigerant absorbs heat from the air or water and transforms into superheated vapor. The enthalpy and pressure of the superheated refrigerant continue to increase in the compressor. Then, the heat is transferred to the cold water in the condenser. Finally, the cooled refrigerant flows back to the evaporator through the throttle valve.

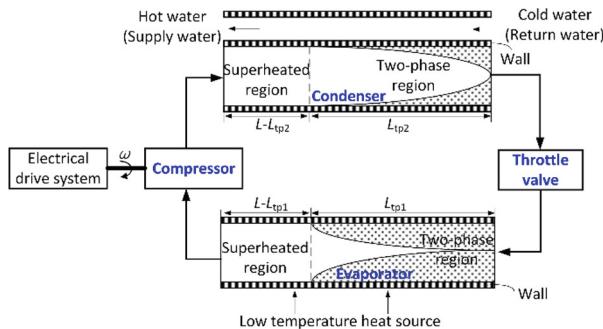


Fig. 2. General construction of the variable-speed heat pump

Model of Evaporator/Condenser

To capture the nonlinear speed-power response of the heat pump, the moving boundary method is used to model the evaporator and the condenser. The refrigerant flows through different regions (eg. Two-phase and superheated zones) and the boundaries between these regions shift according to the variations in heat transfer, mass flowrate and pressure. One-dimensional flow is assumed, with negligible pressure drop. R134a is applied as the refrigerant in this study.

The model approach for the evaporator and condenser is similar, differing in the order of the two-phase and superheated regions. Due to the space limitation, the evaporator model is explained here as an example. The inlet refrigerant of the evaporator is assumed to be the saturated mixture of liquid and vapor. The evaporator model is developed based on the mass conservation of the refrigerant, and energy conservations of the refrigerant

and pipe metal wall [9], expressed as:

$$\begin{cases} A_{\text{in-c}} \frac{\partial \rho_{\text{re}}}{\partial t} + \frac{\partial m_{\text{re}}}{\partial x} = 0 \\ A_{\text{in-c}} \frac{\partial(\rho_{\text{re}} h_{\text{re}} - p)}{\partial t} + \frac{\partial(m_{\text{re}} h_{\text{re}})}{\partial x} + k_{\text{in}} A_{\text{in-s}} (T_{\text{re}} - T_{\text{we}}) = 0 \\ c_{\text{we}} \rho_{\text{we}} (A_{\text{out-c}} - A_{\text{in-c}}) \frac{\partial T_{\text{we}}}{\partial t} = k_{\text{in}} A_{\text{in-s}} (T_{\text{we}} - T_{\text{re}}) + A_{\text{out-s}} k_{\text{out}} (T_0 - T_{\text{we}}) \end{cases} \quad (1)$$

where ρ_{re} is the refrigerant density, $\text{kg}\cdot\text{m}^{-3}$; m_{re} is the mass flowrate of the refrigerant, $\text{kg}\cdot\text{s}^{-1}$; p is the pressure, Pa ; h_{re} is the enthalpy of the refrigerant, $\text{J}\cdot\text{kg}^{-1}$; T_{re} , T_{we} and T_0 are the temperatures of the refrigerant, the metal wall and the low-temperature heat source, K ; $A_{\text{in-c}}$ and $A_{\text{in-s}}$ are the cross-section and surface areas of the inner pipe, m^2 ; $A_{\text{out-s}}$ is the pipe outer surface area, m^2 ; k_{in} and k_{out} are heat transfer coefficients of the metal wall to the refrigerant and to the heat source, $\text{W}\cdot\text{K}^{-1}$; c_{we} is the specific heat capacity of the metal pipe wall, $\text{J}\cdot(\text{kg}\cdot\text{K})^{-1}$; ρ_{we} is the metal pipe wall density, $\text{kg}\cdot\text{m}^{-3}$.

The partial differential equations are integrated over the two-phase and superheated regions, respectively, and subsequently converted to the differential algebraic equations that can be solved directly. The energy and mass conservation of water are considered in the condenser as well.

Model of Compressor

The compressor stage is modeled as a standard isentropic compression process. The mathematical model of the compressor [9] is as:

$$m_{\text{cp}} = \omega_{\text{cp}} V_{\text{cp}} \rho_{\text{cp}} \eta_{\text{cp}} \quad (2)$$

$$h_{\text{cp-out}} = \frac{K_1 R T_{\text{cp-in}}}{K_1 - 1} \left[\left(\frac{p_{\text{cp-out}}}{p_{\text{cp-in}}} \right)^{\frac{K_1 - 1}{K_1}} - 1 \right] + h_{\text{cp-in}} \quad (3)$$

$$P_{\text{HP}} = m_{\text{cp}} (h_{\text{cp-out}} - h_{\text{cp-in}}) \quad (4)$$

where m_{cp} is the refrigerant flowrate through the compressor, $\text{kg}\cdot\text{s}^{-1}$; V_{cp} is theoretical volume displacement per revolution, $\text{m}^3\cdot\text{s}^{-1}$; ρ_{cp} is the refrigerant density at the compressor inlet, $\text{kg}\cdot\text{m}^{-3}$; η_{cp} is the compressor gas transmission coefficient; $h_{\text{cp-in}}$ and $h_{\text{cp-out}}$ are the refrigerant enthalpy at the compressor inlet and outlet, $\text{J}\cdot\text{kg}^{-1}$; K_1 is an empirical constant; R is the gas constant, $8.314 \text{ J}\cdot(\text{mol}\cdot\text{K})^{-1}$; $p_{\text{cp-in}}$ and $p_{\text{cp-out}}$ are the inlet and outlet pressures, Pa ; $T_{\text{cp-in}}$ is the inlet refrigerant temperature of the compressor, K ; P_{HP} is the power consumption of the compressor, W .

Model of Throttle Valve

The throttle valve follows the Bernoulli's principle. The model [10] is expressed as:

$$m_{\text{vlv}} = C_{\text{vlv}} \sqrt{\rho_{\text{vlv}} (p_{\text{vlv-out}} - p_{\text{vlv-in}})} \quad (5)$$

$$h_{\text{vlv-out}} = h_{\text{vlv-in}} \quad (6)$$

where m_{vlv} is the mass flowrate of the throttle valve, $\text{kg}\cdot\text{s}^{-1}$; C_{vlv} is the flow characteristic parameter; ρ_{vlv} is the refrigerant density at the inlet of the valve, $\text{kg}\cdot\text{m}^{-3}$; $p_{\text{vlv-in}}$ and $p_{\text{vlv-out}}$ are the inlet and outlet pressure at of the throttle valve, Pa ; $h_{\text{vlv-in}}$ and $h_{\text{vlv-out}}$ are the refrigerant enthalpy at inlet and outlet of the throttle valve, $\text{J}\cdot\text{kg}^{-1}$.

Model of Electrical Drive System

The electrical drive system has the permanent magnetic synchronous machine connecting to the compressor shaft and driven by the AC-DC-AC power converters. The rotational speed is controlled. As the electrical drive system responses much quickly than the thermal system, detailed internal states and switching are not considered. The transfer function method is applied to describe the response of the rotational speed to a given setpoint [9].

2.3 Dynamic Model of the District Heating System

In the secondary heat network, the hot water produced by the heat pump is delivered to the heat-users through pipelines. To assess the thermal inertia of the heat-pump driven district heating network, it is necessary to consider not only the heat pump but also the heat demand and storage of the buildings and pipelines.

Model of Building

Considering the heat storage and demand of the buildings, the building model consists of the radiator, rooms and the envelope structure [11]. The model is developed based on the energy balance equations, which are as:

$$\begin{cases} C_d \frac{dT_w}{dt} = c_w m_w (T_{\text{win}} - T_{\text{wout}}) - k_d A_d \left(\frac{T_{\text{win}} + T_{\text{wout}}}{2} - T_{\text{rm}} \right) \\ C_{\text{rm}} \frac{dT_{\text{rm}}}{dt} = k_d F_d \left(\frac{T_{\text{win}} + T_{\text{wout}}}{2} - T_{\text{rm}} \right) - k_{\text{wli}} A_{\text{wli}} (T_{\text{rm}} - T_{\text{wl}}) - k_{\text{en}} A_{\text{en}} (T_{\text{rm}} - T_0) \\ C_{\text{wl}} \frac{dT_{\text{wl}}}{dt} = k_{\text{wli}} A_{\text{wli}} (T_{\text{rm}} - T_{\text{wl}}) - k_{\text{wlo}} A_{\text{wlo}} (T_{\text{rm}} - T_0) \end{cases} \quad (7)$$

where C_d , C_{rm} , C_{wl} are the aggregated heat capacities of the radiator, rooms, and room walls, $\text{J}\cdot\text{K}^{-1}$; c_w is the specific heat capacity of water, $\text{J}\cdot(\text{kg}\cdot\text{K})^{-1}$; m_w is the water mass flowrate through the radiator, $\text{kg}\cdot\text{s}^{-1}$; T_w , T_{win} and T_{wout} are the water average temperature in the radiator, supply water temperature and return water temperature, K ; T_{rm} , T_{wl} and T_0 are the temperatures of the inside room, room wall, and outside environment, K ; k_d , k_{wli} , k_{en} and k_{wlo} are the heat transfer coefficients of the radiator, inner wall, window, and outer wall of the room, $\text{W}\cdot\text{K}^{-1}$; A_d , A_{wli} , A_{en} and A_{wlo} are the heat transfer areas of the radiator, inner wall, window, and outer wall of the room, m^2 .

Model of the Pipe Network

The pipe network is modeled as the combination of pipe elements and water mixture nodes, which can represent the network constructed by the direct buried laying method. Each pipe element consists of the metal wall and thermal insulation layer. As water flows through the pipes, heat is transferred to the wall and subsequently dissipated to

the surrounding soil. The energy balance equations of the pipelines are as:

$$\begin{cases} C_{\text{pip}} \frac{dT_{\text{pout}}}{dt} = c_w m_p (T_{\text{pin}} - T_{\text{pout}}) - k_{\text{pip}} A_{\text{pip}} \left(\frac{T_{\text{pin}} + T_{\text{pout}}}{2} - T_m \right) \\ C_m \frac{dT_m}{dt} = k_{\text{pip}} A_{\text{pip}} \left(\frac{T_{\text{pin}} + T_{\text{pout}}}{2} - T_m \right) - \frac{T_m - T_{\text{so}}}{R_m + R_{\text{so}}} (1 + \alpha) \cdot L_p \end{cases} \quad (8)$$

where C_{pip} and C_m are the heat capacities of the pipe and the metal wall of the pipe, $\text{J}\cdot\text{K}^{-1}$; T_{pin} and T_{pout} are the temperatures of the inlet and outlet water of the pipe, K ; T_m is the metal wall temperature of the pipe, K ; k_{pip} is the pipe heat transfer coefficient; A_{pip} is the heat transfer area of the pipe, m^2 ; T_{so} is the soil temperature around the pipe, K ; R_m and R_{so} are the thermal resistances of the thermal insulation layer and the soil, $(\text{m}\cdot\text{K})\cdot\text{W}^{-1}$; α is the conversion coefficient; L_p is the pipe length, m .

The water mixture node is modeled as:

$$T_{\text{wmix}} = \frac{1}{m_w} \sum_i T_{\text{wi}} m_{\text{wi}} \quad (9)$$

where T_{wmix} is the water temperature of the mix backwater pipe, K ; T_{wi} is the temperature of the outlet water of the i th user radiator, K ; m_{wi} is the mass flowrate of the water through the i th user radiator, $\text{kg}\cdot\text{s}^{-1}$.

2.4 Demand-Side Frequency Response

As shown in Fig. 3, the conventional thermal power generators are initially considered as the main assets to regulate the grid frequency in this study. Primary frequency control is implemented through droop control, while secondary frequency control is achieved via automatic generation control (AGC). α_1 and $(1 - \alpha_1)$ are the proportion factors of heat pump and generator in secondary frequency control, which assume 0.5 in this study. The equivalent grid model [12] is simplified as the transfer function G_{gd} :

$$G_{\text{gd}} = \frac{1 + F_{\text{HP}} T_{\text{RH}} s}{(1 + T_{\text{ST}} s)(1 + T_{\text{RH}} s)(1 + T_g s)} \quad (10)$$

where T_{ST} and T_{RH} are the steam turbine and reheater time constants; T_g is the reheater gain; F_{HP} is the generator governor constant. D and M represent the load damping coefficient and the rotational inertia of thermal power units, respectively. $\Delta P_L(s)$ and $\Delta f(s)$ are the load disturbance and frequency deviation, respectively.

The heat pump is proposed as a demand-side frequency response unit participating in both primary and secondary frequency responses by regulating its power consumption. The heat pump power is controlled by adjusting the rotational speed. The set point of the rotational speed (ω^*) is determined by a PID controller based on the supply water temperature deviation between the real value and setpoint (T_{sup}^*). T_{sup}^* responds to the aggregation signal of the droop and AGC control. The PID controller gains can initially be obtained by Ziegler-Nichols tuning method under step disturbance.

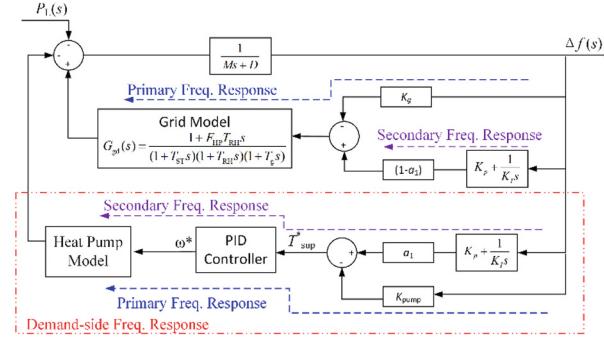


Fig. 3. Demand-side frequency response model based on heat-pump driven DHS

3 Results and Discussion

In this section, the dynamic performance of a town-level DHS with heat pump is studied and the effectiveness of the proposed demand-side frequency control is analyzed. The nadir and frequency settling time are used to assess the frequency stability.

3.1 Dynamic Performance of Heat Pump and Heat Network

Heat Pump

Due to the thermal inertia of the heat pump, the temperature of the generated hot water cannot change transiently. The rotational speed (ω), as the regulating measure to adjust heat pump power consumption, influences the dynamic performance of the heat pump. The main parameters of the heat pump are shown in Table 1.

Table 1. Main parameters of the heat pump.

Parameter	Value	Unit	Parameter	Value	Unit
Rated Power	760	kW	COP	3.49	-
Evaporator Pressure	225.5	kPa	Condenser Pressure	1491.5	kPa
Heat source temperature	10	°C	Cold water temperature	25.0	°C
Evaporator/Condenser length	15	m	Hot water temperature	50.2	°C
Evaporator refrigerant flowrate	15.53	kg·s ⁻¹	Condenser refrigerant flowrate	15.53	kg·s ⁻¹
Compressor rotational speed	3000	r·min ⁻¹	Water flowrate	25.1	kg·s ⁻¹

Figure 4 shows the dynamic performance of the heat pump under rotational speed disturbances. When ω step increases or decreases at 500 s, the outlet hot water temperature increases or decreases accordingly with a gradually decreasing rate. The outlet hot water temperature changes from 50.2 °C to between 41.4 °C and 52.0 °C as ω changes within the range of -50% to $+15\%$. As shown in Fig. 4 (b), the settling time increases with the rotational speed change ratio, ranging from 80.9 s to 129.0 s.

As shown in Fig. 4 (c), when ω step increases, the heat pump power consumption increases rapidly at first and then gradually to a new steady-state value higher than the initial level. The power consumption changes faster than the hot water temperature, which provides the potential for heat pumps to participate in the frequency response service with minimal impact on the heat network water supply. When ω varies $-50\% \sim +15\%$, the power consumption variation ranges between 368.0 kW and 94.3 kW.

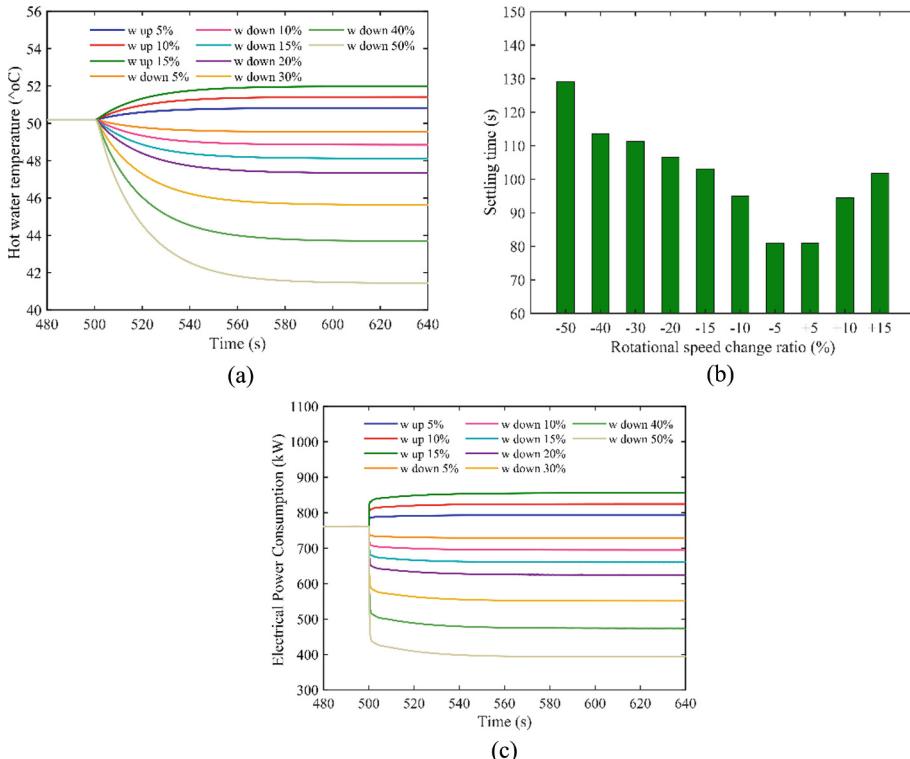


Fig. 4. Dynamic performance of the heat pump under step rotational speed disturbances (a) hot water temperature (b) settling time (c) power consumption

District Heating System with Heat Pump

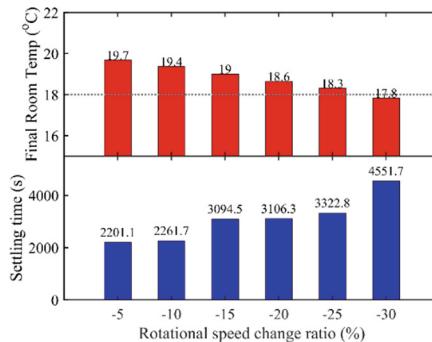
The heat pump provides the hot supply water for a town-level DHS, whose rated heat load is 2.66 MW. The main DHS parameters are shown in Table 2.

When the rotational speed step changes, due to the working medium and heat storage of the heat pump, building and pipelines, the room temperature changes gradually with a transient period. Settling time is the indicator to assess the time taking for room temperature changing from the initial state and the final state. Figure 5 compares the final room temperature and its settling time when rotation speed step decreases by $5\% \sim 30\%$. It can be shown that the final room temperature can be kept above 18 °C if the

Table 2. Main parameters of the DHS with heat pump.

Parameter	Value	Unit	Parameter	Value	Unit
Supply water temperature	50.2	°C	Diameter of the network pipe	0.126	m
Return water temperature	25.0	°C	Total length of the network pipe	1600	m
Flowrate of supply water	25.1	kg·s ⁻¹	Rated room temperature	20	°C

rotational speed changes by less than 25%, while the settling times in such cases range between 2201.1 s and 3322.8 s.

**Fig. 5.** Final room temperature and settling time of the DHS when rotational speed step changes

3.2 Frequency Stability Improvement Using Heating System with Heat Pump

This section tests the ability of the DHS with heat pump in demand-side frequency response process and its impact on the heat users. Assuming the heat pump connects to a distribution electrical system, the maximum total load consumption is 4720 kW. At the benchmark condition, the heat pump and other load power consumptions are 760 kW and 2960 kW, respectively. The generation capacity is 5700 kW (generation margin: 20% of total demand). The base value of the frequency is 50 Hz.

Figure 6 compares the frequency responses under load step disturbances of 0.05 p.u. – 0.25 p.u. with and without the proposed frequency control of heat-pump driven DHS. In reference cases, only synchronous generators (SG) regulate the frequency.

As shown in Fig. 6 (a), when the load step increases at 3000 s, the frequency deviation drops first and then recovers to zero. The maximum frequency deviation (called nadir) is effectively reduced by the heat-pump driven DHS, as the proposed control decreases the heat pump power consumption by reducing the water supply temperature responding to negative frequency deviations. The absolute value of the nadir increases with the load disturbance magnitude. Figure 6 (b) summarizes the frequency nadir and settling time. The heat pump better reduces the frequency deviation when larger disturbances happen. The nadir improves from -0.53×10^{-2} p.u. and -2.65×10^{-2}

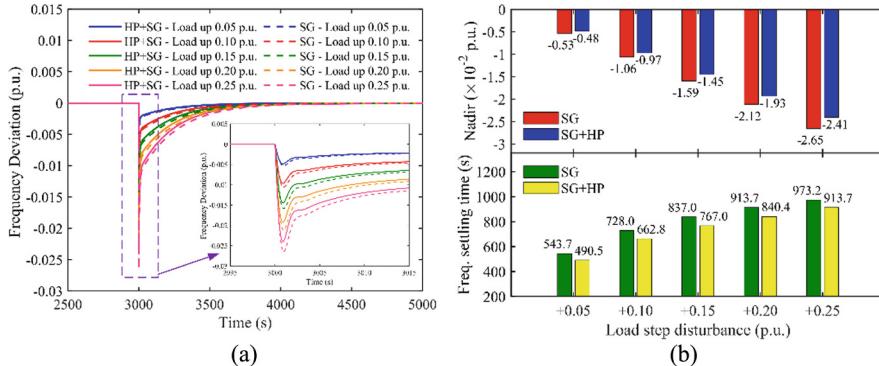


Fig. 6. Frequency response comparisons with/without demand-side response under load step disturbances (a) frequency response curves (b) Frequency nadir and settling time comparisons

p.u. to -0.48×10^{-2} p.u. and -2.41×10^{-2} p.u., when the load step increases by 0.05 p.u. and 0.25 p.u., respectively. Further, the proposed demand-side frequency control reduces the settling time for frequency deviation recovering to zero. The settling time decreases from 543.7 s – 973.2 s to 490.5 s – 913.7 s.

Figure 7 shows the final value and settling time of the room temperature under the above load step increase disturbances. When DHS with heat pump participates in frequency regulation, the final room temperature decreases and the settling time increases with the disturbance magnitude, as the water supply temperature is controlled to decrease when the frequency deviation is negative. The final room temperature decreases from 20 °C to 19.8 °C and 19.2 °C in 797.5 s and 1972.7 s. All the simulated cases satisfy the requirement of human thermal comfort defined in ASHRAE 55 that the acceptable temperature ramp should limit to 1.1 °C within 1 h.

Therefore, the demand-side frequency control by heat-pump driven DHS can effectively enhance the overall frequency stability and guarantee the user thermal comfort.

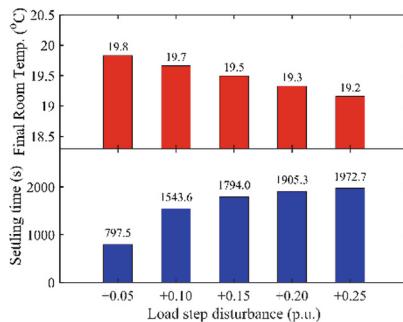


Fig. 7. Impact of the demand-side frequency response on the room temperature

4 Conclusion

In this study, a heat-pump driven DHS as a demand-side resource is used to participate in frequency response while respecting the end-user heating demand. The dynamic model of the DHS with variable-speed heat pump is developed considering its thermal inertia. Dynamic performance of the heat pump and DHS are analyzed. The demand-side frequency control is proposed to enhance the frequency stability and the effectiveness is verified by a town-level case study. Main conclusions are as follows.

- (1) When the rotational speed step decreases, the outlet hot water temperature of the heat pump drops at a gradually decreasing rate. The room temperature changes more slowly due to the thermal inertia of the heat pump, building and pipelines.
- (2) Regulating the rotational speed, depending on the control of the supply water temperature, enables the heat pump to participate in primary and secondary frequency response. The proposed demand-side frequency control effectively reduces the nadir and accelerates frequency recovery, especially for large disturbances. The nadir is improved from -2.65×10^{-2} p.u. to -2.41×10^{-2} p.u and the frequency settling time is shortened from 973.2 s to 913.7 s under step increase load disturbance of 0.25 p.u.
- (3) The heat-pump driven DHS can effectively enhance the frequency stability while guaranteeing the user thermal comfort. Under load disturbance ranging from 0.05 p.u. and 0.25 p.u., the room temperature variation is less than 0.8 °C taking 2000 s, which is within the acceptable temperature ramp (1.1 °C within 1 h).

Acknowledgement. This work is part of Energy Demand Research Centre programme, funded by UKRI (grant EP/Y010078/1).

References

1. Gov.Uk: Electricity statistics. <https://www.gov.uk/government/collections/electricity-statistics>
2. Alotaibi, I.M., Abido, M.A., Khalid, M.: Primary Frequency Regulation by Demand Side Response. *Arab. J. Sci. Eng.* **46**, 9627–9637 (2021)
3. Xu, X., Lyu, Q., Qadrda, M., et al.: Quantification of flexibility of a district heating system for the power grid. *IEEE Trans. Sustain. Energy* **11**, 2617–2630 (2020)
4. Li, X., Li, W., Zhang, R., et al.: Collaborative scheduling and flexibility assessment of integrated electricity and district heating systems utilizing thermal inertia of district heating network and aggregated buildings. *Appl Energ* **258**, 114021 (2020)
5. Dai, Y., Chen, L., Min, Y., et al.: Dispatch model for chp with pipeline and building thermal energy storage considering heat transfer process. *IEEE Trans. Sustain. Energy* **10**, 192–203 (2019)
6. Kim, Y.-J., Norford, L.K.: Kirtley JL modeling and analysis of a variable speed heat pump for frequency regulation through direct load control. *IEEE Trans. Power Syst.* **30**, 397–408 (2015)
7. Rasmussen, T.B.H., Wu, Q.: Primary frequency support from local control of large-scale heat pumps. *Int. J. Electr. Power Energy Syst.* **133**, 107270 (2021)

8. Rasmussen, T.B.H., Wu, Q., Moller, J.G., et al.: MPC coordinated primary frequency support of small- and large-scale heat pumps. *IEEE Transactions on Smart Grid* **13**, 2000–2010 (2022)
9. Song, R., Terzija, V., Hamacher, T., et al.: Integrating air-source heat pumps into the demand-side fast frequency response service: a study based on thermal dynamic uncertainty. *IEEE Transactions on Sustainable Energy* **16**, 323–335 (2025)
10. Wang, L., Zhang, S., Fu, Y., et al.: Heat–power decoupling for the CHP unit by utilizing heat storage in the district heating system integrated with heat pumps: dynamic modeling and performance analysis. *Energy* **306**, 132485 (2024)
11. Wang, L., Zhang, S., Fu, Y., et al.: Usable thermal energy stored in the district heating system for peak shaving: a dynamic simulation study. *Appl. Therm. Eng.* **239**, 122169 (2024)
12. Li, P., Tan, Z., Zhou, Y., et al.: Secondary frequency regulation strategy with fuzzy logic method and self-adaptive modification of state of charge. *IEEE Access* **6**, 43575–43585 (2018)



A Parametric Optimization Approach for Enviro-Economic Evaluation of Energy Renovation Strategies – A Case Study on a Congregation House in Southern Sweden

Md Parvaz¹ (✉) , Einar Örn Porvaldsson¹ , Jesper Engström¹, Dennis Johansson² , and Henrik Davidsson¹ 

¹ Division of Energy and Building Design, Department of Building and Environmental Technology, Faculty of Engineering, Lund University, 221 00 Lund, Sweden
md6480pa-s@student.lu.se

² Division of Building Services, Department of Building and Environmental Technology, Faculty of Engineering, Lund University, 221 00 Lund, Sweden

Abstract. This study explores sustainable energy renovation strategies for a congregation building in southern Sweden using a parametric optimization approach to balance passive and active measures. Passive measures include envelope and window improvements and leakage control, while active measures involve photovoltaic integration and ventilation enhancements. The renovation aims to optimize energy use, thermal comfort, indoor air quality (IAQ), life cycle cost (LCC), and life cycle assessment (LCA). However, individual measures often conflict with other objectives. By integrating the optimization tool Opossum with ClimateStudio, which simulates energy, thermal, and photovoltaic panels (PV) performance, along with programmed LCC and LCA calculations, the study identifies optimal solutions. Results indicate a 24% energy reduction, 23% lower LCC, and 47% less CO₂ emissions without changing the existing ventilation system, though IAQ and comfort remain inadequate. With new ventilation, energy use increases, LCC rises by 53%, but CO₂ emissions reduce by 11%, resolving IAQ issues and improving comfort by 74%. The study acknowledges limitations due to data assumptions and simulation constraints; future work could include full building LCA and moisture assessments to enhance the findings.

Keywords: Energy renovation · Parametric optimization · Life cycle costing · Photovoltaic systems · Life cycle assessment

1 Introduction

The European Union (EU) aims to achieve climate neutrality by 2050, with an interim target of reducing GHG emissions by at least 55% by 2030 compared to 1990 levels [1]. Sweden, in particular, has committed to achieving net-zero emissions by 2045, one of the most ambitious climate targets globally [2]. Given that buildings account for approximately 40% of the EU's energy use and 36% of its CO₂ emissions, the transformation of the sector is imperative [3].

Over 50% of buildings in Sweden and the EU are over 40 years old, with ~75% of the EU's stock being energy inefficient, underscoring the need for renovation to meet climate targets [3]. Older buildings typically exhibit poor insulation, high levels of air leakage, and inefficient energy systems. Without significant intervention, the building sector risks failing to meet the climate neutrality goals of both the EU and Sweden. In response, the EU launched the Renovation Wave initiative, which aims to renovate 35 million buildings by 2030, with the goal of reducing building-related emissions by 60% and lowering energy consumption by 14% [4]. In Sweden, the national strategy includes investments through the Recovery and Resilience Plan, as well as programs such as *Klimatkivet*, which focus on improving energy efficiency and decarbonizing the building stock [5].

Achieving climate-neutral cities necessitates a paradigm shift in the renovation of existing buildings, particularly through the integration of life cycle thinking and decarbonization strategies. Energy renovation is especially critical in cold climatic regions, where space heating accounts for a significant portion of energy use. However, the embodied carbon associated with energy efficiency measures presents a potential trade-off, emphasizing the need for sustainable, low-carbon renovation approaches.

Furthermore, renovation strategies in cold climates often prioritize winter thermal comfort, resulting in thermal discomfort during summer months [6, 7]. A significant portion of older Swedish buildings also lack active ventilation systems, contributing to poor indoor air quality [8]. Previous studies have predominantly focused on improving energy performance and reducing GHG emissions, frequently overlooking the indoor environmental quality [9, 10]. Additionally, the integration of renewable energy systems and energy storage technologies in retrofits has been identified as a research gap [10]. Addressing these challenges requires interdisciplinary research to develop holistic renovation strategies that combine emissions and cost reduction with improvements in indoor environmental quality.

This study investigates energy renovation strategies for a congregation building located in southern Sweden, with a particular focus on solar energy integration. The case study involves a two-story parish house with a basement, affiliated with the Västra Karup Church in Båstad Municipality, Skåne County. The geographical coordinates of the site are 56.4117057°N, 12.7426024°E. The objective is to evaluate combinations of energy efficiency and renewable energy measures that minimize energy consumption, reduce life cycle CO₂ emissions and costs, and enhance indoor environmental quality.

2 Methodologies

The workflow in Fig. 1 began with model verification, active and passive measures for alternative inputs into energy renovation. Based on the simulation results and objectives, the optimizer cycled through all the possible energy renovation scenarios to find the optimal cases.

Rhino 7 and ClimateStudio tools inside Grasshopper were utilized to create the geometrical and thermal models [11, 12]. Assumptions of building construction and usage were based on a site visit, and the model was verified against one year of actual electricity usage. Thermal comfort was assessed using the same tool, assuming a metabolic rate of

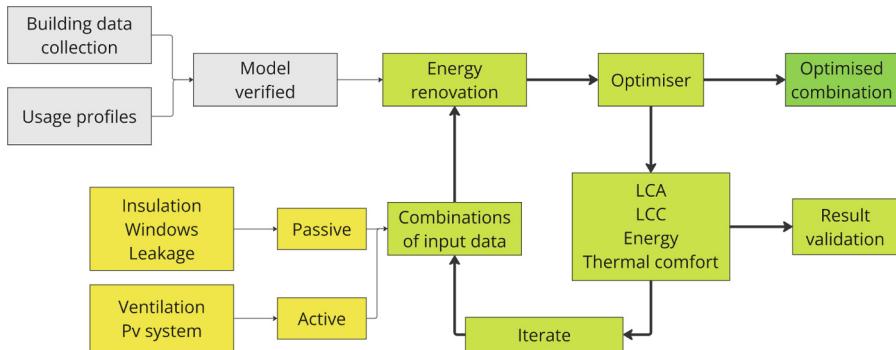


Fig. 1. Workflow stages: base case verification (grey), improvement of input data (yellow). Optimizer iteration cycle and optimized combination (green). (Color figure online)

1.2 MET and dynamically adjusted clothing insulation. Several energy-efficient measures were considered to reduce the building's energy consumption. Measures included reduction of transmission losses through the building envelope, through windows, and also reduction of infiltration losses. A new ventilation system was considered to mitigate poor IAQ and to comply with local regulations. Furthermore, photovoltaic systems were considered for the integration of green electricity generation.

2.1 Passive Measures

Envelope insulation was chosen based on which options would be least invasive to the exterior building architecture, due to possible cultural protection limitations. However, moisture aspects as well as applicability were also taken into consideration. The suggestion was to insulate the preexisting concrete constructions in the cellar and crawlspaces (see Fig. 2) with inputs in Table 1. The exterior placement was deemed necessary due to the existing moisture problems currently limiting the cellar space usage. The chosen insulation material was motivated by its combined capillary breaching, draining, and insulating properties [13].

Two options were suggested for window improvements, as stated in Table 2. Upgrading the existing two-pane windows with an extra pane was considered sustainable due to the extended lifespan. While the interior energy glass would decrease transmission losses, infiltration would also be limited since covering all mullions and glasses with one whole pane is financially motivated. Window renovation could be hard to estimate due to existing conditions, and it was therefore compared to new window replacements. The new window was assumed based on a non-openable model, which should be considered a best-case scenario for thermal performance. The need for emergency evacuation and natural ventilation would need to be taken into consideration.

An average infiltration leakage of 1.30 l/s/m^2 (at 50 Pa) was calculated based on measured CO_2 decay, which was, however, also validated by a statistical analysis [16]. The statistical analysis estimated an average leakage of 1.27 l/s/m^2 (at 50 Pa). Several literature studies showed that leakage could be reduced by up to 55% by regular leakage control measures. Within those studies, common measures were sealing of window and

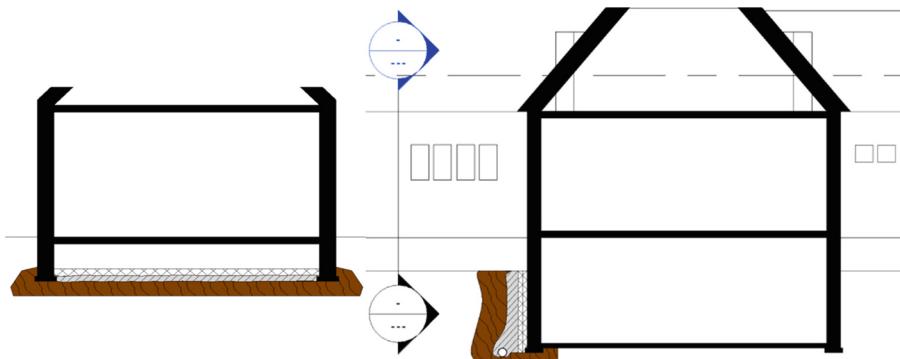


Fig. 2. Placement of envelope insulation measures. The left side shows a section of the kitchen with insulation on the crawl space floor. The right side shows the section through offices with insulation and drainage of cellar walls.

Table 1. Overview of construction thermal performance input data for insulation measures.

Alternatives	Cellar wall	Crawlspace wall	Crawlspace roof
	U-value W/m ² /K	U-value W/m ² /K	U-value W/m ² /K
Existing	0.544	0.544	3.509
Insulation 0.1 m	0.213	0.213	0.321
Insulation 0.2 m	0.134	0.134	0.170

Table 2. Overview of window performance inputs for the base case and renovation.

Alternative	U-value W/m ² /K	G-value	T-vis	Source
Existing	2.69	0.703	0.774	[11]
Upgrade	1.30	0.616	0.688	[14]
Replacement	0.91	0.530	0.740	[15]

door frames, but also building envelope penetrations due to cracks or building services. Draught-proofing measures to the external building envelope are comparatively cost-effective [17]. Self-expanding polyurethane foam and sealing tapes were considered for sealing. Considering the similar leakage-controlling measures taken in the studies mentioned in Table 3, reduction of 25% and 50% leakage was estimated for this study.

Table 3. Leakage reduction studies with respective potential for reduction.

Source	Year	Number of buildings	Location	Leakage reduction	Remarks
[18]	2024	1	Spain	44%	–
[17]	2016	1	UK	52%	–
[19]	2011	11	UK	50%	Highest
[20]	2003	26	UK	55%	Highest
[21]	1998	465	US	25%	Average

2.2 Active Measures

Since the building was unequipped with solar energy generation, PV systems with various capacities were considered along with optional storage (battery) systems. The PV production was simulated using ClimateStudio, and the hourly production data was used for parametric analysis for system optimization [11]. The southern part of the roof had a slope of $\sim 47^\circ$, an azimuth of 15° (South-East), and a tree shading from the south side. For impact assessment, a shading analysis was performed with Ladybug tools in Grasshopper [22]. The analysis showed considerable shading problems on the proposed PV module placement, especially during winter when the solar altitude is low. With PV optimizers, it would be possible to mitigate the shading effects on the whole system. However, the tree was assumed to be cut down to eliminate the shading problem. The PV production outputs were then validated by comparison with simulations in SAM (The System Advisor Model) [23]. The surplus energy produced from PV, after deducting the building's use, was considered to be sold to the grid. Therefore, a grid-tied solar system was designed.

The building was equipped with a mechanical exhaust ventilation system without heat recovery, and only one office on the top floor had a supply and exhaust room unit with heat recovery. The system was, however, quite small, while the building was facing serious IAQ issues. Measurements of CO_2 concentrations were collected for 54 days to evaluate the quality of indoor air. According to the design of the Swedish workplace regulation (AFS 2020:1), it should not exceed 1000 ppm [24]. The measurements showed that the CO_2 threshold exceeded 44.4% of the days in the main hall, compared to 16.7% and 14.8% in the top-floor offices and the basement, respectively. To ensure good IAQ, a mechanical supply and exhaust with heat recovery (FTX) was proposed in this study. Two different FTX systems were proposed: a demand-controlled ventilation (DCV) system and a constant air volume (CAV) system. According to the Swedish Housing Authority's building regulation (BBR29) and AFS 2020:1, a minimum air supply of 0.35 l/s/m^2 and 7 l/s per person should be provided [24, 25]. The airflow requirement in each room was calculated according to the regulations and the regular occupancy schedule. The existing exhaust ducts were reused, aiming for a cost-effective solution. The size of the ducts, associated pressure drops, and velocity according to the airflow were calculated using the ASHRAE circular duct chart [26]. The specific fan power (SFP) was calculated from the pressure drop, airflow, and the efficiency of the system. The Swedish Energy

Agency's network for energy-efficient premises (BELOK) suggested limiting the SFP to 1.3 kW/m³/s, and BBR29 suggested limiting it to 1.6 kW/m³/s [25, 27]. Therefore, the goal was to design an energy-efficient ventilation system to comply with those regulations, including the CO₂ limit.

2.3 Life Cycle Cost

The LCC was calculated using the net present value (NPV) method with the following Eqs. (1)–(3).

$$LCC = \sum_{i=1}^I NPV_i \quad (1)$$

NPV of expenditure occurring after n years –

$$NPV_i = C_i / (1 + r_{di})^n \quad (2)$$

NPV of a recurring cost that happens at the end of every year until year N –

$$NPV_i = C_i \cdot ((1 + r_{di})^N - 1) / (r_{di} \cdot (1 + r_{di})^N) \quad (3)$$

where:

LCC = life cycle cost [SEK]

NPV_i = net present value [SEK]

n = number of years after which the one-time cost occurs [years]

N = total number of years the recurring cost is incurred [years]

C_i = expenditure in today's value [SEK]

r_{di} = discount interest rate [%]

The material and manpower costs were collected from Wikells Sektionsfakta and validated with the quotations collected from different vendors [28]. Costs for primary and secondary material, equipment, manpower, project management, including administration and associated VATs, were considered. The costs for investment, annual maintenance, periodic maintenance, and replacement were calculated using the Excel program. The NPV and LCC calculations were programmed in the energy simulation script in Grasshopper and used for parametric optimization. The scrap or the end-of-life value of the materials was excluded from the LCC.

Discount interest rates (Table 4) were estimated for specific items to get a fair result. Initially, historical data for material price change was collected, and then future price changes were estimated following the price indexes. The inflation rate was considered 2.56%, taking the last 35 years' average inflation in Sweden, according to Statistics Sweden [29]. A nominal interest rate of 2.40% was considered according to Swedbank's one-year fixed interest rate [30]. A sensitivity analysis was carried out considering a nominal interest rate of 3.90%, 2.40%, and 0.90% due to the uncertainty of the future financial scenario.

The hourly electricity spot prices for 2023 were collected from Energi Data Service to estimate the operational energy cost of the building [31]. The monthly average spot prices were compared to the actual energy bills of the building, and an average variation of 6% was found. However, hourly prices were not available in the energy bills, and the data from Energi Data Service were considered in the calculation. Due to the significant fluctuations in the electricity price index, four different nominal price changes (1.0%, 0.5%, 0%, and -0.5%) were considered for a sensitivity analysis, and a 0.5% nominal price change was used in the parametric simulation. The electricity buying prices from the grid and selling prices to the grid were calculated on an hourly basis according to the following Eqs. (4) and (5).

Electricity buying price

$$\begin{aligned}
 &= \text{Spot price} + \text{Energy tax} + \text{Transmission fee} \\
 &+ \text{Electricity certificate fee} + \text{VAT}
 \end{aligned} \quad (4)$$

$$\text{Electricity selling price} = \text{Spot price} + \text{Grid benefit} + \text{Tax deduction} \quad (5)$$

Table 4. Discount interest rates for life cycle cost calculations.

Material	Insulation	Window	PV System
$r_{di}/\%$	1.39	0.89	3.64
Material	Li-ion Battery	Foam Sealant	HVAC Materials
$r_{di}/\%$	7.79	3.64	1.39

Estimated life cycle costs for each energy renovation measure are given in Table 5.

Table 5. Estimated life cycle costs for each energy renovation measure.

Description	Alternatives	Investment	Maintenance		Life Span
			Annual	Periodic	
		SEK	SEK/year	SEK/years	/Year
Insulation Cellar wall	0.1 m	148 000	–	–	50
	0.2 m	160 000	–	–	50
Insulation Crawlspace	0.1 m	40 000	–	–	50
	0.2 m	71 000	–	–	50
Window	Upgrade	153 000	–	42 000/8	25
	Replacement	395 000	–	4 000/10	40
Infiltration Reduction	25%	27 000	–	13 500/10	40
	50%	27 000	–	13 500/5	40

(continued)

Table 5. (continued)

Description	Alternatives	Investment	Maintenance		Life Span
			Annual	Periodic	
		SEK	SEK/year	SEK/years	/Year
Ventilation	DCV	554 000	10 000	20 000/5	40
	CAV	445 000	6 000	12 000/5	40
PV	10 kW _p	143 500	–	–	40
	18 kW _p	270 500	–	–	40
Battery	10 kWh	50 000	–	–	10
	20 kWh	100 000	–	–	10

2.4 Life Cycle Assessment

Aiming for the climate neutrality goal of Sweden by 2045, the Global Warming Potential (GWP) was assessed for the energy renovation measures and associated operational energy use. The environmental product declaration data (EPD) of primary materials were collected from several EPD databases. The GWP potential related to operational energy use was collected from the ELCD database [32]. The LCA was completed for a building lifespan of 40 years from now matching the lifespan of most of the products and systems. A cradle-to-grave model considering the product stage (A1–A3), construction stage (A4–A5), use stage (B1–B6), and end-of-life stages (C1–C5) was used for the assessment. The GWP for each item was first calculated in Excel, and the rest of the LCA was programmed in Grasshopper, which was later connected to parametric simulation for optimization.

2.5 Parametric Optimization

There were seven different potential improvement categories with three alternatives each, as mentioned in Table 6. There were also four different objectives – energy use intensity (EUI), LCC, LCA, and thermal comfort. Working with an individual variable could lead to a contradictory output for various objectives. Therefore, it was necessary to consider all 2 187 possible combinations to get the exact scenario which would help to evaluate the most optimized combinations and decision making. This parametric optimization was done using Opossum, an open-source Grasshopper plugin [33].

Table 6. Parameter optimization variables.

Parameters	Alternatives		
Insulation – Cellar wall	0 m	0.1 m	0.2 m
Insulation – Crawlspace	0 m	0.1 m	0.2 m

(continued)

Table 6. (continued)

Parameters	Alternatives		
Window	Existing	Upgrade	Replacement
Infiltration Reduction	0%	25%	50%
Ventilation	Existing	DCV	CAV
PV	N/A	10 kW _p	18 kW _p
Battery	N/A	10 kWh	20 kWh

2.6 Regulations and Environmental Certification

Buildings with a floor area of more than 250 m² that are frequently visited by the public are recommended to have an energy declaration in Sweden [25]. The declaration provides information on the building's primary energy use and energy class (A–F). The primary energy number (EP_{pet}) was calculated according to BBR29 [25]. Additionally, the building was studied for environmental certification, Miljöbyggnad 4.0 [34]. To get a gold certification in "indicator 3 – energy use", the requirements are to limit the EP_{pet} within 70% of the BBR recommendation, and more than 5% of the property energy should be locally generated with renewable energy. For silver rating, the EP_{pet} requirement is 80% and for bronze, the BBR requirement should be met at least. There are no renewable energy requirements for silver and bronze.

3 Results and Discussion

3.1 Parametric Optimization

Based on the simulation results and objectives, the optimizer cycled through all the possible energy renovation scenarios. Of the 2 187 cases considered, 112 combinations were suggested by the optimizer, with the rest deemed worse by the optimizer algorithm. Figure 3 shows the LCA vs. LCC for each simulated case. Thermal comfort was evaluated as the percentage of people dissatisfied (PPD). Clustered in red are the cases with worse thermal comfort and consequently lower EUI, LCC, and LCA. Yellow and orange cases had better thermal comfort, with higher EUI, LCC, and LCA. Figure 4 illustrates the ventilation system installed in each of the simulated cases, with the current system both better in terms of LCC and LCA, but at the cost of thermal comfort, as seen in Fig. 3. As evident by comparing Fig. 3 and 4, cheaper cases in terms of both LCC and LCA did not have a new ventilation system installation, and more resource-intensive cases did. All cases with better thermal comfort and good IAQ had a new ventilation system installed.

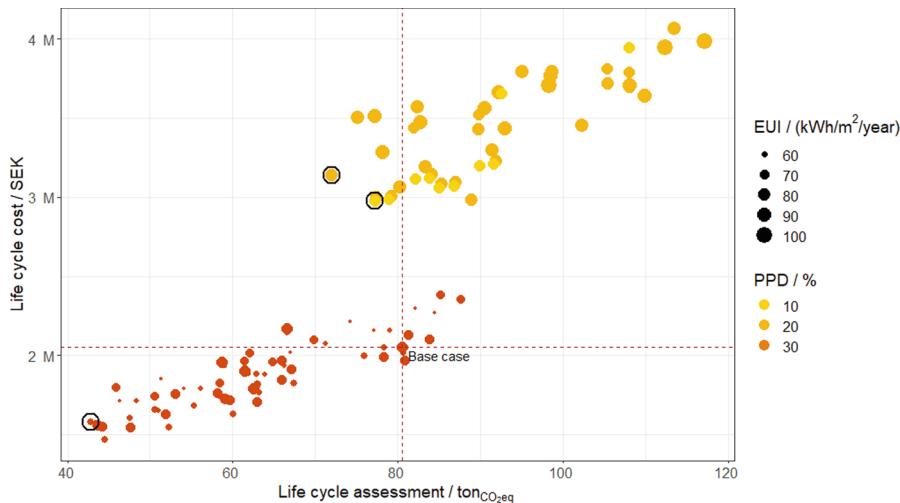


Fig. 3. LCA vs. LCC with color as thermal comfort (lower is better) and point size as the EUI. Each point represents a simulated case. Recommended cases are circled.

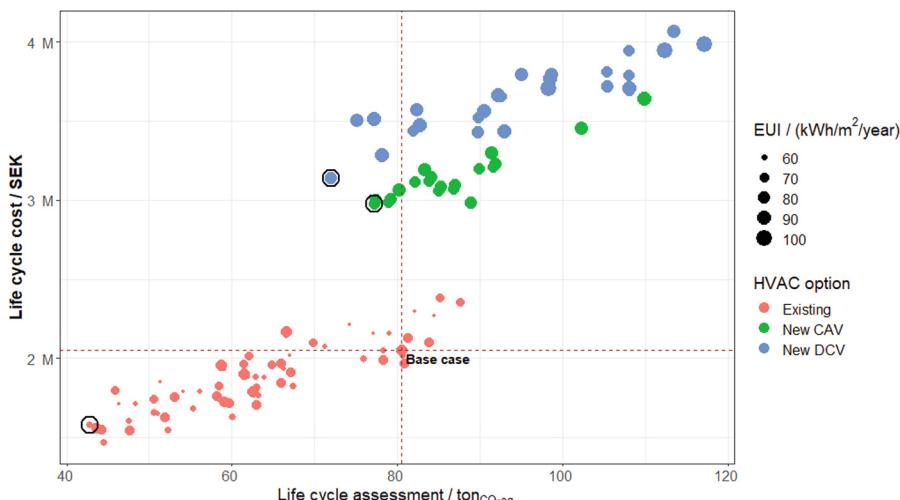


Fig. 4. LCA vs. LCC with color as the ventilation system type and point size as the EUI. Each point represents a simulated case. Recommended cases are circled.

Fig. 5 analyses the PV system of each case simulated. The lower cluster with the current ventilation system highly preferred the larger system without a battery, while the higher cluster with new ventilation systems preferred the larger system, but both without and with the smaller battery.

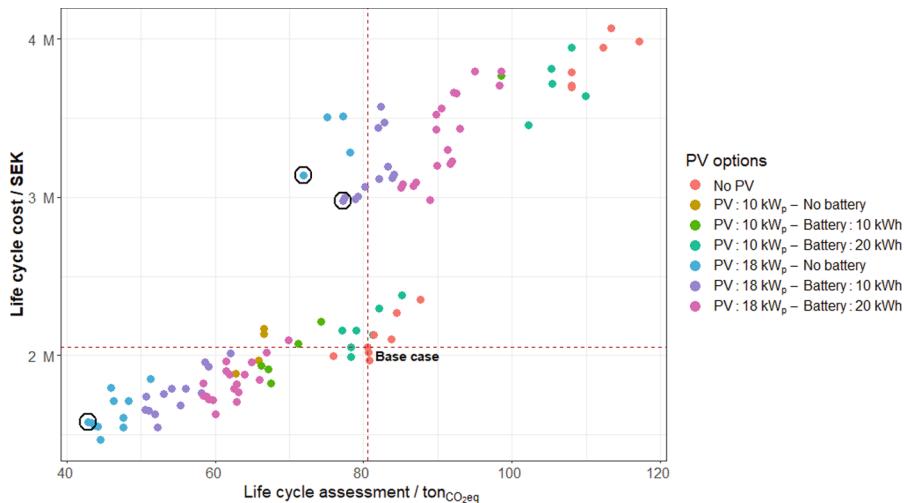


Fig. 5. LCA vs. LCC with color as PV system. Each point represents a simulated case. Recommended cases are circled

3.2 Sensitivity Analysis

Figure 6 explores the effect of the yearly price increase of electricity being different from what was considered during the life cycle costing. As expected, the lower cluster was less sensitive to price increases due to less operational electricity use. Figure 7 illustrates the impact of the nominal interest rate on each simulated scenario. With higher interest rates, some of the cases with new ventilation systems became more financially feasible.

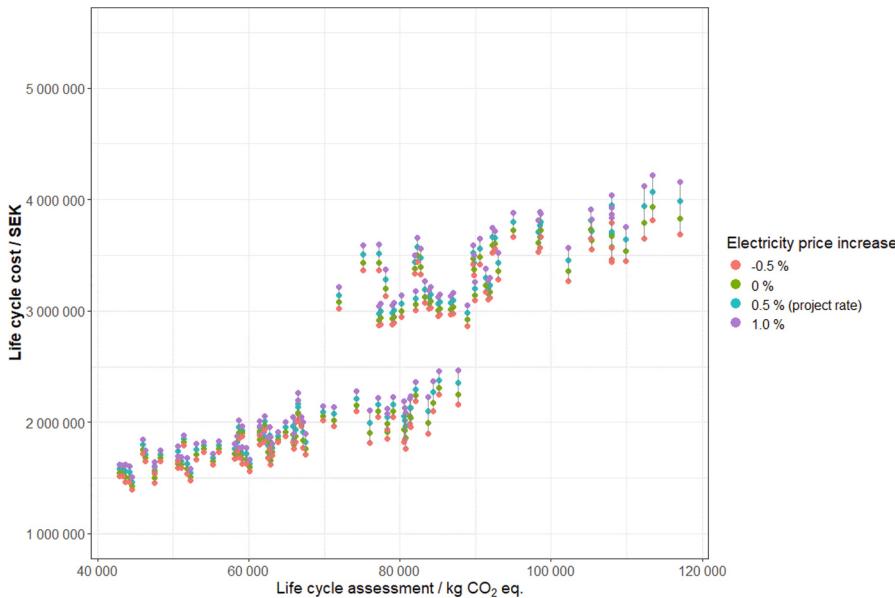


Fig. 6. LCA vs. LCC with color as the expected electricity price increase over the analyzed period.

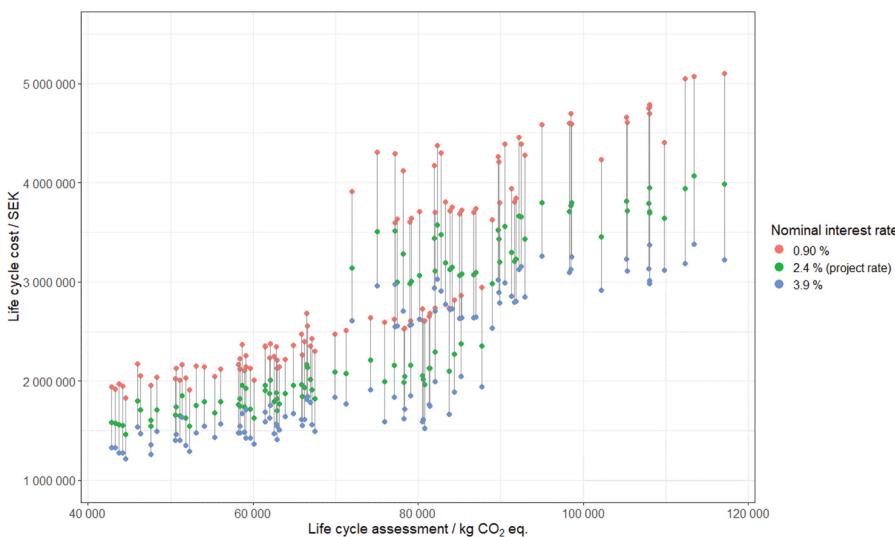


Fig. 7. LCA vs. LCC with color as the expected nominal interest rate.

The most optimized and recommended combinations are illustrated in Table 7. Most optimized and recommended combinations., and marked with circles in Figs. 3, 4, and

5. From the sensitivity analysis of electricity price changes and nominal interest rate changes, the recommended cases did not change.

Table 7. Most optimized and recommended combinations.

Parameter	Base Case	Case 983	Case 543	Case 1059
Ventilation	Existing	Existing	CAV	DCV
PV System/kW _p	N/A	18	18	18
Battery/kWh	N/A	N/A	10	N/A
Windows	Existing	Upgrade	Upgrade	Existing
Cellar insulation/m	N/A	N/A	0.1	N/A
Crawlspacel insulation/m	N/A	0.1	0.2	0.2
Leakage Reduction/%	N/A	50	50	50
EUI/kWh/m ² A _{temp} /year	80	61	77	80
LCC/MSEK	2	1.6	3	3.1
LCA/ton CO ₂ eq.	81	43	77	72
PPD/%	38	35	10.5	10

3.3 Ventilation

The calculated highest CO₂ concentrations in different rooms were 733 PPM to 988 PPM with the proposed new ventilation systems, which complied with the Swedish workplace regulation (AFS 2020:1) [24]. The calculated SFP of the system was 0.96 kW/m³/s which complied with BELOK and BBR regulations [25, 27].

3.4 Regulations and Environmental Certification

The calculated EP_{pet} of the existing building was 144 kWh/m²A_{temp}/year (BBR reference value: 70 kWh/m²A_{temp}/year), which fell into energy class F. Without considering a new ventilation system, the most optimized combination had 72 kWh/m²A_{temp}/year (BBR reference value: 70 kWh/m²A_{temp}/year) EP_{pet}, corresponding to energy class D. Considering the most optimized case with a new ventilation system, the value was 94 kWh/m²A_{temp}/yr (BBR reference value: 96 kWh/m²A_{temp}/yr), corresponding to energy class C [25]. With the new PV system, the building met the Miljöbyggnad gold requirement with a PV energy usage of more than 5% of total energy use for the optimized cases, but the EP_{pet} could not reach beyond the bronze rating [34]. With other combinations, it had the potential to meet the gold criteria for energy requirements, but the LCC, LCA, thermal comfort, and IAQ were not optimized, corresponding to the energy savings.

4 Conclusions

With the proposed optimized energy renovation plans with the existing ventilation system, it was possible to reduce the EUI by 24%. Besides, it was also possible to reduce the LCC by 23% and carbon emissions by 47%. However, the existing IAQ problem could not be solved with those measures. Moreover, there was a negligible improvement in thermal comfort, with only a 7% reduction in PPD from the base case.

With the proposed optimized solution, including a new ventilation system, there was no considerable improvement in the EUI due to the increased energy consumption for the fan operation. The LCC went up because of the high investment and maintenance costs of the new ventilation system. The LCC increased by 53% and 45% for the optimized DCV and CAV systems, respectively. However, the carbon emissions could be reduced by 11% and 4% for the optimized DCV and CAV systems, respectively. However, the new ventilation system could solve the IAQ problem, and the thermal comfort improved by 72% and 74% with the optimized DCV and CAV systems, respectively. It was also possible to reach up to energy class C according to Boverket and to achieve a bronze rating in “indicator 3 – energy use” according to Miljöbyggnad. The methodology used in this study offers a holistic, multi-objective approach that is applicable to all building types in cold climate regions.

There were some limitations in the study, which included multiple assumptions due to data unavailability. A simplified approach with limited variables was used to run the parametric optimization simulator; there were also simulation program limitations. There was a lack of environmental declarations of secondary materials for LCA and uncertainties with financial parameters for LCC. The 54-day CO₂ measurements offer useful insights but may not reflect full seasonal variation. However, this limitation does not impact the study’s outcome, as the proposed ventilation system effectively addresses the IAQ problem. Future studies that could be done from this study include a PV production analysis with shading from the tree (a solution without cutting down the tree), LCA of the full building to meet the BBR and Miljöbyggnad recommendations, and a moisture analysis for the basement and the crawlspace with optimized cases.

Acknowledgements. The project received no external funding. The authors would like to express their gratitude to Joakim Olsson and Ricardo Bernardo for their valuable expert insights during dedicated consultation sessions.

References

1. European Commission. Overall Targets and Reporting. European Commission – Energy, Climate Change & Environment. https://commission.europa.eu/energy-climate-change-environment/overall-targets-and-reporting_en. Accessed 24 June 2025
2. Government Offices of Sweden. Sweden’s Climate Policy Framework. Ministry of Climate and Enterprise, Stockholm (2021). <https://www.government.se/articles/2021/03/swedens-climate-policy-framework/>. Accessed 24 June 2025
3. European Commission. Energy Performance of Buildings Directive. European Commission – Energy Efficiency (2024). https://energy.ec.europa.eu/topics/energy-efficiency/energy-performance-buildings/energy-performance-buildings-directive_en. Accessed 07 July 2025

4. European Commission. Renovation Wave. European Commission – Energy (2020). https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficient-buildings/renovation-wave_en. Accessed 24 June 2025
5. European Commission. Sweden's Recovery and Resilience Plan. European Commission – Recovery and Resilience Facility (2023). https://commission.europa.eu/business-economy-euro/economic-recovery/recovery-and-resilience-facility/country-pages/swedens-recovery-and-resilience-plan_en. Accessed 07 July 2025
6. Sukanen, H., Taylor, J., Castaño-Rosa, R., Pelsmakers, S., Lehtinen, T., Kaasalainen, T.: Passive mitigation of overheating in Finnish apartments under current and future climates. *Indoor Built Environ.* **32**(7), 1372–1392 (2023). <https://doi.org/10.1177/1420326X231160977>
7. Joshi, K., Khan, A., Anand, P., et al.: Understanding the synergy between heat waves and the built environment: a three-decade systematic review informing policies for mitigating urban heat island in cities. *Sustainable Earth Rev.* **7**, 25 (2024). <https://doi.org/10.1186/s42055-024-00094-7>
8. Abdul Hamid, A., Johansson, D., Bagge, H.: Ventilation measures for heritage office buildings in temperate climate for improvement of energy performance and IEQ. *Energy Build.* **211**, 109822 (2020). <https://doi.org/10.1016/j.enbuild.2020.109822>
9. Abdul Hamid, A., Farsäter, K., Wahlström, Å., Wallentén, P.: Literature review on renovation of multifamily buildings in temperate climate conditions. *Energy Build.* **172**, 414–431 (2018). <https://doi.org/10.1016/j.enbuild.2018.04.032>
10. Chen, W., Lai, J.: Performance assessment of residential building renovation: a scientometric analysis and qualitative review of literature. *Smart Sustainable Built Environ.* **14**(3), 625–648 (2025). <https://doi.org/10.1108/SASBE-09-2023-0276>
11. Solemma LLC. ClimateStudio. <https://www.solemma.com/climatestudio>. Accessed 25 June 2025
12. Robert McNeel and Associates. Grasshopper – Algorithmic Modeling for Rhino.” Rhino3D.com. <https://www.rhino3d.com/#:~:text=Kay%2C%20Kiteboard%20Shaper-,Grasshopper,very%20rich%20add%20Don%20library>. Accessed 07 July 2025
13. Isodrän. Teknisk fakta. <https://www.isodran.se/dokument>. Accessed 15 December 2024
14. Grundels. Energiglas eller energifönster – Vad är skillnaden? Grundels (2024). <https://www.grundels.se/fabriksproduktion/energiglas-energifonster/>. Accessed 15 December 2024
15. Svenska Fönster. Fast fönster aluminiumbeklätt trä 3-glas: Teknisk specifikation. Svenska Fönster (2024). <https://svenskafonster.se/aluminiumfonster/fasta-fonster/teknisk-specifikation>. Accessed 15 December 2024
16. Zou, Y.: Classification of Buildings with Regard to Airtightness, M.S. Thesis. Chalmers University of Technology, Gothenburg, Sweden (2010). <https://odr.chalmers.se/bitstreams/ba91ebb4-b79c-4d9f-ae95-4bc633522464/download>. Accessed 07 July 2025
17. Gillott, M.C., Loveday, D.L., White, J., Wood, C.J., Chmutina, K., Vadodaria, K.: Improving the airtightness in an existing UK dwelling: the challenges, the measures and their effectiveness. *Build. Environ.* **95**, 227–239 (2016). <https://doi.org/10.1016/j.buildenv.2015.08.017>
18. Echarri-Iribarren, V., Gómez-Val, R., Ugalde-Blázquez, I.: Energy losses or savings due to air infiltration and envelope sealing costs in the Passivhaus standard: a review on the Mediterranean coast. *Buildings* **14**(7), 2158 (2024). <https://doi.org/10.3390/buildings14072158>
19. Johnston, D., Miles-Shenton, D., Bell, M., Wingfield, J.: Airtightness of Buildings – Towards Higher Performance: Final Report – Domestic Sector Airtightness, Project Report. Leeds Beckett University, Leeds, UK (2011)
20. Borland, S., Bell, M.: Airtightness of Buildings – Towards Higher Performance: A Project Proposal to Communities and Local Government Building Regulations Division under the

Building Operational Performance Framework, Report No. CI 61/6/16 (BD2429), Building Sciences Ltd. (2003)

21. Sherman, M.H., Dickerhoff, D.J.: Airtightness of U.S. dwellings. *ASHRAE Trans.* **104**, 1359–1367 (1998). https://www.aivc.org/sites/default/files/airbase_11887.pdf. Accessed 07 July 2025
22. Ladybug Tools. Home – Ladybug Tools. <https://www.ladybug.tools/>. Accessed 07 July 2025
23. National Renewable Energy Laboratory. System Advisor Model (SAM) – Welcome. SAM. <https://sam.nrel.gov/>. Accessed 07 July 2025
24. Arbetsmiljöverket. Arbetsplatsens utformning (AFS 2020:1) – Föreskrifter. <https://www.av.se/arbetsmiljoarbete-och-inspektioner/publikationer/foreskrifter/upphavda-foreskrifter/arbetsplatsens-utformning-afs-20201-foreskrifter/>. Accessed 07 July 2025
25. Boverket. Boverkets byggregler, BBR, BFS 2011:6 med ändringar till och med BFS 2020:4, 2020. <https://www.boverket.se/sv/lag--ratt/forfatningssamling/gallande-foreskrifter-och-allmanna-rad/bfs-20116-med-andringar-till-och-med-bfs-20204/>. Accessed 07 July 2025
26. ASHRAE. ASHRAE Handbook – Fundamentals (SI Edition), ASHRAE, GA (1997)
27. Belok. Beloks Energikrav Version 8 (2022). <https://belok.se/wp-content/uploads/2022/12/Beloks-Energikrav-Version-8.pdf>. Accessed 07 July 2025
28. Wikells Byggdetaljer AB. Sektionsfakta: Construction Price Database. <https://www.wikells.se/sektionsfakta>. Accessed 24 June 2025
29. Statistikmyndigheten SCB. Statistikmyndigheten SCB. <https://www.scb.se>. Accessed 07 July 2025
30. Swedbank. Base Interest Rates (2024). <https://www.swedbank.ee/private/home/more/prices/rates/loaninterests>. Accessed 21 December 2024
31. Energi Data Service. Open Energy Data from Energinet to Society (2024). <https://www.energidataservice.dk>. Accessed: 21 December 2024
32. OpenLCA. ELCD Database (2024). <https://nex-us.openlca.org/database/ELCD>. Accessed 21 December 2024
33. Opossum. OOptimizatiOn Solver with SURrogate Models (2024). <https://opossum-optimizer.com/>. Accessed 07 July 2025
34. Sweden Green Building Council. Miljöbyggnad – Environmental Certification of Buildings (2023). <https://www.sgbc.se/certifiering/miljobyggnad/>. Accessed 07 July 2025

Building Simulation, Urban Energy, and Environmental Sensing



Development of an Automated 3D Building Modeling Method and Urban-Scale Analysis of Heating and Cooling Loads

Hisato Osawa^(✉) , Taro Mori , and Konatsu Suzuki 

Faculty of Engineering, Hokkaido University, Sapporo 060-6826, Japan
osawa.hisato.z7@eng.hokudai.ac.jp

Abstract. To achieve carbon neutrality (CN) by 2050, it is essential to develop systems that enable detailed and spatially explicit regional analyses, allowing for an accurate understanding of current urban energy consumption. In this study, we aim to develop a system that constructs a wide-area building database by utilizing geospatial information and enables the visualization and utilization of this data through a WebGIS platform. This paper proposes a method for the automatic generation of 3D building models based on GIS data, which are compatible with air-conditioning load calculations, and applies this approach to energy analysis on an urban scale. For the building types of office, accommodation, dedicated store, and detached house, we defined the conditions necessary for the automatic generation of EnergyPlus-compatible 3D models and determined their geometry based on the arrangement of surrounding buildings. By analyzing the heating and cooling loads of each building using the generated 3D models and visualizing the results on a map, our findings suggest that combining a single office building with multiple detached houses may contribute to the leveling of heating and cooling demand.

Keywords: Geographic Information System (GIS) · Simulation · Automation Modeling · Air Conditioning Load · Carbon Neutral

1 Introduction

Achieving carbon neutrality (CN) by 2050 stands as a globally shared objective. The Japanese government, under its “GX (Green Transformation) Promotion Strategy” [1], aims to simultaneously realize stable energy supply, economic growth, and decarbonization. The Cabinet Office and the Ministry of the Environment have formulated a strategy to promote region-led decarbonization, aiming to achieve successful decarbonization in at least 100 municipalities by fiscal year 2030 and to disseminate these pioneering initiatives nationwide by fiscal year 2050 [2]. However, small municipalities face substantial barriers due to shortages in technical expertise and human resources, limiting their capacity for independent decarbonization efforts. Furthermore, many existing success stories involve costly new technologies or large-scale facility renovations, making direct replication by other municipalities financially impractical.

While current support tools like the Cabinet Office’s “RESAS” [3] and the Ministry of the Environment’s “Japan Energy Database” [4] provide macroeconomic and energy consumption insights, these platforms primarily offer statistical overviews. To create effective action plans for carbon neutrality, we need more detailed and location-specific regional analyses. However, current systems do not have this capability.

This study addresses this gap by developing a geospatial information system for constructing wide-area building databases and visualizing them through WebGIS. Previous research demonstrated energy consumption estimation using simplified floor plans derived from GIS building footprints and positional data [5], as well as thermal environment simulations integrating aerial photography-based tree identification with GIS data to generate building-specific meteorological profiles [6]. As demand control becomes critical for city-scale decarbonization, this paper proposes a novel methodology for automatically generating 3D models from GIS data to enable air-conditioning load calculations. We further present urban-scale energy analysis results derived from this approach, advancing the discourse on spatially resolved decarbonization strategies.

2 Method

This chapter proposes a method for automatically generating EnergyPlus [7] simulation models using geographic spatial information from PLATEAU [8]. PLATEAU is a 3D urban model development, utilization, and open data project led by the Ministry of Land, Infrastructure, Transport and Tourism of Japan. It enables visualization of urban structures through 3D models based on urban planning fundamental survey information maintained by local municipalities. The platform is being utilized for urban disaster prevention and as a metaverse integration platform, with development progressing across municipalities throughout Japan.

2.1 Target Area

The buildings analyzed in this study are in Sapporo City, Hokkaido, Japan. Sapporo is the prefectural capital of Japan’s northernmost prefecture, Hokkaido, with a population of approximately 2 million residents. The city experiences cool summers and snowy winters. Buildings within the red boundary shown in Fig. 1 were selected as calculation targets. Table 1 presents the distribution of building uses within this boundary, with “NA” indicating buildings of unknown use. Excluding buildings with undefined uses, this area is predominantly residential with commercial buildings constituting the second largest category, characterizing it as a typical residential district. While the PLATEAU dataset includes auxiliary structures such as warehouses and garages, only primary buildings were considered for calculations. This study specifically targets business facilities, accommodation facilities, and dedicated retail facilities classified as commercial buildings, and detached houses from the residential classification.



Fig. 1. Target Area (Color figure online)

Table 1. Building Use and Number

Residential	Commercial	Industrial	Public	Other	NA
17,508	2,093	320	516	5	9,954

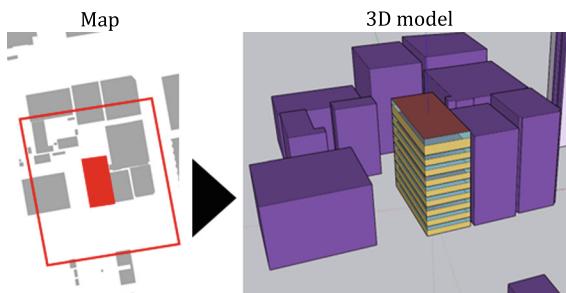
2.2 3D Modeling of Commercial Buildings

Among the commercial building classification, this study specifically targeted office, accommodation, and dedicated store for calculation purposes. The dataset comprised 955 offices, 144 accommodations, and 710 dedicated stores. This section details the methodology employed for automatic 3D model generation of these building types.

Generation of Building Geometry. Simulation models were automatically generated based on building polygon information and height data from PLATEAU. Floor count was referenced from PLATEAU's floor information; when unavailable, it was calculated by dividing the measured height by a minimum of 4 m per floor. Window area was determined using the window-to-wall ratio values presented in Table 2, which vary according to building use and wall orientation, based on previous research [9]. Windows were omitted for wall surfaces with adjacent buildings within 4 m proximity or for wall widths less than 3 m. Surrounding structures within 50 m of each target building were also modeled to account for shading effects. Figure 2 illustrates an example of the automatically generated model.

Table 2. Assigned Window Area Ratios by Building Use and Orientation

Direction	Office	Accommodation	Dedicated Store	Others
West	0.332	0.673	0.56	0.006
East	0.298	0.045	0.59	0.368
South	0.407	0.05	1	0.341
North	0.463	0.081	0.51	0.333

**Fig. 2.** Automatically generated 3D model

Calculation Parameters. In Japan, WEBPRO is utilized for evaluating energy consumption performance of new constructions or major renovations in compliance with the Building Energy Efficiency Act. The calculation load parameters for commercial buildings in this study were based on WEBPRO calculation conditions [10], as shown in Table 3 and Table 4. As indicated in Table 4, annual schedules are categorized into patterns A-F, each composed of three distinct 24-h operational schedules with different room specifications. These 24-h schedules vary according to the use classification numbers presented in Table 3, with detailed specifications referenced from technical documentation [11].

Table 3. Calendar Patterns and Corresponding Room Usage Schedules

Calendar Pattern	Room Usage Pattern			Remarks
	1	2	3	
A	Weekdays	Saturdays	Year-end/New Year National Holidays	
B	Weekdays	Weekends National Holidays Long Vacations	Year-end/New Year	Assumed for Junior High Schools

(continued)

Table 3. (continued)

Calendar Pattern	Room Usage Pattern			Remarks
	1	2	3	
C	Weekdays	Weekends National Holidays Long Vacations	Year-end/New Year	Assumed for Universities
D	Weekdays	Weekends National Holidays	Year-end/New Year 1 weekday off per month	
E	Weekdays	Weekends National Holidays	Year-end/New Year 1 weekday off per week	
F	Weekdays	Weekends National Holidays	2 consecutive weekdays off per month	

Table 4. Internal Load and Occupancy Parameters by Building Use

Building Use	Use No	Calendar Pattern	Lighting Heat Load [W/m ²]	Human Heat Load [W/m ²]	Equipment Heat Load [W/m ²]	People [persons/m ²]
Office	O-1	A	12	11.9	12	0.1
Municipal	O-1	A	12	11.9	12	0.1
Assembly/Sales	H-1	A	15	24.4	40	0.2
Accommodation	H-1	A	15	6.44	0.7	0.07
Dedicated Store	S-2	A	30	23.8	40	0.2
Local/National	O-1	A	12	11.9	12	0.1
Bathhouse Business	R-4	A	10	23.8	0	0.2
Amusement	M-42	A	60	72.5	100	0.5
Medical	D-1	A	12	7.36	3	0.08
Gym	M-6	D	60	14.5	0	0.1
Memorial	M-29	E	15	5.7	0	0.03
Educational	E-8	A	20	23	8	0.3
Research	E-8	A	20	23	8	0.3
Welfare	H-1	A	15	6.44	0.7	0.07
Social Welfare	M-49	D	10	5	0.1	0.1
Religious	M-29	E	15	5.7	0	0.03
Cultural	M-29	E	15	5.7	0	0.03

The weather data was created by modifying EnergyPlus Weather Data (epw) parameters based on hourly meteorological values for Sapporo City published by the Japan Meteorological Agency.

As detailed in Table 5, exterior walls were modeled as a composite of 50 mm thermal insulation material with a thermal conductivity of 0.038 W/(m·K) and double-glazed windows. Ventilation rate was uniformly set at 0.5 air changes per hour (ac/h) for all simulations.

Table 5. Specifications of Building Envelope Components

Wall	Floor, Ceiling	Window
Outside: GW 50 mm	GW 100 mm	Low-e Double Glazing
Inside: Air Cavity		

2.3 3D Modeling of Residential Buildings

The residential buildings in Fig. 1 include subcategories such as detached and multi-family houses. This study focused on detached houses for automated model generation. Figure 3 illustrates the duration curve of building heights for detached houses in the target area. As shown, the dataset contains structures with extremely low and high heights, which may represent non-residential buildings. To focus on typical detached houses, only buildings with heights between 6.5 m and 10.5 m were selected, resulting in a total of 7,473 detached houses. This section describes the configuration methodology for their automated 3D modeling.

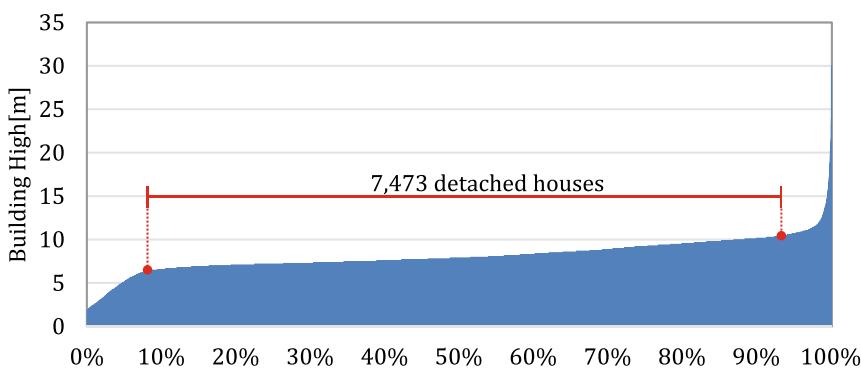


Fig. 3. Duration Curve of Building Heights in the Target Area: Focus on Detached Houses (6.5 m–10.5 m)

Generation of Building Geometry. Simulation models were automatically generated using PLATEAU polygon data and measured heights. Each house was modeled with the

first floor designated as a living area and the second floor as bedrooms. Floor heights were calculated by dividing the total measured height by the number of floors derived from PLATEAU. Following the commercial building methodology, windows were added to walls facing unobstructed spaces (no adjacent buildings within 4 m) with widths exceeding 3 m. Window area was set to 50% of the wall surface. Surrounding structures within 50 m of each target building were modeled to account for spatial context.

Organization of Household Composition and Daily Activity Schedules. To assess the impact of household composition on occupant-related loads, we analyzed the distribution of household types and their proportions in Sapporo's Chuo Ward, which corresponds to the study area, based on the 2020 National Census [12]. The results are summarized in Table 6. Couples only accounted for 43.51%, couples with children for 37.01%, and single parents with children for 17.18%, indicating that the majority of households in Chuo Ward fall into these three categories. In this study, the five most prevalent household types (No. 1, 2, 3, 5, and 7) were randomly assigned to each detached house according to their statistical ratios. For the purpose of modeling, couples were assumed to be male and female in their forties, children as male teenagers, and grandparents as males aged 70 or older. Individual daily activity schedules were defined with reference to the 2020 National Time Use Survey [13] provided by the NHK Broadcasting Culture Research Institute, as shown in Table 7.

Table 6. Household types and their proportions in Sapporo's Chuo Ward

No	No.1	No.2	No.3	No.4	No.5	No.6	No.7	No.8	No.9	No.10	No.11
Household Composition Pattern (Assumed Number of People)	Married Couple (2 people)	●	●		●	●	●	●	●	●	●
	Single Parent (1 person)			●							
	Grandparents (2 people)				●		●			●	●
	Grandfather or Grandmother (1 person)					●		●			
	Child (1 person)		●	●			●	●			●
	Other Relatives (1 person)							●	●	●	●
Number in Household	2	3	3	4	3	5	4	5	4	5	6
Number of Households	24,467	20,812	9662	31	450	52	374	111	234	20	23
Household Proportion	43.5%	37.0%	17.2%	0.1%	0.8%	0.1%	0.7%	0.2%	0.4%	0.0%	0.0%

Table 7. Individual daily activity schedules

		Weekday	Saturday	Sunday
Male Teens	Sleep	0:00–6:00	0:00–7:00	0:00–8:00
				23:00–24:00
	At-home Activities	6:00–8:00	7:00–11:00	8:00–14:00
		18:00–24:00	14:00–24:00	16:00–23:00
	Outing	8:00–18:00	11:00–14:00	14:00–16:00
Female 40s	Sleep	0:00–6:00	0:00–6:00	0:00–7:00
	At-home Activities	6:00–9:00	6:00–11:00	7:00–11:00
		18:00–24:00	12:00–14:00	17:00–24:00
			15:00–24:00	
	Outing	9:00–18:00	11:00–12:00	11:00–17:00
			14:00–15:00	
Male 40s	Sleep	0:00–6:00	0:00–7:00	0:00–7:00
	At-home Activities	6:00–7:00	7:00–9:00	7:00–11:00
		19:00–24:00	18:00–24:00	17:00–24:00
	Outing	7:00–19:00	9:00–18:00	11:00–17:00
Male 70 and above	Sleep	0:00–6:00	0:00–6:00	0:00–6:00
		23:00–24:00	23:00–24:00	23:00–24:00
	At-home Activities	6:00–10:00	6:00–23:00	6:00–10:00
		12:00–23:00		11:00–23:00
	Outing	10:00–12:00	なし	10:00–11:00

Calculation Conditions for Detached Houses. Internal loads were varied hourly based on the household compositions in Table 6 and the daily schedules in Table 7. Sensible heat gains from occupants were set at 100 W per person, and equipment heat gains in living rooms were set at 1,150 W/m². Lighting heat gains were set at 2 W/m², with the lighting schedule established according to the technical guidelines for residential energy efficiency standards provided by the Building Research Institute [14]. The building envelope composition and meteorological data were set in accordance with the parameters used for commercial buildings.

3 Results and Discussion

Heating and cooling load simulations were conducted using the automatically generated 3D models produced by the proposed methodology, and the results were analyzed accordingly.

Analysis of Heating and Cooling Loads. Figure 4 and Fig. 5 present boxplots of the annual total heating and cooling loads by building use. As shown in Fig. 4, when focusing on the annual load per unit area, offices exhibited the lowest values across the target area, while dedicated stores showed the highest values. These differences can be attributed to variations in building operation types, occupancy hours, and internal heat gain settings. In contrast, detached houses displayed greater variability compared to other building types, which is likely due to the diversity in household compositions. When examining the annual total heating and cooling load per building, as shown in Fig. 5, accommodation had the highest loads, whereas detached houses had significantly lower values than the other categories. This is primarily a result of building scale, as larger commercial buildings tend to have greater heating and cooling demands than residential buildings. While the annual load per unit area is effective for evaluating individual building performance, as efforts shift toward city-scale decarbonization, it becomes necessary to analyze annual total loads on a per-building basis.

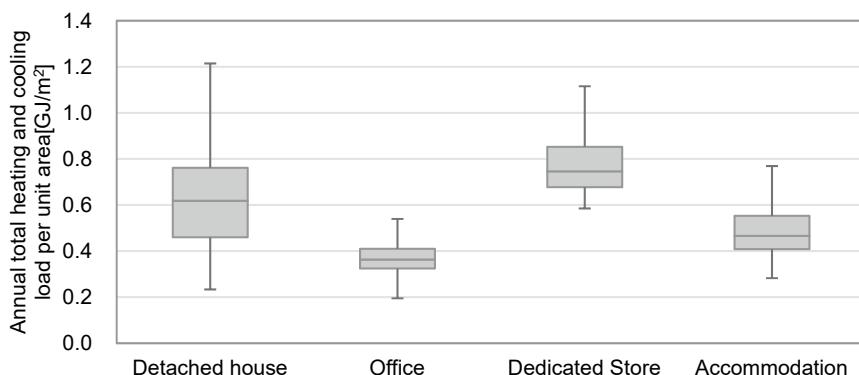


Fig. 4. Annual total heating and cooling load per unit area by building use

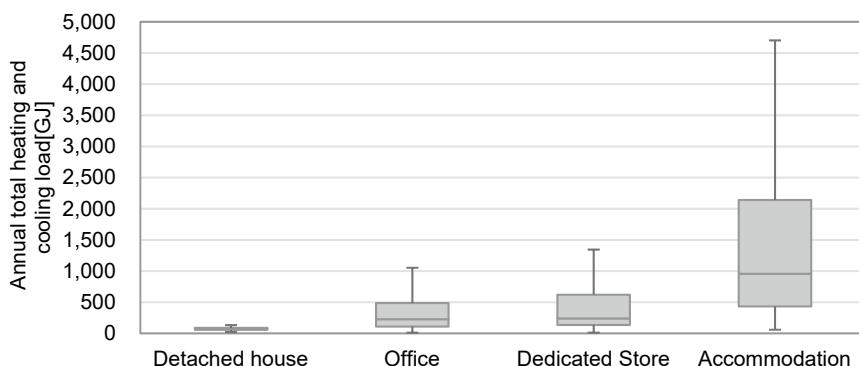


Fig. 5. Annual total heating and cooling load

Spatial Distribution Analysis. Figure 6 shows the spatial distribution of buildings by use within the study area. The northern part of the map, where non-residential buildings are concentrated, corresponds to an entertainment district, while other areas are predominantly composed of detached houses with other building types dispersed throughout. Given that the annual total heating and cooling load for detached houses is about one-tenth that of offices (see Fig. 5), there is potential for load leveling by coordinating electricity use between a single commercial building and multiple detached houses. To further investigate this, Fig. 7 shows the heating and cooling load results for the area highlighted in red in Fig. 6 at 6:00 a.m., 1:00 p.m., and 9:00 p.m. on January 17th. The magnitude of the heating and cooling load is indicated by points plotted above each building. The distribution of building uses is also depicted. In Fig. 7, there are several dozen detached houses surrounding each office building, and it is evident that the heating and cooling loads of detached houses are much lower than those of offices. Some detached houses exhibit higher loads at times other than 1:00 p.m., whereas offices often show the opposite trend. These results suggest that combining one office building with multiple detached houses could contribute to the leveling of heating and cooling demand. Realizing such a scheme could reduce infrastructure costs and help cut peak energy demand, even on a small scale.



Fig. 6. Distribution Map of Buildings by Usage

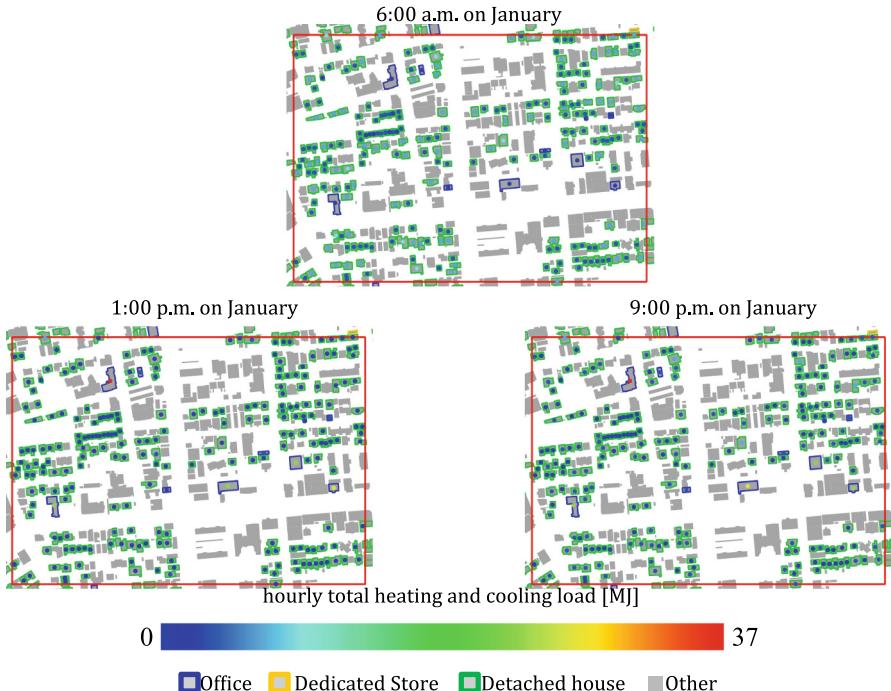


Fig. 7. The heating and cooling load at 6:00 a.m., 1:00 p.m., 9:00 p.m. on January 17th for the area highlighted in red in Fig. 6. (Color figure online)

4 Conclusion

Achieving CN by 2050 has become an internationally recognized goal, and it is essential for each municipality to accomplish decarbonization at the local level. However, small municipalities often face significant challenges due to a lack of specialized knowledge and human resources, which limits their ability to independently advance decarbonization initiatives. Consequently, there is a growing need for systems that can facilitate detailed and spatially explicit regional analyses.

In this study, we aim to develop a system that constructs a comprehensive building database covering a wide area by utilizing geospatial information, and enables its visualization and utilization as a WebGIS. This paper proposes a method for the automatic generation of three-dimensional building models based on GIS data, specialized for air-conditioning load calculations, and demonstrates their application to urban-scale energy analysis.

The automatically generated 3D models are designed to be compatible with EnergyPlus simulations. Building information was sourced from PLATEAU, a GIS platform that has rapidly gained popularity throughout Japan. The target building types for automatic generation included offices, accommodations, dedicated stores, and detached houses. Building geometry was determined using polygon data from PLATEAU, building

heights, and the distances to surrounding structures. For accommodations and dedicated stores, the calculation conditions for load estimation were constructed with reference to WEBPRO. For detached houses, the calculation conditions were established based on assumed family compositions and daily schedules, derived from the 2020 National Census and the 2020 Survey on Time Use and Leisure Activities.

Using the automatically generated 3D models, we calculated and analyzed the heating and cooling loads for each building. The results revealed that offices, accommodations, dedicated stores, and detached houses exhibited differences in heating and cooling loads due to variations in usage patterns, operating hours, and internal heat generation settings. In particular, detached houses showed greater variability in heating and cooling loads compared to other building types, which can be attributed to the diversity in household compositions assumed in the calculations. When focusing on the annual cumulative heating and cooling load per building, it was observed that the loads increased with building size, resulting in relatively lower loads for residential buildings compared to other types. While annual cumulative heating and cooling load per unit area is effective for evaluating the performance of individual buildings, the current phase of decarbonization requires analysis based on the annual cumulative load per building, especially as efforts shift toward city-level implementation. Furthermore, an analysis of the spatial distribution of heating and cooling loads using mapping techniques suggested that combining a single office building with multiple detached houses could help to level out heating and cooling demand. This approach has the potential to reduce infrastructure costs and cut energy demand peaks, even on a small scale.

In summary, the system developed in this study enables the automatic generation of 3D building models for urban-scale energy analysis, providing valuable insights for the formulation of decarbonization strategies involving multiple buildings. Existing research has also utilized GIS-based urban energy simulations and map displays to identify energy-inefficient buildings and clarify issues in specific applications [15–18]. However, these preceding studies often rely on analyses of annual cumulative energy values. In contrast, our methodology performs dynamic simulations with EnergyPlus for every building, calculating and visualizing heating and cooling loads on an hourly basis. This detailed time-series data allows for the precise capture of energy demand peaks and fluctuations, enabling its integration into more realistic urban planning. While the current analysis is specialized for heating and cooling loads, the detailed internal load settings, such as human activity schedules, allow for future expansion to analyze total energy consumption, including hot water supply, lighting, and appliances. This is expected to lead to a granular understanding of city-wide energy dynamics and contribute to the formulation of more advanced energy plans.

5 Limitations and Future Directions

This section summarizes the limitations of the present study and outlines future directions.

1. **Model Simplification and Accuracy:** A limitation of this study is the simplification of floor plans, envelope geometries, and internal loads during the automatic generation of 3D models from GIS data. Although a formal accuracy validation has not yet been

conducted due to a lack of measured field data, the calculated heating and cooling loads appear to be higher than standard reference values. This discrepancy is likely a consequence of the model's simplification, and it is necessary to improve the accuracy of elements where higher fidelity is achievable. Regarding envelope geometry and floor plans, the increasing availability of Level of Detail 3 (LOD3) data is expected to enable more precise representations in the future. Furthermore, research is being conducted on methods to infer internal layouts from building footprints using machine learning [19]. If room uses can be assigned based on GIS footprints by referencing such existing studies, the accuracy of energy load prediction could be significantly improved by applying use-specific load settings.

2. **Applicability to Other Regions:** As the proposed methodology is based on GIS data, it can be readily applied to other regions where such data is available. In recent years, GIS data has become increasingly common, particularly in major cities. Even in unmapped areas, building geometries can be determined using resources like OpenStreetMap or commercial Digital Elevation Models (DEMs). Differences in meteorological conditions can be addressed by utilizing region-specific weather data within EnergyPlus, allowing for simulations that consider local characteristics. However, a significant challenge lies in accurately reflecting regional differences in lifestyles and activity schedules, as creating a perfect replication is often impractical. Therefore, we believe that the application of this energy simulation method can be expanded more efficiently by using machine learning to correct for discrepancies between the simulation results from the auto-generated 3D models and actual measured regional data.
3. **Integration with WebGIS:** The outcomes of this research will be utilized for the energy data in a WebGIS currently under development. The WebGIS consists of a main window and a sub-window, as shown in Fig. 8. In the main window, energy values for different types are displayed color-coded on a 2D map, providing a comprehensive overview of the energy totals and trends for individual or multiple buildings. This is also useful when considering countermeasures for the fluctuations of unstable renewable energy sources like solar power. The sub-window allows for detailed parameter settings, the display of building attribute information, and the graphical representation of energy quantities, including total values for a single building or a wide-area selection. Furthermore, users can export the energy data of selected buildings in CSV format, enabling detailed offline analysis. This WebGIS combines intuitive usability with advanced analytical capabilities, contributing to the formulation of decarbonization plans for entire cities. Specifically, the visualization of time-series energy data enables the identification of instantaneous peak values and energy trends, leading to more realistic urban energy analysis.

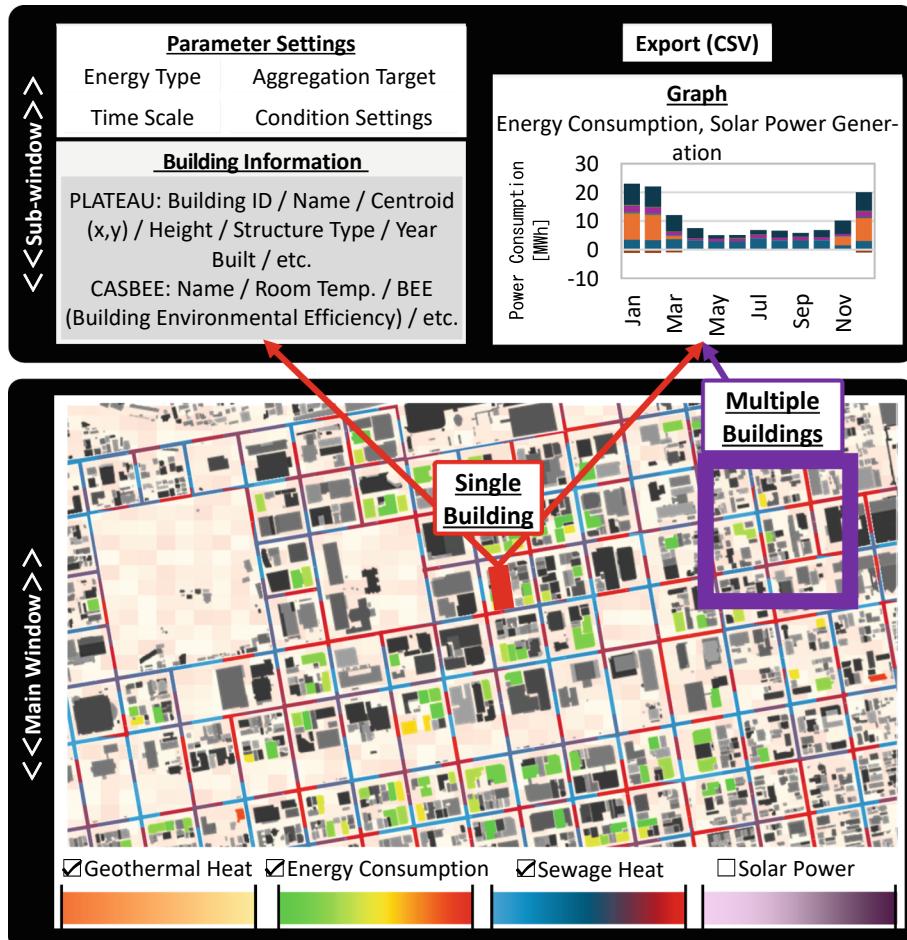


Fig. 8. User Interface of the WebGIS

Acknowledgments. This work was supported by the Grant-in-Aid for Early-Career Scientists (Grant Number 25K21423), the Grant-in-Aid for Scientific Research (A) (Grant Number 25H00763), and the project JPNP24010, commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ministry of Economy, Trade and Industry: Pathways to Japan's Green Transformation (GX) Provisional Translation

2. Ministry of the Environment Government of Japan Ministry of the Environment Government of Japan (in Japanese). https://ondankataisaku.env.go.jp/carbon_neutral/road-to-carbon-neutral/. Accessed 8 May 2024
3. RESAS (in Japanese). <https://resas.go.jp/>. Accessed 10 Apr 2025
4. Japan Energy Database (in Japanese). <https://energy-sustainability.jp/>. Accessed 11 Apr 2024
5. Hisato, O., Mori, T., Shinagawa, K., Nakayama, S., Hosobuchi, H., Mushtaha, E.: Analysis of requirements for achieving carbon neutrality in a cold dense city using GIS. *Environ. Sustain.* **15**, 7063 (2022). <https://doi.org/10.3390/EN15197063>
6. Hisato, O., Mori, T., Shinagawa, K., Nakayama, S., Hosobuchi, H., Mushtaha, E.: Risk assessment of heat stroke during the marathon of the Tokyo 2020 Olympics in Sapporo, Hokkaido. *Sustainability* **15**, 3997 (2023). <https://doi.org/10.3390/SU15053997>
7. PLATEAU. <https://www.mlit.go.jp/plateau/>. Accessed 18 June 2023
8. EnergyPlus. <https://energyplus.net/?adlt=strict&toWww=1&redig=D44D2187730E4FB0B65986DC894AE58A>. Accessed 18 July 2022
9. Daimaru, S., Mori, T., Osawa, H., Izumi, R., Takahashi, R., Suzuki, K.: Development of Energy Prediction Methods for Wide-Area Buildings Using Geographic Information System Collection and Analysis of Window Information of Non-Residential Buildings in Central Sapporo. <https://www.aij.or.jp/paper/detail.html?productId=701949>. Accessed 18 Mar 2025
10. Program Compliant with Energy Conservation Standards for Non-Residential Buildings (in Japanese). <https://building.lowenergy.jp/program>. Accessed 5 Apr 2024
11. Technical Information (Non-Residential Buildings) (in Japanese). <https://www.kenken.go.jp/becc/building.html>. Accessed 11 Apr 2024
12. E-Stat. <https://www.e-stat.go.jp/dbview?sid=0003445284>. Accessed 5 Apr 2025
13. NHK Broadcasting Culture Research Institute. National Time Use Survey (in Japanese). <https://www.nhk.or.jp/bunken/yoron-jikan/>. Accessed 9 Apr 2025
14. Building Research Institute: Technical Information on the Evaluation of Energy Consumption Performance in Accordance with the 2016 Energy Conservation Standards (Residential), Current Edition (in Japanese). <https://www.kenken.go.jp/becc/house.html#2-1>. Accessed 9 Apr 2024
15. Nouvel, R., Zirak, M., Dastageeri, H., Coors, V., Eicker, U.: Urban energy analysis based on 3D city model for national scale applications. In: Proceedings of the Fifth German-Austrian IBPSA Conference, September 2014
16. Amirth, S., Oraiopoulos, A., Korolija, I., Fennell, P.J.: Developing an open access plugin for urban building energy modelling in QGIS. In: Proceedings of BSO Conference 2022: 6th Conference of IBPSA-England; IBPSA-England, vol. 6, 13 December 2022
17. Ali, U., Shamsi, M.H., Hoare, C., O'Donnell, J.: GIS-based residential building energy modeling at district scale. In: Proceedings of the 4th IBPSA-England Conference on Building Simulation and Optimization, Cambridge, UK: 11–12 September 2018. International Building Performance Simulation Association, 12 September 2018
18. Hosseingholizadeh, M., Coors, V., Ostadabbas, H., Friesecke, F.: Development of QGIS plugin for urban energy simulation using 3D city model at the city district level. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.* **10**, 81–90 (2023). <https://doi.org/10.5194/ISPRS-ANNALS-X-1-W1-2023-81-2023>
19. Abouagour, M., Garyfallidis, E.: GenPlan: Automated Floor Plan Generation. Accessed 7 July 2025



Research on 3D Modeling for Thermal Environment Simulation Using Publicly Available Map Data

Kazuma Otani^(✉) , Taro Mori , Hisato Osawa , and Konatsu Suzuki 

Hokkaido University, Kita 13, Nishi 8, Kita-ku, Sapporo, Hokkaido 060-8628, Japan
otani.kazuma.e2@elms.hokudai.ac.jp

Abstract. In Japan, the pursuit of carbon neutrality necessitates a substantial reduction in building energy consumption alongside the proactive integration of renewable energy sources. This study evaluates the effectiveness of a 3D modeling approach in estimating energy demand across multiple buildings.

A 3D building model was developed using GIS data, site surveys, and open datasets to extract critical parameters, such as window-to-wall ratio. Subsequently, indoor environmental conditions, HVAC system specifications, and other relevant factors were incorporated to simulate energy consumption under various conditions. Despite employing a consistent methodology, the accuracy of simulations varied among buildings. Further analysis identified key determinants of accuracy, including the precision of 3D modeling, the functional use of buildings, and the definition of envelope performance parameters.

These findings underscore the importance of high-quality data acquisition and parameter calibration in energy simulations. The insights gained contribute to the advancement of simulation methodologies, ultimately supporting more effective energy management and sustainable urban planning.

Keywords: Energy Simulation · PLATEAU · CASBEE

1 Introduction

Global climate change has emerged as a critical issue demanding coordinated international action. Japan has committed to achieving carbon neutrality (CN) by 2050 and is implementing a range of national policies aimed at restructuring energy consumption and promoting renewable energy adoption [1]. In this effort, urban areas—particularly the building sector—have become a major focus, as they account for approximately 30% of the country's total final energy consumption [2]. Urban environments, characterized by high population density and concentrated economic activity, face increasingly complex energy challenges, including demand concentration and peak load fluctuations. Addressing these challenges requires a detailed understanding of how energy is consumed across different parts of a city [3].

The objective of this study is to develop a methodology for estimating building-level energy consumption and to create a simulation map of energy use in Sapporo, Japan.

Such a map provides a foundation for formulating strategies to enhance energy efficiency and promote renewable energy adoption in cities. For the precision of energy consumption simulations, accurate estimation of building envelope performance parameters is extremely important. Among these parameters, the window-to-wall ratio (WWR) has been identified as a key factor influencing cooling energy use in office buildings, making its accurate assessment indispensable for reliable simulations [4]. In this study, 3D models of individual buildings were created using publicly available map data and field surveys to estimate each building's WWR, and then use the resulting data to simulate energy consumption. The methodology for this study consists of the following steps.

1. Generation of 3D building models using publicly available map data to create the building data needed for the simulation.
2. Estimation of energy consumption based on building characteristics and regional climatic conditions.
3. Comparative analysis with actual consumption data to assess accuracy.

2 3D Modeling Using Publicly Available Map Data

2.1 Modeling Outline

Figure 1 presents an overview of the modeling approach adopted in this study. First, a 3D model was developed in SketchUp (3D modeling software) using the publicly available map data described below as a reference. Next, field photographs were imported into SketchUp and applied as textures to express the building façades. Finally, envelope composition parameters such as WWR were organized and compiled in Excel.

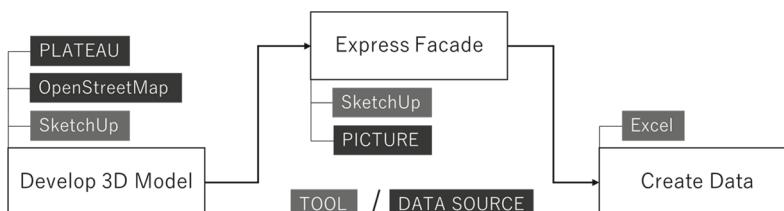


Fig. 1. Flow Chart of Modeling Procedure

2.2 Detailed Modeling Procedure

(1) Develop 3D Model

To develop the 3D models, the following open-source datasets was used to obtain building footprints and heights. (Fig. 2).

- OpenStreetMap (OSM):

OSM is an open-access geographic information dataset that allows users to freely utilize and edit map data [5]. In this study, building footprints were extracted from OSM to define the planar geometry of structures.

- Project PLATEAU:

Project PLATEAU is a web-based 3D urban dataset provided by the Japanese government, constructed from aerial surveys and other sources [6]. In this study, building heights were referenced from PLATEAU.

(2) Express Façade

As shown in Fig. 3, photographs of building façades were taken on site with care to avoid distortions caused by reflections or obstructions. These images were then set as textures onto the corresponding walls of the 3D models created in Step (1) (left side of Fig. 4). The final model is illustrated in right side of right side of Fig. 4.

(3) Create Data

Based on the textured 3D model from Step (2), individual windows were identified and selected (Fig. 5), and the WWR was calculated. This data was organized into a database on a floor-by-floor basis. The wall facing the adjacent building typically has fewer windows. Therefore, it was treated as a windowless facade.

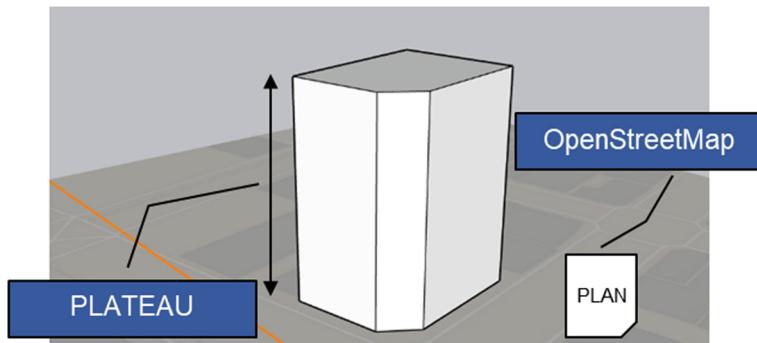


Fig. 2. Develop 3D Model



Fig. 3. Photos Used as Texture

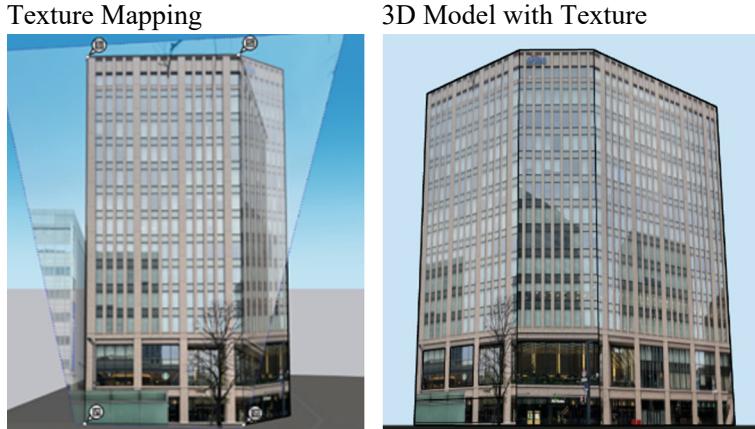


Fig. 4. Texture Mapping and 3D Model with Texture

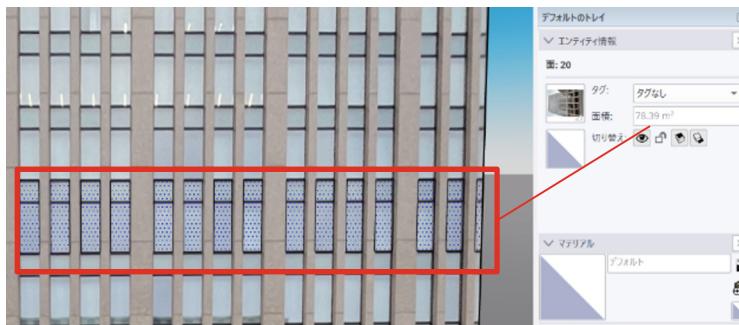


Fig. 5. Selection of the Windows

3 Simulation of Building Energy Consumption

3.1 Method

This study performed energy consumption simulations for 13 buildings in Sapporo, Japan, for which actual energy use data were available. The simulation employed the Simplified Calculation Methods of Cooling and Heating Loads (SHASE-S112-2019) by The Society of Heating, Air-Conditioning and Sanitary Engineers of Japan, a tool implemented in Microsoft Excel. By inputting building-specific parameters (see Table 1), along with indoor and HVAC conditions, the program calculates annual primary energy consumption, CO₂ emissions, peak thermal loads, and total heating/cooling loads.

The input parameters were derived from the 3D modeling process described above. The 13 buildings were divided into two categories based on the type of reference data and simulation process. Simulations were conducted for each building group (A and B) to estimate the annual primary energy consumption. In addition, in this study, the energy consumption for lighting and outlet loads was calculated uniformly based on the conditions shown in Table 2, and these values were added to the simulation results.

- Category A: Buildings using uniform assumptions for envelope data and HVAC.
- Category B: Buildings with case specific parameters sourced from the CASBEE Sapporo database.

CASBEE Sapporo is the City of Sapporo's implementation of the national Comprehensive Assessment System for Built Environment Efficiency (CASBEE), adapted to municipal needs. Under this scheme, building owners of facilities above a specified size must submit environmental performance reports, covering metrics such as energy consumption, envelope specifications, and operational data, to the city. The collected data are published and visualized to support facility managers in optimizing energy use, shifting peak loads, and advancing decarbonization efforts at the local level [7].

Table 1. Parameters of Models

Heading	Unit
Number of Floors	Floor
Floor Type	Top Floor / Intermediate Floor
Orientation of Exterior Wall with Windows	[°]
Zone Type	Interior / Perimeter
Zone Floor Area	[m ²]
Window-to-Wall Ratio	[%]
Room Depth	[m]

Table 2. Energy Consumption from Lighting and Outlets

Heading	Value
Lighting Load Capacity [W/m ²]	10
Plug Load Capacity [W/m ²]	15
Annual Operating Days [day]	245
Daily Operating Hours [h]	13
Primary Energy Conversion Factor [MJ/kWh]	8.64
Primary Energy Consumption from Lighting and Plug Loads [MJ/(m²·year)]	687

3.2 Simulation and Validation for Category a Buildings

6 buildings of the 13 buildings were simulated using standard envelope performance values. For these buildings, exterior wall thermal transmittance was uniformly set based on conditions listed in Table 3. The calculation result was 1.243 [W/m²·K]. Double-glazed windows with thermal transmittance of 2.7 [W/m²·K] were assumed. HVAC systems were assigned based on available design information.

These buildings are all public facilities owned by the City of Sapporo. The calculated energy consumption was compared against open data provided by the municipality (Fig. 6).

Table 3. Conditions for Calculating Thermal Transmittance of Walls

Building Material	Thickness [m]	Thermal Conductivity [W/(m·K)]	Thermal Resistance [m ² ·K/W]
Gypsum Board	0.008	0.22	0.036
Glass Wool	0.025	0.05	0.5
Concrete	0.15	1.6	0.094
Cement Mortar	0.025	1.5	0.017
Tile	0.01	1.3	0.008

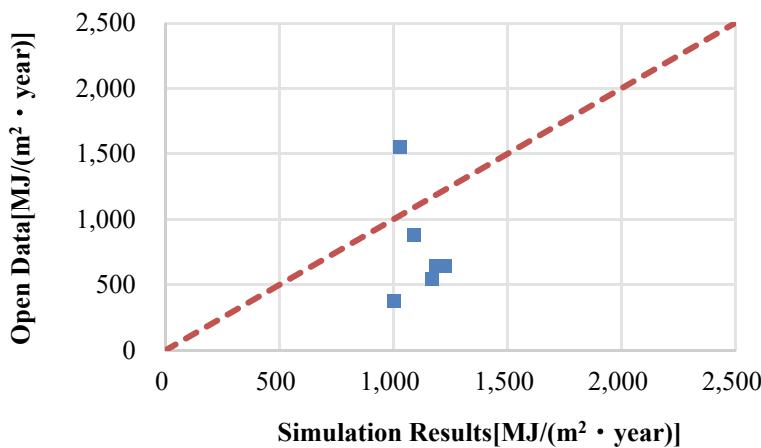


Fig. 6. Comparison of Simulation Results and Open Data (Group A)

In general, large errors were observed between the simulation results and the open data. The Mean Absolute Percentage Error (MAPE) for this simulation was calculated to be 85%. In five of the six cases, the simulation overestimated energy consumption. This suggests that actual operation times for heating/cooling systems, lighting, and outlet loads may be shorter than those assumed in the simulation. Additional factors, such as unconditioned zones or limited night-time operation, may also have contributed to the differences. Conversely, one building—the fire station—showed higher actual energy use than the simulation, likely due to unaccounted factors such as conditioned garages or thermal performance of rolling shutters. The model of the Chuo Ward Office showed a high degree of accuracy. This was largely due to the fact that the estimated envelope performance values and air conditioning systems were close to the actual ones.

3.3 Simulation and Validation for Category B Buildings

The remaining seven buildings were selected from those for which detailed environmental performance data were published by CASBEE Sapporo. Parameters such as the type of air conditioning system, thermal transmittance of walls, and thermal transmittance of windows were referenced separately for each building (see Table 4).

Building Energy Index (BEI) values reported in CASBEE Sapporo were used to compute the design primary energy consumption, which was then compared to simulation results (Fig. 7) [8].

Table 4. Detailed Environmental Performance Data

Building	Air Conditioner Type	Wall thermal transmittance [W/m ² · K]	Window thermal transmittance [W/m ² · K]
Office A	AHU	2	4
Office B	FCU	0.7	3
Office C	AHU	0.5	2.2
Office D	AHU	2	4
Hotel A	FCU	1	3
Hotel B	AHU	2	4
Hotel C	FCU	1	3

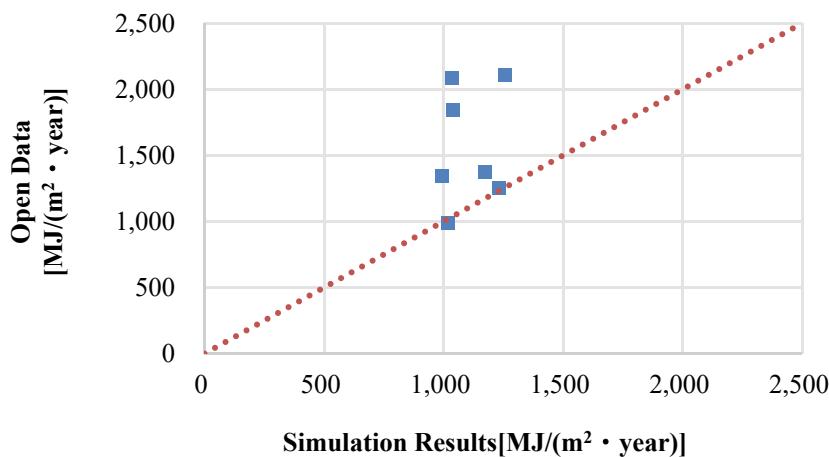


Fig. 7. Comparison of Simulation Results and Open Data (Group B)

Among the seven buildings, four showed good agreement between simulated and actual energy consumption. These four were all office buildings. The other three, all hotels, exhibited significant underestimation by the simulation. The overall MAPE for this simulation was 26%, while the MAPE for the office buildings alone was 11%.

This is likely due to the longer operation hours of hotel HVAC and lighting systems, which were not fully captured by the office-based simulation assumptions. Although the energy consumption associated with electrical outlets and lighting was uniformly set, building operating hours are expected to vary considerably depending on their intended use. Therefore, the energy consumption should be adjusted based on the building type.

The higher accuracy in simulating office buildings is attributed to the detailed parameter inputs enabled by CASBEE Sapporo data. These results indicate that by adapting use type settings and envelope scores more precisely, the reliability of energy consumption simulations can be significantly improved.

4 Discussion

In this study, energy simulations were successfully performed only for buildings whose envelope performance and HVAC systems referenced CASBEE Sapporo data. The results underscore the importance of accurately specifying envelope performance and HVAC conditions. The four buildings with the highest simulation accuracy were all office buildings, indicating that building use has a significant impact on simulation precision. These findings demonstrate that the proposed simulation method can produce accurate energy consumption estimates when building use and envelope specifications are clearly defined. Currently, many buildings in Sapporo are not registered under CASBEE Sapporo. CASBEE Sapporo is mainly implemented for large-scale buildings. For buildings without detailed data on envelope performance and air conditioning systems, simulation methods need to be considered. Furthermore, a simulation method for energy consumption of buildings used for purposes other than offices should also be considered.

5 Conclusion and Future Perspectives

This study examined the effectiveness of estimating building energy consumption using 3D building models developed from publicly available map data. Geometric information such as WWR was extracted by integrating photo-based façade textures, and building-level energy demand was simulated accordingly. The results showed that simulations were highly accurate for office buildings for which detailed envelope and air conditioning data were available through CASBEE Sapporo, whereas generalized assumptions often led to overestimations. A major limitation of this study was the small number of buildings analyzed. In particular, buildings with complex wall geometries were excluded from the simulation because the façade could not be properly reconstructed in the 3D model. The limited dataset made it difficult to identify whether discrepancies between the simulation and actual energy consumption were due to issues in the 3D modeling, inaccurate envelope inputs, or incorrect HVAC assumptions.

Future research should focus on the following challenges and directions:

- Introduce modeling techniques that can capture complex façade geometries, such as those using image recognition and LiDAR data, to improve the accuracy of 3D models and to increase the number of buildings to simulate.
- Develop application methods for buildings with different operational characteristics, such as hotels and residential buildings, by establishing use-specific parameter sets.

- Examine efficient methods for increasing the number of buildings to be simulated, such as recording video with a camera mounted on a vehicle and extracting still images from the footage.

By implementing these improvements, it will be possible to achieve high-precision

energy assessments at the urban scale, thereby contributing to the design of Zero Energy.

Buildings and the formulation of effective energy policies.

Acknowledgments. This work was supported by the Grant-in-Aid for Early-Career Scientists (Grant Number 25K21423), the Grant-in-Aid for Scientific Research (A) (Grant Number 25H00763), and the project JPNP24010, commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ministry of Land, Infrastructure, Transport and Tourism, Initiatives toward Realizing Carbon Neutrality in the Land, Infrastructure, Transport and Tourism Sector. https://www.mlit.go.jp/policy/shingikai/carbon_neutral.html. Accessed 18 May 2025
2. Agency for Natural Resources and Energy, Energy White Paper 2022, Part 2, Chapter 1, Section 2. <https://www.enecho.meti.go.jp/about/whitepaper/2022html/2-1-2.html>. Accessed 18 May 2025
3. Reinhart, C.F., Davila, C.C.: Urban building energy modeling – a review of a nascent field. *Build. Environ.* **97**, 196–202 (2016)
4. Troup, L., Phillips, R., Eckelman, M.J., Fannon, D.: Effect of window-to-wall ratio on measured energy consumption in US office buildings. *Energy Build.* **203**, 109434 (2019)
5. OpenStreetMap Japan. <https://openstreetmap.jp/>. Accessed 16 May 2025
6. Ministry of Land, Infrastructure, Transport and Tourism, PLATEAU. <https://www.mlit.go.jp/plateau/about/>. Accessed 16 May 2025
7. City of Sapporo, CASBEE Sapporo. <https://www.city.sapporo.jp/kankyo/casbee/>. Accessed 16 May 2025
8. National Institute for Land and Infrastructure Management, Technical Information on the Evaluation of Energy Consumption Performance Based on the 2016 Energy Conservation Standards. https://www.kenken.go.jp/becc/documents/building/Definitions/modelBuilding_RferenceValues_20200918.pdf. Accessed 16 May 2025



Rooftop Irregularity Segmentation in Aerial Imagery Using Deep Learning

Konatsu Suzuki^(✉) , Taro Mori , Hisato Osawa , and Kazuma Otani 

Hokkaido University, Kita 13, Nishi 8, Kita-ku, Sapporo, Hokkaido 060-8628, Japan
suzuki.konatsu.f1@elms.hokudai.ac.jp

Abstract. To achieve carbon neutrality (CN) by 2050, efforts have primarily focused on CN at the level of individual buildings. However, it is essential to consider how CN can be achieved for entire cities, including these buildings. This study aims to estimate the available space for solar panel installation on buildings in urban areas to enhance the utilization of photovoltaic systems. In this study, we propose a semantic segmentation method for estimating building shape (protruding parts and depressed parts relative to the reference plane), which significantly influences the available space for solar panel installation. We conducted semantic segmentation under five different conditions, varying the training and testing datasets to analyze the impact of roads in aerial imagery. Furthermore, to improve the extraction accuracy of protruding parts, this study proposes a weighted overlay method, which overlays five different inference images using various weight combinations. For protruding part extraction using the semantic segmentation method, an F-Score of 30.5% was achieved when roads were fully retained during training and inference, whereas an F-Score of 41.4% was obtained when roads were entirely removed during training and inference. However, the extraction of depressed parts was not successful. Additionally, the weighted overlay method improved the F-Score, achieving a maximum value of 72.5%.

Keywords: Rooftop · Deep Learning · Semantic Segmentation

1 Introduction

With the international goal of achieving carbon neutrality (CN) by 2050, Japan has outlined a policy to realize nationwide decarbonization through the formulation of a decarbonization roadmap, enabling municipalities across the country to progressively implement carbon reduction measures. Traditionally, the focus has been on achieving CN at the level of individual buildings, such as Net Zero Energy Buildings (ZEBs). However, moving forward, it is essential to consider how CN can be achieved for entire cities, including these buildings. Achieving city-scale carbon neutrality necessitates the development of sophisticated analytical tools designed to optimize energy supply and demand. These tools are pivotal in facilitating the efficient integration of renewable energy sources into the urban energy infrastructure.

In this context, research on the effective use of renewable energy in urban areas has been actively pursued. For example, Ali et al. [1] proposed a method to enhance grid-connected PV system by using an Adaptive Fuzzy Logic Controller to stabilize output under fluctuating solar conditions. Nevertheless, relatively few studies have investigated the potential for introducing renewable energy, particularly solar power, in urban environments, where space constraints and building diversity pose unique challenges. This study aims to estimate the available space for PV installation on rooftops in urban areas to enhance the utilization of photovoltaic systems.

Detecting building outlines from high-resolution satellite or aerial images has been an important and challenging research topic. In particular, the combination of deep learning methods and GIS data (e.g., satellite and aerial images, map data, etc.) has been widely used for building outline detection [2–5]. For example, Li et al. [2] proposed a U-Net-based semantic segmentation method using public GIS datasets and satellite imagery. However, these studies do not account for building shapes. To better characterize building shapes, many studies have adopted light detection and ranging (LiDAR) data [6–8]. However, public LiDAR datasets are limited, and LiDAR data treat rooftop equipment the same as rooftop irregularities. In contrast, aerial imagery is widely available across Japan and allow for a clear distinction between rooftop equipment and irregularities.

This report presents research on building shape estimation, which significantly influences the available space for solar panel installation. The study was conducted using aerial imagery and deep learning techniques, including semantic segmentation, which enables precise pixel-level classification of objects such as buildings and surrounding infrastructure. This study creates and evaluates deep learning semantic segmentation models using approximately 1,500 buildings in central Sapporo, as shown in Fig. 1. Throughout our research, we conducted training and inference under five different conditions (see Table 1) using the same hyperparameters, varying the training and test settings with aerial imagery at a resolution of 10 cm to assess the influence of roads in aerial images. We used AP (Aerial Photograph), AP_B (Images with only buildings extracted from AP), GT (Ground Truth), and GT_B (Images with only buildings extracted from GT) for training and testing. Furthermore, to improve extraction accuracy, this study proposes a weighted overlay method, which overlays five different inference images using various weight combinations.

In this report, the term “protruding part” refers to a section of the rooftop plane that is elevated above the reference plane, which is the largest continuous surface area of each building. Similarly, the term “depressed part” refers to a section of the rooftop plane that is lower than the reference plane.

The rest of the report is organized as follows. Section 2 introduces the methodology, including hyperparameters, data preparation and augmentation, and the weighted overlay method. Section 3 presents the results of protruding and depressed parts extraction, as well as overlayed results. Section 4 discusses and analyzes the extraction results obtained from each semantic segmentation model and the weighted overlay method. Finally, Sect. 5 summarizes the conclusions of this research.



Fig. 1. Study area (red dotted line), covering Sapporo’s business district with 1472 mainly non-residential buildings. Left: building outlines. Right: aerial image. (Color figure online)

Table 1. Training and Inference Conditions

Condition Name	Learning Condition		Test Condition
	Input Image	Ground Truth	
A	AP	GT	AP
B	AP _B	GT	AP _B
C	AP _B	GT	AP
D	AP _B	GT _B	AP _B
E	AP _B	GT _B	AP

2 Method

2.1 Hyperparameters

Table 2 presents the hyperparameters used in this study. In this study, we employed U-Net [9] as the architecture for the deep learning model. U-Net is a type of convolutional neural network (CNN) designed for pixel-wise image segmentation. The U-Net encoder extracts features through convolution, while the decoder restores them to their original image size via upsampling. Additionally, skip connections help preserve fine details and effectively transfer them to the decoder. Due to its ability to achieve high segmentation accuracy even with limited training data, U-Net is well-suited for this study. To reduce computational load, the input images (512×512) were resized to 256×256 . Initially, training was conducted for 6 epochs, but in some cases, convergence was insufficient. Therefore, the training was extended to 12 epochs where necessary. Binary Cross Entropy was selected as the loss function since this segmentation task involves distinguishing between background and target areas. Batch size determines the number of samples processed per update step. A smaller batch size introduces greater variability in gradient updates, which can enhance generalization and prevent overfitting. Consequently, a batch size of 2 was chosen. To ensure sufficient training despite the small batch size, a higher number of update steps per epoch were required. Thus, the number of steps per epoch

was set to 2,000 to allow for adequate parameter updates and ensure effective learning. In this study, the Adam optimizer was used to optimize the deep learning model. This choice was made because its adaptive learning rate adjustment facilitates fast and stable convergence. The initial learning rate was set to 0.0001.

Table 2. Hyperparameters

Hyperparameters	Value
Model Architecture	U-Net
Input Image Size	512 × 512 (Resized to 256 × 256)
Epochs	6 (12)
Loss Function	Binary Cross Entropy
Batch Size	2
Steps per Epoch	2,000
Learning Rate	0.0001
Optimizer	Adam

2.2 Preparation of Input and Test Images

For the input images (hereafter referred to as “Input Image”) and the test images (hereafter referred to as “Test Image”), we utilized 10 cm resolution aerial imagery [10] (AP) taken in August 2022, which were commercially provided by Kokusai Kogyo Co., Ltd.

Additionally, as shown in Fig. 2, we used processed images (AP_B) where only buildings were extracted from aerial imagery. The extraction of building areas was performed using a pre-trained model (Buildings Segmentation) [11] implemented as a QGIS plugin developed by Przemysław Aszkowski and Bartosz Ptak. This model employs Xunet for segmentation between buildings and the background. It was trained on the RampDataset [12], achieving an F-Score of 81% during validation. Both AP and AP_B images have dimensions of 13,312 × 14,848 pixels.

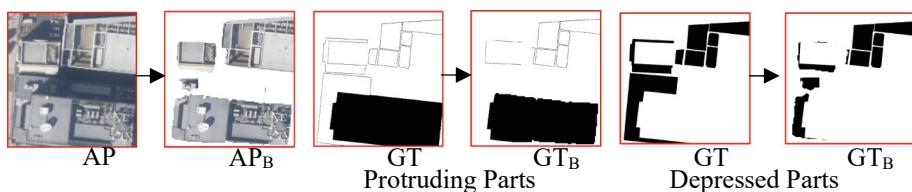


Fig. 2. Extraction of building area from the aerial imagery and the ground truth image.

2.3 Creation of Ground Truth

Annotations were manually performed using QGIS. The annotation process involved manually identifying and labeling protruding and depressed parts. The resulting images

consist of GT, where protruding or depressed parts are black, and the background is white. Additionally, we used another type of image, GT_B , which was extracted from GT and contains only the overlapping areas of protruding or depressed parts and AP_B .

2.4 Training and Inference Conditions

To facilitate resizing to 256×256 , the images were initially divided into 512×512 segments, allowing downscaling by a factor of two while minimizing interpolation artifacts. Additionally, rather than directly splitting the images into 256×256 sections, the 512×512 segmentation approach was chosen to preserve the structural integrity of individual buildings, ensuring that they remain as intact as possible within each processed image. As shown in Fig. 3, both the Input Image and Ground Truth were divided into 512×512 pixels. Of the 754 image pairs, 749 were augmented through rotation, translation, shear transformation, zoom, horizontal flipping, and filling, generating 3,995 training pairs. For inference, five unseen images were used as Test Images. The output consists of an image representing the probability of protruding parts in grayscale (Output Image) and an image where pixels above a 50% threshold are classified as protruding or depressed parts (Binary Image). Evaluation metrics were calculated for all conditions using GT and Binary Image.

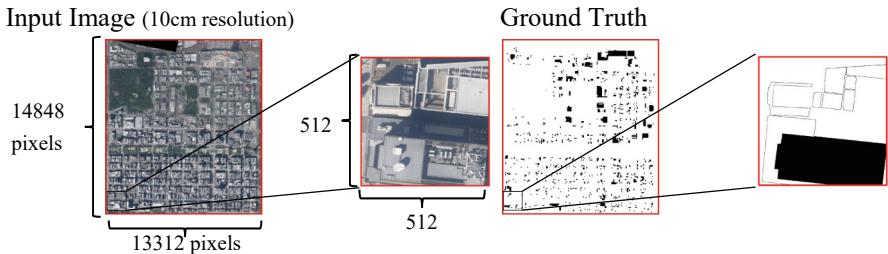


Fig. 3. Division of Input Image and Ground Truth

2.5 Weighted Overlay Method

A weighted overlay was applied to five inference images of protruding parts. Using the Output Images generated under five conditions, weights were varied from 0.0 to 1.0 in steps of 0.1, while keeping the total weight at 1.0, resulting in 1001 different weighting combinations. Each inference image was then overlaid using these combinations, with a 50% threshold applied to generate the corresponding images (Blended Image). Evaluation metrics were calculated for all 1001 conditions.

2.6 Evaluation Metrics and Methods

The evaluation metrics used in this study include precision, recall, F-Score, and Intersection over Union (IoU). Precision represents the proportion of data predicted as positive by the model that is actually positive. On the other hand, recall indicates the proportion of actual positive data that the model correctly predicts as positive. The F-Score is the

harmonic mean of precision and recall. IoU measures the overlap between the ground truth and the predicted area, calculated as the ratio of their intersection to their union. In all conditions, evaluation metrics were calculated based on GT.

3 Results

3.1 Inference Results for Each Condition

Protruding Parts. Figure 4 presents the inference results for protruding parts, and Table 3 lists the evaluation metrics. Across all conditions, both the Output Image and Binary Image showed differed notably from the Ground Truth. The maximum F-Score reached approximately 40%, and the highest IoU was around 25%, indicating insufficient learning. Conditions A, B, and D contained some correctly extracted areas, as suggested by their relatively high precision values. However, their very low recall values indicate that many true protruding areas were missed. Conversely, conditions C and E extracted excessively large areas compared to their GT images. Their high recall values further suggest a high rate of false positives. The F-Score and IoU values in conditions B and C were significantly lower than those in condition A. In contrast, condition D achieved higher values than condition A, while condition E showed comparable results. Notably, in Images 1 and 2, protruding parts located in the shadow of other buildings were not extracted. Additionally, in Image 5, the equipment beneath the structural frame supporting the signboard was erroneously recognized as a protruding part.

Depressed Parts. Figure 5 presents the inference results for depressed parts. In condition A, some areas were inferred; however, the inferred areas differed significantly from the true areas. Additionally, the evaluation metrics were low, indicating insufficient learning. In conditions B and C, areas corresponding to building shadows in the AP_B were inferred as depressed parts with high probability. As a result, the extracted areas deviated significantly from the actual depressed parts. Despite this, their evaluation metrics were slightly better than those of condition A. In conditions D and E, all pixels were inferred as 100% depressed parts, suggesting that the model failed to learn properly.

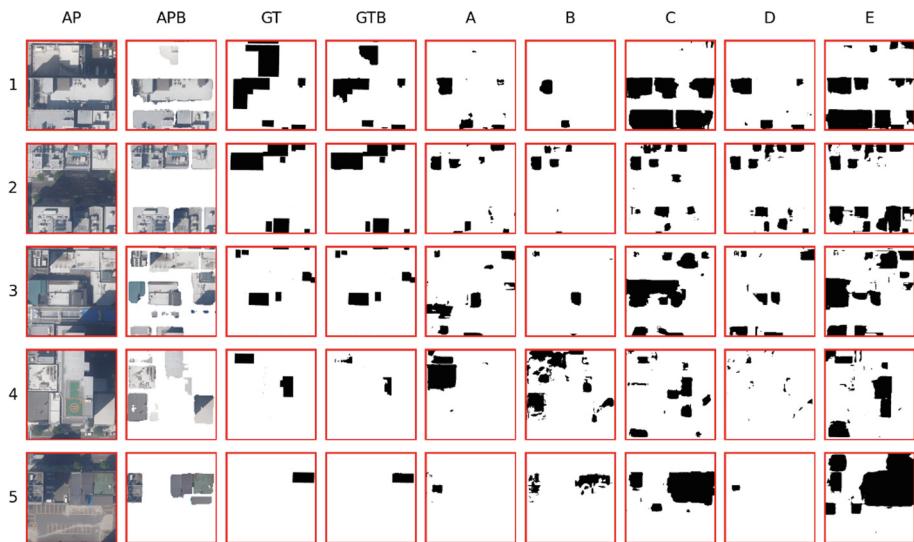


Fig. 4. From left to right are the AP, AP_B, GT, and GT_B used for inference and Binary Images for conditions A through E when the protruding parts were inferred.

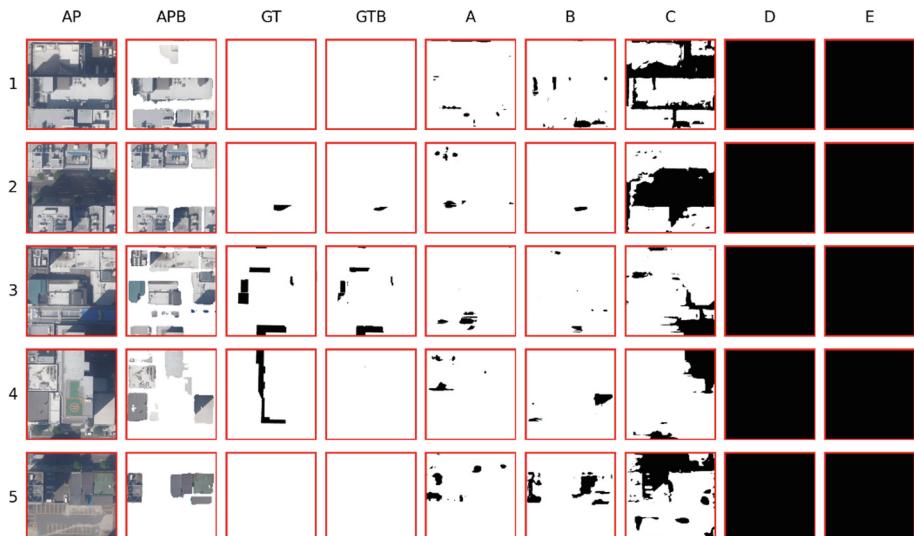


Fig. 5. From left to right are the AP, AP_B, GT, and GT_B used for inference and Binary Images for conditions A through E when the depressed parts were inferred.

Table 3. Training and Inference Conditions

	Conditions	Precision	Recall	F-Score	IoU
Protruding Part	A	38.8%	25.2%	30.5%	18.0%
	B	34.7%	19.9%	25.3%	14.5%
	C	21.9%	41.2%	28.6%	16.7%
	D	60.7%	31.4%	41.4%	26.1%
	E	22.0%	52.4%	31.0%	18.3%
Depressed Part	A	1.9%	1.6%	1.8%	0.9%
	B	6.2%	6.3%	6.2%	3.2%
	C	1.3%	12.0%	2.3%	1.2%
	D	***	***	***	***
	E	***	***	***	***

3.2 Results of Overlaying Inference Images

Table 4 lists the optimal weighting combinations for each test image that resulted in the highest F-Score and IoU, along with the evaluation metrics for the Blended Images. Additionally, Fig. 6 presents the output images generated for each test image and condition, as well as the Blended Images.

In Images 1, 4, and 5, the Blended Images more closely resembled the Ground Truths than individual conditions, with improved F-Score and IoU values. Conversely, in Images 2 and 3, the optimal weighting assigned condition D a weight of 1.0 while all other conditions had weights of 0.0, meaning overlaying had no significant effect on these images. Moreover, in Images 1 and 2, for buildings where protruding parts were stacked on top of one another, the lower protruding parts were not extracted—even in the Blended Image. This is likely because those areas had low probability values in individual condition outputs, preventing them from being captured through overlaying. For Image 5, false extractions persisted in a portion of the building on the left side of the image, even in the Blended Image. This is further indicated by the high recall value.

Regarding weighting combinations, the optimal weights varied significantly across test images, with no consistent trend observed.

Table 4. Training and Inference Conditions

Test Image	Weight Combination					Evaluation Metrics			
	A	B	C	D	E	Precision	Recall	F-Score	IoU
1	0.7	0.0	0.0	0.1	0.2	66.4%	36.8%	47.3%	31.0%
2	0.0	0.0	0.0	1.0	0.0	65.9%	51.3%	57.7%	40.5%
3	0.0	0.0	0.0	1.0	0.0	30.8%	41.2%	35.3%	21.4%

(continued)

Table 4. (continued)

Test Image	Weight Combination					Evaluation Metrics			
	A	B	C	D	E	Precision	Recall	F-Score	IoU
4	0.1	0.6	0.0	0.0	0.3	77.4%	68.1%	72.5%	56.8%
5	0.3	0.3	0.2	0.1	0.1	37.9%	79.3%	51.3%	34.5%

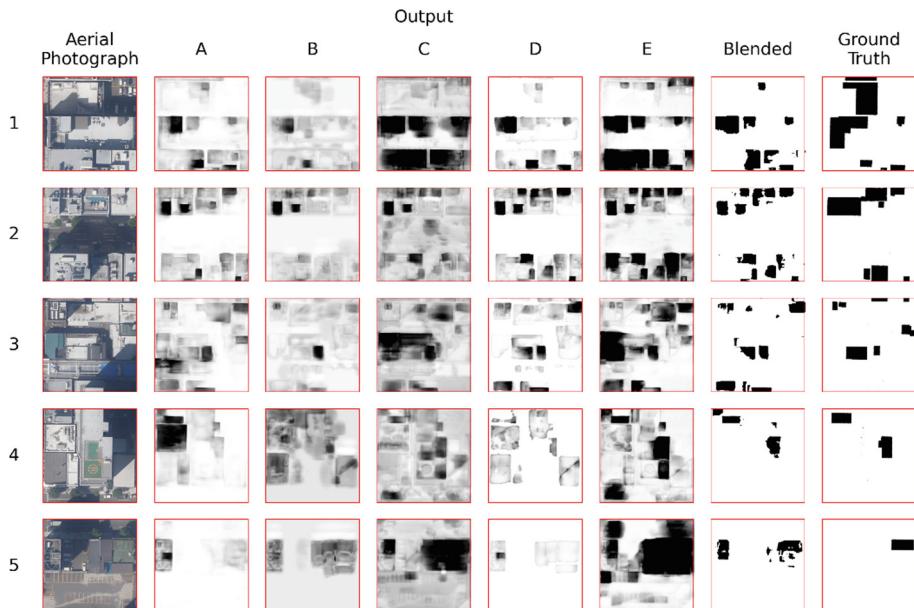


Fig. 6. From left to right are the aerial imagery, Output Images (images that represent the probability of protruding parts in grayscale) for each condition, the Blended Image with the best F-Score and IoU values, and the Ground Truth, which is the correct image.

4 Discussion

4.1 Inference of Protruding Parts and Depressed Parts

Protruding Parts. Conditions C and E exhibited high Recall, while B and D showed higher Precision. In C and E, the models frequently extracted entire buildings rather than just protrusions, likely because they learned a heuristic such as “areas adjacent to shadows indicate protrusions.” Since these models used AP images for inference without prior exposure to building shadows, they misinterpreted building shadows as protrusions, increasing false positives.

In contrast, B and D used AP_B images for both training and testing, which lack building shadows. Thus, their models applied the shadow-adjacency rule more accurately, leading to fewer false positives and better Precision.

Condition A was also tested on AP images, but its model was trained on AP, allowing it to distinguish building shadows from those of protrusions. As a result, it avoided the over-extraction seen in C and E.

Conditions B and C, as well as D and E, shared training parameters but differed in test data. In B and C, Ground Truth (AP_B) included inconsistencies due to imperfect building extraction, causing the model to learn some background areas as protrusions. D and E had more consistent data, and D, which maintained full consistency between training and testing (both using AP_B), achieved the best F-score and IoU among B to E.

Comparing A and D, A required the model to distinguish between different shadow types, whereas D allowed all shadows to be treated as protrusions. This simpler task in D likely contributed to its superior performance.

Although AP_B was chosen to avoid road-related noise, results show that the type of shadow learned by the model had a greater influence. In C and E, building shadows present in Test Images (AP) led to over-extraction, despite the training aim to remove such noise. Since building shadows often fall on roads, their absence in training (as in AP_B) critically impacted inference accuracy.

Depressed Parts. The inference results for depressed parts were worse than for protruding parts. One possible reason is that depressed parts constitute an extremely small proportion of the image compared to the background, making effective model learning difficult. Additionally, in conditions D and E, the already limited depressed parts may have been further diminished due to inaccuracies in building extraction, potentially leading to complete learning failure. Another contributing factor could be improper data augmentation. Given the small size of depressed areas, data augmentation should have been carefully tailored; however, some of the augmentation techniques used may have produced inappropriate training data.

4.2 Overlaying Inference Results

Blended Image 2 achieved relatively high evaluation metrics without overlaying. In this case, overlaying was deemed unnecessary because one of the conditions already had high evaluation metrics. Additionally, in Blended Images 1 and 4, condition A and B respectively had significant weight, while condition A, D, and E complemented the main condition. This indicates that overlaying helped highlight extractable areas within each condition while mitigating false extractions. Therefore, overlaying was successful in these images.

Conversely, overlaying was ineffective for Images 3 and 5. In Image 3, where condition D had a weight of 1.0 and the evaluation metrics were relatively low, overlaying was ineffective, indicating that the protruding parts in the image were difficult to extract. In Image 5, although the extracted areas varied slightly across conditions, overlaying expanded the extraction range, ultimately amplifying false extractions.

5 Conclusion

This study analyzed the protruding and depressed parts of buildings in Sapporo City using aerial imagery and semantic segmentation. Additionally, it examined whether the presence of roads in the photographs influences training and whether the overlaying method enhances the extraction of protruding parts. The key findings are as follows:

1. Despite the low evaluation metrics, some protruding parts were correctly extracted, demonstrating the potential of semantic segmentation for protrusion detection.
2. Extracting depressed parts proved challenging under the current conditions.
3. The removal of roads through the use of AP_B images, while intended to reduce noise, indirectly affected extraction accuracy by eliminating building shadows that often fall on road surfaces.
4. In certain cases, overlaying methods improved the accuracy of protruding part extraction. However, in other instances, overlaying amplified false extractions, resulting in lower evaluation metrics.

Enhancing the performance and applicability of our building protrusion and depression extraction model will involve focusing on three key areas in future work. Firstly, we will optimize model performance by exploring a wider range of hyperparameter configurations, investigating alternative loss functions (e.g., focal loss, Dice loss, IoU loss), and evaluating different model architectures such as SegNet and DeepLabV3+. We also plan to assess the impact of batch size and the number of training steps on accuracy and generalization. Secondly, acknowledging the significant impact of building shadows on model inference, we aim to utilize alternative types of input imagery and refine data augmentation strategies. These efforts are intended to reduce the model's dependency on shadow features and promote the learning of intrinsic protrusion characteristics. Finally, to address the limitations of manual annotation, we intend to explore semi-supervised learning approaches. This will help improve data scalability and enable model applications beyond the Sapporo region.

In addition to the extraction of protruding and depressed parts, we plan to evaluate the effectiveness of the weighted overlay method under various conditions to investigate its capability in distinguishing building forms. This will contribute to more comprehensive morphological classification of urban structures based on aerial imagery.

Acknowledgments. This work was supported by the Grant-in-Aid for Early-Career Scientists (Grant Number 25K21423), the Grant-in-Aid for Scientific Research (A) (Grant Number 25H00763), and the project JPNP24010, commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ali, M., Ahmad, M., Koondhar, M.A., Akram, M.S., Verma, A., Khan, B.: Maximum power point tracking for grid-connected photovoltaic system using adaptive fuzzy logic controller. *Comput. Electr. Eng.* **110**, 108879 (2023). <https://doi.org/10.1016/j.compeleceng.2023.108879>
2. Li, W., He, C., Fang, J., Zheng, J., Fu, H., Yu, L.: Semantic segmentation-based building footprint extraction using very high-resolution satellite images and multi-source GIS data. *Remote Sens.* **11**(4), 403 (2019). <https://doi.org/10.3390/rs11040403>
3. Ekiz, S., Acar, U.: Improving building extraction from high-resolution aerial images: error correction and performance enhancement using deep learning on the Inria dataset. *Sci. Prog.* **108**(1) (2025). <https://doi.org/10.1177/00368504251318202>
4. Shrestha, S., Vanneschi, L.: Improved fully convolutional network with conditional random fields for building extraction. *Remote Sens.* **10**(7), 1135 (2018). <https://doi.org/10.3390/rs10071135>
5. Yi, Y., Zhang, Z., Zhang, W., Zhang, C., Li, W., Zhao, T.: Semantic segmentation of urban buildings from VHR remote sensing imagery using a deep convolutional neural network. *Remote Sens.* **11**(15), 1774 (2019). <https://doi.org/10.3390/rs11151774>
6. Castagno, J., Atkins, E.: Roof Shape Classification from LiDAR and satellite image data fusion using supervised learning. *Sensors* **18**(11), 3960 (2018). <https://doi.org/10.3390/s18113960>
7. Bizjak, M., Žalik, B., Štumberger, G., Lukač, N.: Large-scale estimation of buildings' thermal load using LiDAR data. *Energy Build.* **231**, 110626 (2021). <https://doi.org/10.1016/j.enbuild.2020.110626>
8. Zhang, C., Cui, Y., Zhu, Z., Jiang, S., Jiang, W.: Building height extraction from GF-7 satellite images based on roof contour constrained stereo matching. *Remote Sens.* **14**(7), 1566 (2022). <https://doi.org/10.3390/rs14071566>
9. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W., Frangi, A. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
10. Kokusai Kogyo Co., Ltd. <https://www.kkc.co.jp/>. Accessed 13 May 2025
11. Aszkowski, P., Ptak, B., Kraft, M., Pieczyński, D., Drapikowski, P.: Deepness: deep neural remote sensing plugin for QGIS. *SoftwareX* **23**, 101495 (2023). <https://doi.org/10.1016/j.softx.2023.101495>
12. DevGlobal: ramp Building Footprint Training Dataset - Paris, France, Version 1.0. Radianit MLHub (2022). <https://doi.org/10.34911/rdnt.t86thc>



Autonomous Air Quality Monitoring System with Photovoltaic Energy Harvesting: A Sustainable Approach for Public Policies

Emerson Santana^(✉)  and Eduardo Liberado 

Institute of Science and Technology, São Paulo State University (UNESP), Sorocaba, Brazil
emerson.santana@unesp.br

Abstract. Air quality measurement is essential for decision-making and the establishment of public policies that promote health, well-being, and sustainability. This work proposes an autonomous air quality monitoring system powered by photovoltaic energy, aiming to contribute to the expansion of sustainable environmental monitoring networks. The developed system incorporates energy harvesting elements, using solar panels and power management integrated circuits (PMICs) to ensure energy autonomy and portability. A prototype was built to monitor carbon monoxide (CO) and transmit measured data through Wi-Fi. This prototype kept continuous operation for four days having only solar power and batteries as energy sources. These results highlight the feasibility of using photovoltaic energy in low-power systems, especially in remote or hard-to-access locations, reinforcing the importance of sustainable technologies for air quality monitoring. A methodology to design air quality measurement projects was also developed, enabling the replication of the system in different contexts.

Keywords: Air quality · energy harvesting · photovoltaic energy · autonomous monitoring · PMIC

1 Introduction

Air quality measurement is a crucial action for decision-making and the establishment of public policies aimed at improving health, well-being, and sustainability. In 2022, the United Nations (UN) [1] adopted Resolution 73/300, recognizing the right of all people to a clean and healthy environment. This initiative aligns with the Sustainable Development Goals (SDGs), which seek to eradicate poverty, protect the environment, and ensure global prosperity. This work is aligned with SDG 7 (Affordable and Clean Energy), 11 (Sustainable Cities and Communities), and 3 (Good Health and Well-being), focusing on the development of autonomous and sustainable air quality monitoring systems.

Air pollution is one of the main environmental challenges, causing significant impacts on human health and ecosystems. Gases such as carbon dioxide (CO₂), [2] carbon monoxide (CO), particulate matter (PM2.5 and PM10), and ozone (O₃) are common pollutants that require constant monitoring [3].

However, the implementation of large-scale measurement networks faces challenges, such as the need for energy and communication infrastructure in remote areas. In this context, energy harvesting, especially photovoltaic, emerges as a viable solution to power low-consumption monitoring systems.

Therefore, this work proposes an autonomous air quality monitoring system using photovoltaic energy and power management integrated circuits (PMICs). The developed prototype demonstrates the feasibility of the approach, with promising results in terms of energy efficiency and portability. Additionally, a sizing methodology [2] for air quality measurement projects is presented, facilitating the replication of the system in different contexts [3].

In this initial prototype, a single air quality parameter (CO) and a specific transmission technology (Wi-Fi) were selected to validate the energy model and system design under low-power constraints. This choice was made to allow future scalability, as monitoring additional parameters or adopting alternative communication technologies would primarily require adjustments to energy provisioning, without significant changes to the system's architecture.

2 Energy Harvesting and Air Quality Monitoring Technologies and Applications

This section presents a brief review of energy harvesting applications in air quality monitoring systems reported in the literature. Table 1 summarizes technical information for each application. Additionally, the two main technologies applied in this work (PMIC and air quality sensors) are briefly described.

P. Das et al. [4] introduced an air quality monitoring system composed of various sensors to monitor NO₂, O₃, CO, SO₂, particulate matter (PM1, PM10, PM2.5), as well as temperature and humidity. The system is also equipped with GPS and transmits data via Wi-Fi and NB-IoT. It features a backup battery and is powered by a photovoltaic (PV) module. Energy management is performed by a BQ25506 PMIC, which does not include a boost converter and requires an external voltage regulator. The system was designed with low energy consumption in mind, which motivated the adoption of the NB-IoT module due to its reduced power usage both during data transmission and in “sleep mode” periods. This work closely resembles the system developed in this article, except for the use of the BQ25504 in this study and the exclusive measurement of CO with Wi-Fi transmission only.

Touati et al. [5] proposed a prototype capable of managing different energy sources, such as PZT (piezoelectric), RF (radio frequency), PV (photovoltaic), and thermal energy. The system includes sensors to measure temperature, CO, and NO₂. The power capacity provided by the sources is 38.3 mW, which significantly exceeds the energy consumption of the nodes (approximately 0.85 mW). The prototype includes a star-shaped printed circuit board, allowing the insertion of various sensors at the front tips and energy harvesting components at the rear tips, except for the core, where a PV cell was installed. This development employed multiple energy sources, unlike the PV-only approach adopted in this study. Moreover, the demonstrated power output was substantially lower than that achieved in the prototype presented in this article.

Crescini et al. [6] introduced a multiparametric sensor node that integrates measurements of NH₃, NO₂, CO, H₂S, NO, Cl₂, temperature, and humidity. Multiple energy harvesting sources are also employed: vibration, using a piezoelectric material coupled to the system; a thermal energy harvesting material with a capacity of 0.5 mW; a PV cell for indoor use and another for outdoor application; and RF energy harvesting through electromagnetic waves.

Yue et al. [7] developed a CO₂ meter for indoor environments aimed at ensuring air quality control. Measurements were performed by a node powered by a photovoltaic cell (50 mm × 20 mm) and a supercapacitor as the energy storage system. The CO₂ sensor can measure concentrations ranging from 0 to 5,000 ppm (parts per million) and must transmit data every 30 s throughout the week. An important feature of this project is the use of a PV cell in indoor environments. Considering the system's energy requirements, the study defined a minimum illuminance of 200 lx to ensure system operation, which is feasible in office buildings. A low-power controller, the Moteino—a derivative of the conventional Arduino with reduced energy consumption—was used in this system. Additionally, the PMIC adopted was the LTC4071. In these two studies, the energy management system differed from the BQ25504 employed in this work. Furthermore, RF transmission was used instead of the Wi-Fi implementation adopted in this study.

The monitoring system presented in [8] aims to measure a specific particle responsible for COVID-19 transmission at two concentration levels: PM2.5 and PM10. The system performs indoor and outdoor measurements, transmitting data via the internet to a central unit for extended analysis. The total power consumption is 1.68 mW, powered by a supercapacitor designed to provide sufficient energy for reading and transmission tasks. Although initially powered by a 5 V DC source, the system can be adapted to use alternative power sources, such as a PV cell. Compared to the prototype developed in this study, a lithium battery was employed instead of a supercapacitor, and a 3.3 V operating voltage was used, already incorporating a PV panel.

Table 1. Summary of technical information from literature

Ref	PMIC	SOURCE	POWER	NOMINAL VOLTAGE (V)
[4]	BQ25506	PV	269 mW	5
[5]	LTC3109	PZT, RF, PV	0.8 mW	4.1
[6]	LTC3109	PZT, RF, PV	20 mW	4.1
[7]	LTC4071	PV	118.9 μ w	5.5
[8]	Not used	SUPERCAP	1.68 mW	4.68

2.1 PMIC

The PMIC (Power Management Integrated Circuit) regulates the energy flow between the power source and other system components, including an energy storage system, which may consist of a battery or supercapacitor, and the system controller.

The PMIC's operation primarily involves Maximum Power Point Tracking (MPPT), a technique that maximizes the power extracted from energy sources. This component is also responsible for managing the energy flow, switching between available sources and the energy storage system.

Table 2 provides a list of PMICs commercially available, including specifications extracted from the manufacturers' datasheets. These devices typically measure around 3.5 mm × 3.5 mm. Since PMICs are often integrated into compact circuits, such as prototypes, their reduced size is advantageous for project development [9]. Although some datasheets lack detailed information on power capacity or efficiency, all PMICs follow similar design principles and application guidelines. Notably, the BQ25570 model is frequently used in prototypes due to its broad input range for power sources, higher power capacity, and improved efficiency.

In addition to photovoltaic energy, which serves as the primary energy source in this study, PMICs also support other energy sources identified in the literature.

Table 2. Commercial PMICs adopted in the literature

PMIC	SOURCE	PIN	VIN(V)	VOUT(V)	EFIC (%)
AEM10940	E-PEAS	50 mW	0.1–3.5	1.8–3.3	90
S6E10	CYPRESS	N/A	3.4	2.5–3.3	N/A
BQ25506	TEXAS	400 mW	0.1–5.1	2–5.5	93
MAX14720	MAXIM INT	250 mW	1.8–5.5	2.5–5.5	N/A
SPV1050	STMICRO	N/A	2.2–5.3	2.6–5.3	93
LTC3109	MICROCHIP	50 mW	0.3–6	2.35–5	N/A
LTC4071	LINEAR	N/A	4–4.2	0.3–6	N/A

2.2 Air Quality Sensors

Sensors play a crucial role in control and automation systems [10] and are widely used in industries, automobiles, medical devices, and consumer electronics. In the context of air quality, sensors are designed to measure and monitor the presence of air pollutants, detecting gases [10], particles, and other contaminants. Table 3 presents a selection of commercial sensors and multi-sensor models for air quality monitoring. The table includes parameters such as voltage, current, and power consumption, which were evaluated in this study.

Table 3. Energy consumption of sensors

SENSOR	VOLTAGE (V)	CURRENT (mA)	POWER (mW)
AR MICS6814	4.9–5	32	81
AR ZP97 MP503	5	60	300
AR PMS3003	5	80	400
SGP30	2	48.8	97.6
HDC1000	5	0.22	1.1
XSENSIV	3.3	10	30

The sensors listed in Table 3 are designed to measure parameters such as carbon monoxide, nitrogen dioxide, PM2.5, PM10, temperature, and humidity.

3 Methodology

The proposed system consists of an air quality monitoring architecture composed of sensors, microcontrollers, power management integrated circuits (PMICs), and photovoltaic modules. The sizing methodology includes the following steps:

A - Energy Consumption Assessment

Survey of the power required by sensors and microcontrollers, considering voltage, current and electrical power.

B - PMIC Selection

Choice of the appropriate power management circuit based on supported power and available energy sources [11].

C - Battery Sizing

Calculation of the battery capacity required to ensure autonomy during periods without energy generation.

D - Photovoltaic Module Sizing

Definition of the power and number of solar panels required, considering local solar irradiation.

The prototype developed uses the MiCS-6814 sensor for carbon monoxide (CO) measurement, the ESP8266 microcontroller for data acquisition and transmission, and the BQ25504 PMIC for energy management. Data transmission is performed through Wi-Fi, with information sent every 30 min. A photovoltaic module was chosen as the energy source for this project. The steps of the proposed design methodology are described in the following steps:

Steps A and B – Energy Required by Sensors and Controllers

The first step in assessing the energy required by the system is to survey sensors of various types, such as temperature, humidity, CO₂, and others, to gather information regarding their voltage, current, and power consumption.

The second step involves evaluating the energy consumption of the main microcontrollers responsible for managing sensor data. For example, microcontrollers from the Arduino family, including the ESP8266 and its derivatives, operate at a voltage of 3.3 V and consume up to 500 mA when powered by external sources or USB ports, as per the component datasheets.

To determine the total power required by the system, the combined power consumption of the sensor and controller (including software and transmitter) is calculated using Eq. (1).

$$PR = PS + PC \quad (1)$$

In which:

- PR : Total power required [W]
- PS : Power consumed by the sensor [W]
- PC : Power consumed by the controller [W]

In cases where power consumption is not directly provided in datasheets, it can be calculated as the product of the nominal voltage and current of the component.

The daily energy consumption of the system is then evaluated based on usage levels. The average energy consumption is given by Eq. (2):

$$E = PR \times Td \times Th \quad (2)$$

In which:

- E : Monthly energy consumption of the component or block [Wh/month]
- PR : Power required by the component or block [W]
- Td : Number of days of operation per month
- Th : Number of hours of operation per day

In the sequel, the capabilities of the most common PMICs should be evaluated to ensure they meet the power requirements of the integrated components. The PMICs analyzed typically support input voltages ranging from 100 mV to 6 V, output voltages from 0.3 V to 6 V, and manageable power levels up to 400 mW. There are a lot of PMICs that support the use of batteries or capacitors and can handle various energy harvesting sources.

Step C – Battery Sizing

Battery sizing considers three key aspects [12]:

1. The daily energy requirement, calculated as the sum of the energy consumption of each component using Eq. 4.
2. The number of days the battery must be able to power the system in the absence of energy generation from the primary source.
3. The depth of discharge (DoD).

The daily energy requirement is quantified in Watt-hours (Wh). This value accounts for the power required by the system over the total time the battery must support the load. The depth of discharge is a critical factor in battery durability, as a higher DoD reduces the battery's lifespan. For systems installed in remote or hard-to-access locations, a longer battery life is preferred to minimize maintenance costs.

Commercially available batteries have varying DoD values (e.g., 20%, 50%, 80%). The DoD indicates the remaining charge percentage after the specified period of energy absence. The battery capacity is calculated using Eq. (3).

$$Cb = \frac{E}{Vr \cdot Dod} \quad (3)$$

In which:

- Cb : Battery capacity [Ah]
- E : Monthly energy consumption of the component or block [Wh]
- Vr : Required voltage of component or block [V]
- DoD: Depth of discharge (expressed as a decimal, e.g., 0.5 for 50%)

Step D – Photovoltaic Module Sizing

The sizing of photovoltaic modules requires parameters related to the system architecture and the environmental conditions of the installation site. This step involves two stages [4]:

1. Determining the total power of the photovoltaic module array based on environmental parameters.
2. Selecting a commercial module model and configuring the electrical connections.

The total power of the photovoltaic array is calculated using Eq. (4)

$$PFV = \frac{E}{TD \cdot HSPma} \quad (4)$$

In which:

- PFV : Total power of the photovoltaic array [W]
- E : Monthly energy consumption of the component or block [kWh/month]
- TD : Performance ratio of photovoltaic modules (typically 70% to 80%)
- $HSPma$: Annual average peak sun hours [Wh/m²/day], which represents the daily solar irradiation divided by the Standard Test Conditions value for irradiance (1,000 W/m²). This value is location-specific and can be obtained using software tools such as PVsyst[©] or databases like the [13].

The number of photovoltaic modules required is determined using Eq. (5)

$$Quatmod = \frac{PFV}{Potplaca} \quad (5)$$

In which:

- $Quatmod$: Number of modules required
- PFV : Total power of the photovoltaic array [W]
- $Potplaca$: Power rating of the commercial module [W]

4 Prototype

The prototype, designed according to the proposed methodology, integrates the Sensor + Software + Transmission components as follows:

4.1 Sensor and Microcontroller

The carbon monoxide (CO) sensor selected for the prototype is the MiCS-6814, which operates at 3.3 V or 5 V and has a maximum power consumption of 88 mW.

The ESP8266 microcontroller from the Arduino family was chosen due to its Wi-Fi communication capability, enabling data transmission via telemetry. Despite its integrated Wi-Fi module, the ESP8266 maintains low energy consumption, with 4 MB of flash memory for medium-complexity programming in C/C++. Its specifications include:

- Operating voltage: 3.3 V or 5 V
- Maximum current: 46 mA
- Maximum power consumption (at 3.3 V): 151 mW

4.2 Power Management and Energy Harvesting

The BQ25504 PMIC was selected due to its compatibility with solar energy and battery inputs, providing the necessary voltage and power for the sensor and microcontroller. However, the CJMCU-25504 board which integrates BQ25504 lacks an adjustable DC boost converter. To ensure stable voltage output, an external converter + regulator was added.

4.3 Energy Optimization

To optimize energy consumption, the microcontroller was programmed to enable the “sleep mode” between data transmissions. During this period, all processors are deactivated, except for the Real-Time Protocol (RTP), which tracks the sleep duration. The system wakes up every 30 min, transmits data, and returns to sleep, minimizing energy usage.

4.4 System Architecture

The prototype architecture (Fig. 1) consists of:

1. A solar panel as the primary energy source, connected to the BQ25504 PMIC.
2. A battery to maintain input voltage during periods of low or no solar irradiation with 3.100 mAh
3. The ESP8266 microcontroller, powered by the PMIC, collects data from the sensor and transmits it via Wi-Fi to a web server.

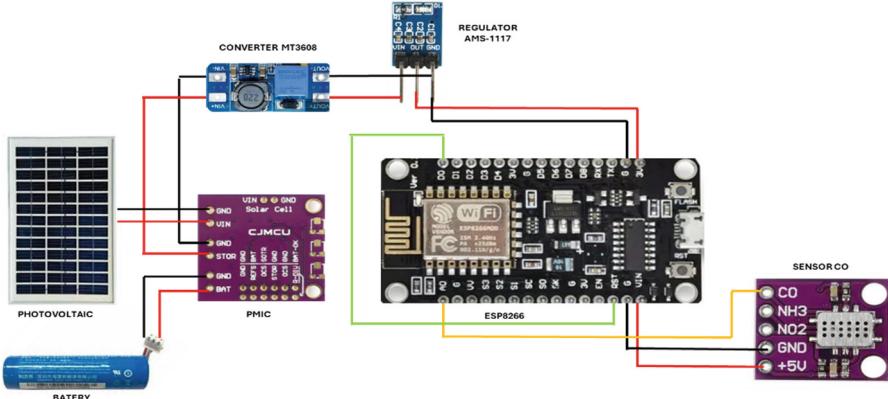


Fig. 1. The prototype circuit.

4.5 Data Transmission and Telemetry Simulation

A web server was configured on a laptop connected to the same Wi-Fi network as the ESP8266 to simulate telemetry. Data is transmitted every 30 min [14], ensuring efficient energy use while maintaining system functionality. Figure 2 presents the data obtained from the ESP8266 readings. It records the date and time of each reading, along with the CO value measured at that moment, at 30 min intervals.

```
ON Prompt de Comando - node server.js
Microsoft Windows [versão 10.0.22631.4249]
(c) Microsoft Corporation. Todos os direitos reservados.

C:\Users\emers>esp32-server
'esp32-server' não é reconhecido como um comando interno
ou externo, um programa operável ou um arquivo em lotes.

C:\Users\emers>cd esp32-server

C:\Users\emers\esp32-server>node server.js
Servidor rodando na porta 3000
Dados recebidos do ESP8266: { Date: '2024-10-5', Hora: '11:54:55', CO: 0.000406536 }
Dados recebidos do ESP8266: { Date: '2024-10-5', Hora: '12:31:12', CO: 0.000406536 }
Dados recebidos do ESP8266: { Date: '2024-10-5', Hora: '12:16:16', CO: 0.0000905875 }
Dados recebidos do ESP8266: { Date: '2024-10-5', Hora: '12:45:15', CO: 0.000638235 }
Dados recebidos do ESP8266: { Date: '2024-10-5', Hora: '13:14:9', CO: 0.001896734 }
Dados recebidos do ESP8266: { Date: '2024-10-5', Hora: '13:42:1', CO: 0.000406536 }
Dados recebidos do ESP8266: { Date: '2024-10-5', Hora: '14:10:39', CO: 0.000215423 }
Dados recebidos do ESP8266: { Date: '2024-10-5', Hora: '14:38:54', CO: 0.000406536 }
```

Fig. 2. Screen results from java notebook

5 Results and Discussion

The prototype was assembled and installed at an urban residential rooftop, as shown in Fig. 3. It was fixed at the rooftop with an inclination of 23°, which is close to the local latitude, to improve the exposition to solar irradiance.

The city chosen for testing the prototype was Guarulhos, state of São Paulo - Brazil, but precisely at latitude and longitude $-23.44, -46.49$.

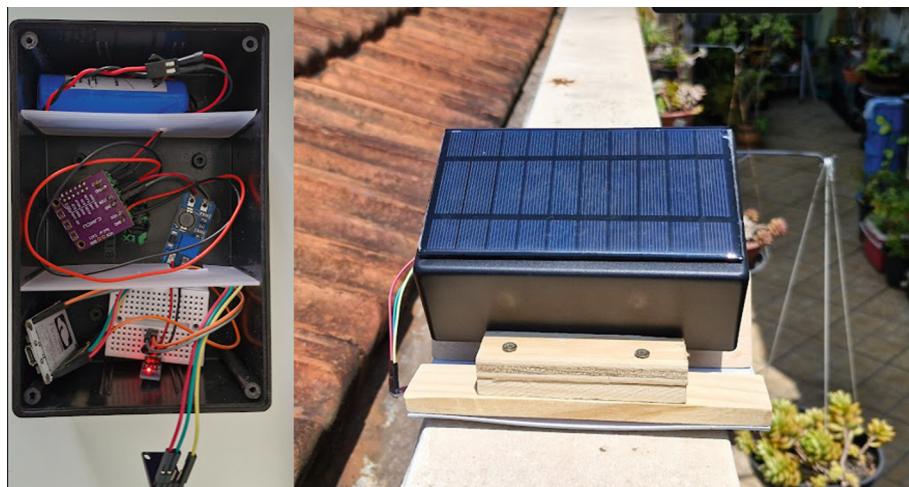


Fig. 3. Prototype

With these conditions, the prototype kept operating continuously for four days. The results demonstrated that the system maintains a 3.3 V supply for the microcontroller, with the battery being charged by the solar panel during the day. The graph shown in Fig. 4 was generated using measurement data collected over a single day for the following variables: battery voltage, voltage supplied by the solar panel, and voltage consumed by the architecture. Table 4 presents the data collected to generate Fig. 4.

It was observed that during the nighttime period, from 18:00 to 21:00 there is a sharp decline in the energy supply provided by the photovoltaic panel. This decrease is reversed at dawn, between 5:00 and 7:00 in the morning.

Between 21:00 and 5:00, the energy supply from the photovoltaic panel shows relative stability. However, the system's architecture voltage is continuously maintained during these variations by the battery integrated into the system, which can sustain the system's operation for up to two days in the absence of solar energy.

Table 4. 24 h voltage variation measured during the prototype operation.

Time	Voltage Supplied Battery (V)	Voltage Supplied by the Solar Panel (V)	Voltage Consumed by the architecture (V)
12	3.7	5.0	3.3
13	3.7	5.0	3.3
14	3.7	5.0	3.3
15	3.7	5.0	3.3
16	3.7	5.0	3.3
17	3.7	5.0	3.3
18	3.7	5.0	3.3

(continued)

Table 4. (*continued*)

Time	Voltage Supplied Battery (V)	Voltage Supplied by the Solar Panel (V)	Voltage Consumed by the architecture (V)
19	3.7	4.5	3.3
20	3.7	3.5	3.3
21	3.7	1.0	3.3
22	3.7	1.0	3.3
23	3.7	1.0	3.3
24	3.7	1.0	3.3
01	3.7	1.0	3.3
02	3.7	1.0	3.3
03	3.7	1.0	3.3
04	3.7	1.0	3.3
05	3.7	1.0	3.3
06	3.7	3.5	3.3
07	3.7	5.0	3.3
08	3.7	5.0	3.3
09	3.7	5.0	3.3
10	3.7	5.0	3.3
11	3.7	5.0	3.3

The maximum power consumed during data transmission was 356 mW, while the average consumption in standby mode was 118 mW.

The proposed sizing methodology was validated, with the calculation spreadsheet allowing precise specification of the necessary components. The use of photovoltaic energy proved efficient, especially in remote or hard-to-access locations where conventional electrical infrastructure is unfeasible.

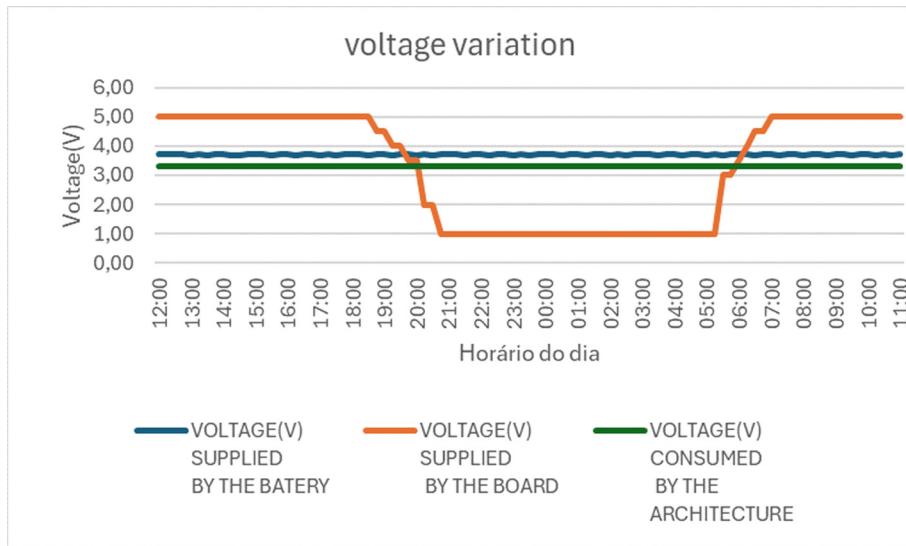


Fig. 4. 24 h Voltage variation during the prototype operation.

6 Conclusion

This work demonstrated the feasibility of autonomous air quality monitoring systems powered by photovoltaic energy. The proposed sizing methodology facilitates the replication of the system in different contexts, contributing to the expansion of sustainable environmental monitoring networks. The use of photovoltaic energy harvesting proved to be efficient for powering low-consumption architecture, especially in remote or hard-to-access locations.

As a future perspective, the investigation of other energy harvesting sources, such as piezoelectric and thermoelectric, is suggested to complement or replace photovoltaic energy under certain conditions. Additionally, the integration of low-power communication networks, such as LoRaWAN, can expand the system's applicability in areas with limited Wi-Fi coverage.

References

1. United Nations: Resolution 73/300 (2019)
2. Heidari, A., et al.: Air Quality Monitoring: A Comprehensive Review. *Environmental Science & Technology* (2021)
3. Prauzek, M., Konecny, J., Borova, M., Janosova, K., Hlavica, J., Musilek, P.: Energy harvesting sources, storage devices and system topologies for environmental wireless sensor networks: a review. *Sensors* **18**, 2446 (2018). <https://doi.org/10.3390/s18082446x>
4. Das, P., et al.: Energy harvesting-enabled 5G advanced air pollution monitoring device. In: 2020 IEEE 3rd 5G World Forum (5GWF), pp. 218–223. IEEE (2020)
5. Touati, F., Legena, C., Gali, A., Crescini, D., Mnaover, A.: Renewable energy-harvested sensor systems for air quality monitoring. In: 2014 26th International Conference on Microelectronics (ICM), pp. 160–163 (2014)

6. Crescini, D., Touati, F., Galli, A.: Multiparametric sensor node for environmental monitoring based on energy harvesting. *Atmosphere* **13**(2), 321 (2022)
7. Yue, X., et al.: Development of an indoor photovoltaic energy harvesting module for autonomous sensors in building air quality applications. *IEEE Internet Things J.* **4**(6), 2092–2103 (2017)
8. Kuncoro, C.B.D., Adristi, C., Asyikin, M.B.Z.: Smart wireless particulate matter sensor node for IoT-based strategic monitoring tool of indoor COVID-19 infection risk via airborne transmission. *Sustainability* **14**(21), 14433 (2022)
9. Kurian, D., et al.: Self-powered IOT system for edge inference. In: 2020 21st International Symposium on Quality Electronic Design (ISQED), pp. 302–305. IEEE (2020)
10. Bhat, G., et al.: Self-powered wearable IoT devices for health and activity monitoring. *Found. Trends® Electron. Des. Autom.* **13**(3), 145–269 (2020)
11. Kjellby, R.A., et al.: Long-range & self-powered IoT devices for agriculture & aquaponics based on multi-hop topology. In: 2019 IEEE 5th World Forum on Internet of Things (WF-IoT), pp. 545–549. IEEE (2019)
12. Villalva, M.G.: Photovoltaic Solar Energy: Concepts and Applications. Editora Érica (2015)
13. CRESESB – Center of Reference to Solar and Eolic Energy Sérgio de S.Brito. <https://www.creesb.cepel.br/index.php?section=sundata>
14. Kjellby, R.A., et al.: Self-powered IoT device based on energy harvesting for remote applications. In: 2018 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS), pp. 1–4. IEEE (2018)



How Clean is the Air You Breathe During Urban Walk? A Case Study of Central London

Nikhil Ravindra¹ , Amin Al-Habaibeh¹ , and Benachir Medjdoub²

¹ Product Innovation Centre, Nottingham Trent University, Nottingham, UK
nikhil.ravindra2024@my.ntu.ac.uk, amin.al-habaibeh@ntu.ac.uk

² Architecture Department, Nottingham Trent University, Nottingham, UK
benachir.medjdoub@ntu.ac.uk

Abstract. This paper presents a case study on air quality in Central London with a focus on particulate matter (PM10 and PM2.5). The relationship between road width, traffic patterns, wind direction, building orientation, and particulate matters (PM10 and PM2.5) concentrations is explored in this experimental study via a walking journey. GPS geographical analysis and portable air quality sensors are used to assess the air quality. The key results indicate that narrower roads with high buildings in general exhibit irregular PM spikes which could be explained by trapped emissions and busy traffic. Larger roads (≥ 20 m in width) have presented a more stable air quality which can be explained by improved air circulation and consistent traffic flow. The results demonstrate the need for enhanced measures with specialized interventions such as the plantation of green walls and trees (green infrastructure) combined with improved traffic control. The results also indicate the importance of traffic management and the importance of policies and strategies such as London's Urban Greening Factor (UGF), London Low Emission Zone (LEZ), Ultra Low Emission Zone (ULEZ), Congestion Charge Zone, and the support of electric transportation. The novel experimental methodology and findings presented in this paper provide a strong foundation for a broader and more comprehensive future study.

Keywords: air quality · green infrastructure · artificial intelligence · electric mobility · smart cities

1 Introduction

It has been established that exposure to air pollution is a key risk factor for human health [1]. According to new figures in 2025, over 90% of neighborhoods in England and Wales have high levels of air pollution, increasing the risk of cancer and heart attacks. Based on the analysis of Friends of the Earth on governmental statistics of Nitrogen Dioxide (NO_2) and particulate matter PM2.5 levels between 2021 and 2023, more than 33,000 neighborhoods had air pollution levels above the World Health Organization (WHO) recommended safe threshold, with London residents being the most severely impacted, making it one of the most polluted places in the UK [2].

This covers every neighborhood in the city of London as well as the six boroughs of London: Hackney, Islington, Kensington and Chelsea, Newham, Tower Hamlets, and Westminster. Fossil fuel burning is found to be the main source of NO₂. Nearly half of the air pollution in London is caused by polluting vehicles followed by power stations, affecting all its boroughs and not just central London [2]. The main sources of PM2.5 pollution are industry and road transportation; with significant contribution from domestic combustion due to the usage of wood-burning stoves [2].

However, because the neighborhoods with the highest pollution levels also had the fewest car owners, the study discovered that many residents of the most polluted regions were dealing with issues brought on by others [2]. Unfortunately, in several locations, the people least responsible for air pollution are the ones who are facing higher negative impacts of it [3].

The first person officially on record to have air pollution to be the cause of death was Ella Adoo-Kissi Debrah, a child who passed away from asthma, in 2013 [2]. A coroner from south London in 2020, made history by citing air pollution as the cause of Ella's death [2]. Her mother fought for years to have the effects of air pollution recognized as a cause of death of the nine-year-old who lived within 30 m of the South Circular, one of London's busiest roads [2].

As a remembrance, her statue was unveiled on the 4th of April 2025 in London, Mountsfield Park, SE6 1AN [4].

A recent shift towards EVs and stricter policies have helped combat air pollution. Transport for London (TfL)'s initiative of an Ultra-Low Level Emission Zone (ULEZ) in April 2019 has already made a substantial difference to the air quality levels [5]. Although in 2023, the city has decided to extend the ULEZ, more needs to be done to combat the global challenge of air pollution.

This paper investigates an experimental pilot case study in Central London during a walking process to explore air quality using a GPS location sensor and a portable PM10 and PM2.5 sensor.

2 Research Outline

Most of the airborne toxic particles are extremely small to be visible to the naked eye. Air pollution can lead to an increased risk of dementia, and other life-changing illnesses including lung diseases and asthma [6]. Nearly half a million Londoners are living with asthma, making them more vulnerable to toxic air impacts [6]. Every year, air pollution leads to the death of thousands of premature in London [6]. The study indirectly addresses the topic of energy informatics linked to green mobility.

This paper aims to sample the current air quality in Central London (PM10 and PM2.5) during a walking journey and discuss the experimental findings; and the potential way forward to enhance air quality via the use of electric vehicles (EV), and promote low-carbon transit in addition to other suitable measures.

3 Literature Study

3.1 Engine Idling and Ultra-Low Emission Zone

Research by the World-Wide Fund for Nature (WWF) highlights the fact that switching off an idling car engine when stopped for 10 s or more has multiple benefits; it is energy efficient, economical, reduces air pollution, slows climate change, and overall better for the planet. Nevertheless, repetitively turning on and off the car's engine can cause wear to some parts, unless it has been built with stop-start technology. To overcome this challenge many city councils around the globe recommend turning off the engine when the vehicle is stopped for a minute or more [5].

To combat engine idling, the London city corporation has come up with measures such as penalties for non-compliance, placing 'no idling' signs in hotspots, and the pan-London idling action project supported by the Mayor's air quality fund [7]. Although numerous efforts including ULEZ have been taken to reduce the city's air pollution and improve air quality, there is more to be achieved considering the sheer volume of road transport emissions [6]. Most importantly, even low levels of emissions can cause significant damage to human health [8].

The other two tiers of emission control regulations that play a vital role are the Congestion Charge Zone (CCZ) and the Expanded Ultra-Low Level Emission Zone (EULEZ) [9]. Since the expansion of ULEZ to all London boroughs, analysis by [10] shows fluctuating traffic congestion levels, given the complexities of the city's antiquated road network. The rush hour traffic is relentlessly dense, leading to slow-moving traffic, which is a core issue causing an increase in transport emissions. In such scenarios, can over-reliance on ULEZ and greater EV uptake reduce transport emissions and improve air quality?

3.2 Electric Vehicles and Integration Using Artificial Intelligence

As EVs are heavier than petrol or diesel cars, they most likely produce more nanoparticle matter pollution from their tires and brakes [11]. The production process is one of the primary reasons why EVs are not environmentally friendly. Emissions from EV production are higher than those from conventional automobiles. The lithium-ion batteries, which require mining operations to obtain their rare earth metals, are the cause of this. Charging EVs requires electricity, which can result in the burning of fossil fuels [12].

Even if the production of EVs may be harmful to the environment in some aspects, it is still far less than the pollutants that traditional cars emit. By 2050, electricity must only come from renewable sources to reach net zero in the UK [12]. The use of AI is essential in the energy management of EV, as it improves user experience, optimizes energy expenditure, enhances coordination, and has advanced user data understanding [13]. Expanded infrastructure incorporating enhanced accessibility of EV chargers, not just improves traffic flow but also cuts down PM2.5 emissions by 1.3 to 2.2% [9]. AI-enabled smart and optimized charging can lower carbon emissions by up to 40% and annually could save nearly 800 lb of CO₂ per vehicle, contributing to sustainable energy for UK EV fleet management [12]. However, reducing the number of automobiles on the road is one of the most effective strategies to lower local air pollution [14].

3.3 Innovative Green Infrastructure

Just shifting to EVs might not significantly help in reducing the air quality problem. Private vehicles have adverse impacts on health and quality of life, particularly considering the emissions, air quality, and congestion caused. Well-connected green mobility corridors can address the need for sustainable mobility means by promoting a shift away from cars to walking, cycling, or micro-mobility [15, 16]. Green mobility corridors use a series of swales, street trees, rain gardens, and pervious pavements, creating a cooler oasis and promoting active mobility (Fig. 1). It is a type of blue-green infrastructure (BGI) that has the potential to reduce the urban heat island (UHI), reduce greenhouse gas (GHG) emissions, and improve air quality, due to the greater presence of vegetation and water [17].

Micro-mobility devices include bicycles, velomobiles, e-scooters [18], e-skateboards, shared bicycle fleets, and electric pedal-assisted bicycles. These operate at speeds not exceeding 28 miles per hour (mph), lightweight mini-vehicles that can be privately owned or shared, and can be human-powered or electric. Micro-mobility is economical, flexible, and sustainable, having the potential to reduce reliance on private vehicles; especially for short-distance travel [17]. Nevertheless, some have traffic risk issues associated with them, including accidents [18]. Urban air quality, particularly in relation to PM10 and PM2.5, could also be influenced by other urban factors [19].



Fig. 1. Photo of an e-bike facility in Nottingham; and green infrastructure along river Manifold in Ilam village.

4 Methodology

The study presents a ground-level air quality monitoring exercise that was carried out in Central London. Data was collected using a hand-held portable battery-powered air quality monitoring device which combines sensor and logger. The device's battery can be electrically charged.

The device (Fig. 2) is a Series 500 – Portable Air Quality Monitor sensor head with a data logger by Aeroqual company. The study focused on the emissions of particulate matter $10\text{ }\mu\text{m}$ (PM10) and particulate matter $2.5\text{ }\mu\text{m}$ (PM2.5) during a walking journey in Central London. To account for interferences (e.g. humidity), the device incorporates algorithms into the laser particle counter (LPC) for PM, which employs optimized signal processing with low-noise electronics. An air pollutant level in this paper refers to the concentration of pollutants at a level of nearly a meter from the ground.

MATLAB software was utilized to visualize and analyze the impact on air quality. Google Maps were used as a platform for spatial mapping of the route taken from start to endpoint. Ethical approval was received in relation to this work.



Fig. 2. Photo of the air quality monitoring device used in this study.

5 Data Collection

A nearly 2 h walk in Central London streets (Fig. 3) in an ULEZ, collecting both, PM10 and PM2.5 data to measure road-side air quality levels in the heart of the smart city of London around noon on a spring day. The data collected was at a frequency of one minute per sample. The walking journey took place as part of The City of London Urban Climate Walking Tour [20].

The starting point was the Martha Smith Memorial water fountain and the finishing point was at Undersholt Street, covering nearly 1.2 miles, on the 1st of March 2025 (Saturday). Global Positioning System (GPS) application was used to record the location. Average road widths were calculated from Google Earth Engine (GEE), by measuring the distance between opposite buildings. GEE as a cloud-based platform is used to process and analyze satellite imagery. Images were taken by the research team or from Google Map [21], Google Earth or Google Street view as referenced on each image.

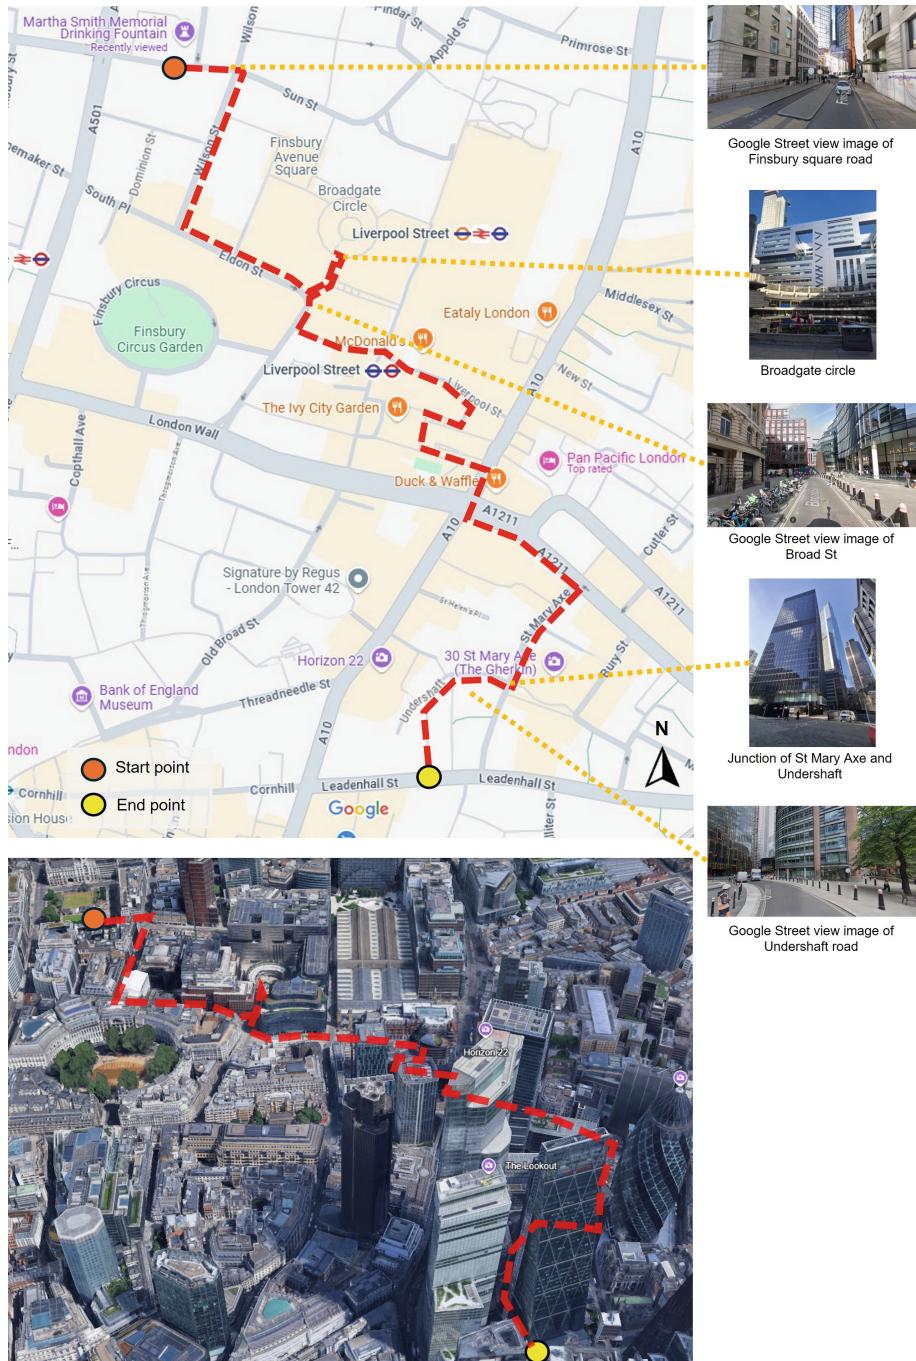


Fig. 3. Route taken from start to end point, highlighted on Google Maps and Google Earth 3D image.

The Met Office [22] reported the atmospheric condition during the walk time was clear sky, and temperature was between 7 to 8 °C. The real feel temperature was between 4 to 5 °C; and relative humidity was between 60 to 67% [23]. Figure 4 presents an example of the analysis of the angle between the wind direction as obtained from the Met office and the street angle as obtained from Google Map [21].

Table 1. Approximate time (rounded off to the nearest minute) and location of the urban walk.

Start time	End time	Minutes	Location	Avg. Road width
11:18 AM	11:32 AM	14	M.S. Memorial water fountain, Finsbury Sq.	19 m
11:32 AM	11:54 AM	22	Wilson St	14 m
11:54 AM	12:02 PM	08	Eldon St	13 m
12:02 PM	12:15 PM	13	Broadgate Circle	Pedestrian-only
12:15 PM	12:16 PM	01	Broad St Pl	17 m
12:16 PM	12:19 PM	03	Liverpool St	24 m
12:19 PM	12:20 PM	01	White Hart Ct	Pedestrian-only
12:20 PM	12:21 PM	01	Alderman's walk	Pedestrian-only
12:21 PM	12:33 PM	12	Passageway from Alderman's Walk to Bishopsgate	Pedestrian-only
12:33 PM	12:34 PM	01	Bishopsgate (A10)	24 m
12:34 PM	12:38 PM	04	Camomile St (A1211)	20 m
12:38 PM	12:48 PM	10	St Mary Axe	13 m
12:48 PM	13:06 PM	18	Undershaft	19 m
Total time	108 min (1 h 48 min)			

Table 2. Comparison between PM levels and the absolute angle difference between the road and the wind direction.

Location	Wind speed	Angle of road	I Angle difference I	PM10 levels	PM2.5 levels
Finsbury Sq	N-E 4.9 mph	97°	52°	25.86 $\mu\text{g}/\text{m}^3$	10.43 $\mu\text{g}/\text{m}^3$
Wilson St		21°	24°	33.86 $\mu\text{g}/\text{m}^3$	10.95 $\mu\text{g}/\text{m}^3$
Eldon St		116°	71°	36.25 $\mu\text{g}/\text{m}^3$	11.00 $\mu\text{g}/\text{m}^3$
Liverpool St	N-E 6.0 mph	113°	68°	24.67 $\mu\text{g}/\text{m}^3$	7.33 $\mu\text{g}/\text{m}^3$
Bishopsgate (A10)		29°	16°	21.00 $\mu\text{g}/\text{m}^3$	5.00 $\mu\text{g}/\text{m}^3$
Camomile St (A1211)		123°	78°	21.75 $\mu\text{g}/\text{m}^3$	6.25 $\mu\text{g}/\text{m}^3$
St Mary Axe		33°	12°	18.30 $\mu\text{g}/\text{m}^3$	5.50 $\mu\text{g}/\text{m}^3$

(continued)

Table 2. (continued)

Location	Wind speed	Angle of road	I Angle difference I	PM10 levels	PM2.5 levels
Undershافت		47°	2°	16.22 $\mu\text{g}/\text{m}^3$	4.39 $\mu\text{g}/\text{m}^3$



Fig. 4. Absolute angle calculation example for Wilson St (presented on background from on Google Maps).

6 Results

Figure 5 presents the concentration of PM10 and PM2.5 over time; visualized using MATLAB with key locations as in Fig. 6. The results are analyzed in conjunction with Table 1. Three scenarios are studied here: wider roads (≥ 20 m), narrower roads (13–19 m), and pedestrian-only zones. The road width is the average width for the corresponding street (Fig. 6).

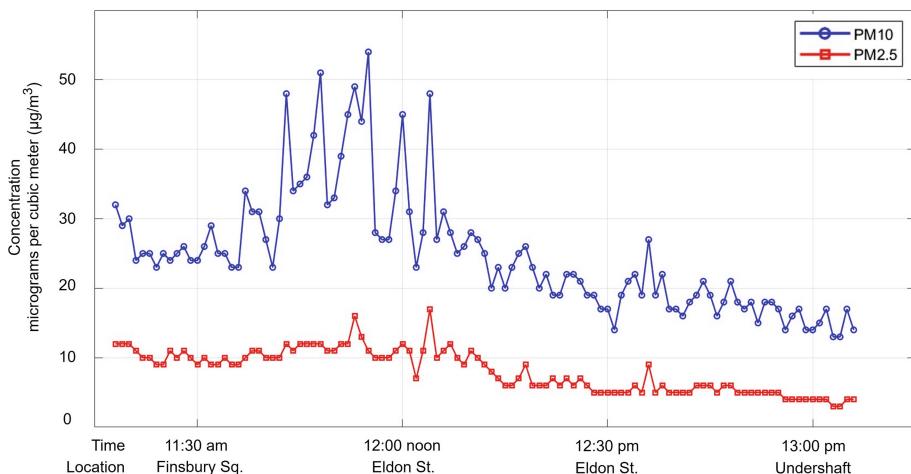


Fig. 5. PM10 and PM2.5 concentration over time visualization using MATLAB.

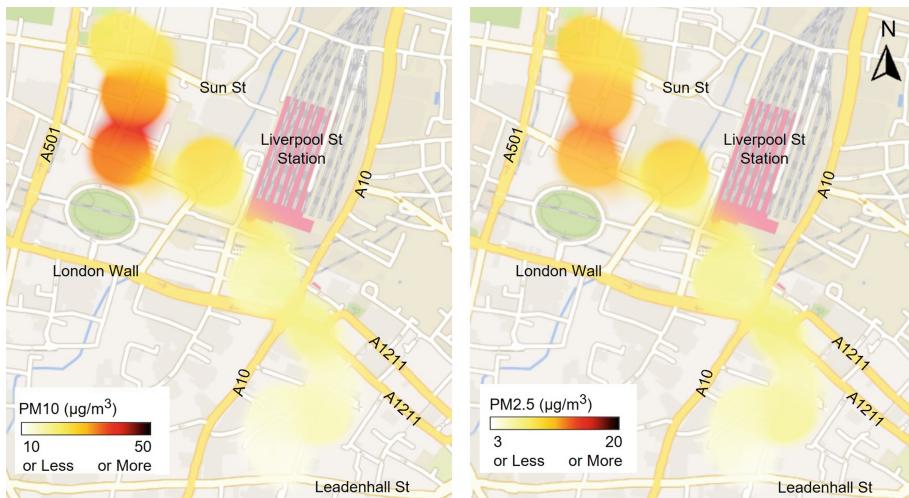


Fig. 6. PM10 and PM2.5 levels at different locations presented on a Google Map background.

Of the total 108 min, 8 min were spent along wider roads, 73 min along narrower roads, and a total of 27 min in the pedestrian zones. The key observations and data measurements for the three scenarios are presented in Tables 3A, 3B and 3C.

Table 3A. Average PM levels in wider roads (≥ 20 m).

Road, width, and time	PM10 levels	PM2.5 levels	Summary
Liverpool St (24 m, 12:16 PM - 12:19 PM)	24.67 $\mu\text{g}/\text{m}^3$	7.33 $\mu\text{g}/\text{m}^3$	Both PM10 and PM2.5 levels are relatively high
Bishopsgate (24 m, 12:33 PM - 12:34 PM)	21.00 $\mu\text{g}/\text{m}^3$	5.00 $\mu\text{g}/\text{m}^3$	Variations in both PM10 and PM2.5 levels, possibly due to increased vehicular activity on A class road
Camomile St (20 m, 12:34 PM - 12:38 PM)	21.75 $\mu\text{g}/\text{m}^3$	6.25 $\mu\text{g}/\text{m}^3$	

Table 3B. Average PM levels in narrower roads (13–19 m).

Road, width, and time	PM10 levels	PM2.5 levels	Summary
M.S. Memorial Water Fountain, Finsbury Sq (19 m, 11:18AM - 11:32 AM)	25.86 $\mu\text{g}/\text{m}^3$	10.43 $\mu\text{g}/\text{m}^3$	Moderate fluctuation of PM10 and PM2.5 levels, however, overall concentrations are relatively stable

(continued)

Table 3B. (*continued*)

Road, width, and time	PM10 levels	PM2.5 levels	Summary
Wilson St (14 m, 11:32AM - 11:54 AM)	33.86 $\mu\text{g}/\text{m}^3$	10.95 $\mu\text{g}/\text{m}^3$	PM10 shows more variations in comparison to PM2.5 which fluctuates slightly. The reason could be because of the longer time spent along this street
Eldon St (13 m, 11:54 AM - 12:02 PM)	36.25 $\mu\text{g}/\text{m}^3$	11.00 $\mu\text{g}/\text{m}^3$	A spike in PM10 levels before declining, but PM2.5 remains relatively stable
St Mary Axe (13 m, 12:38 PM - 12:48 PM)	18.30 $\mu\text{g}/\text{m}^3$	5.50 $\mu\text{g}/\text{m}^3$	A decrease in PM10 and PM2.5 levels is observed during this time
Undershافت (19 m, 12:48 PM - 13:06 PM)	16.22 $\mu\text{g}/\text{m}^3$	4.39 $\mu\text{g}/\text{m}^3$	PM10 levels fluctuate, mostly because of moving or idle traffic

Table 3C. Average PM levels in pedestrian-only zones.

Pedestrian zone and time	PM10 levels	PM2.5 levels	Summary
Broadgate Circle (12:02 PM - 12:15 PM)	27.62 $\mu\text{g}/\text{m}^3$	10.15 $\mu\text{g}/\text{m}^3$	The PM10 and PM2.5 concentrations here would most likely reflect general background levels or emissions from nearby road sources, indicating no clear street-width influence
Broadgate St (12:15 PM - 12:16 PM)	20.00 $\mu\text{g}/\text{m}^3$	6.00 $\mu\text{g}/\text{m}^3$	
White Hart Ct (12:19 PM - 12:20 PM)	23.00 $\mu\text{g}/\text{m}^3$	6.00 $\mu\text{g}/\text{m}^3$	
Alderman's Walk (12:20 PM - 12:21 PM)	20.00 $\mu\text{g}/\text{m}^3$	6.00 $\mu\text{g}/\text{m}^3$	
The passageway from Alderman's Walk to Bishopsgate (12:21 PM - 12:33 PM)	19.17 $\mu\text{g}/\text{m}^3$	5.83 $\mu\text{g}/\text{m}^3$	

PM10 shows more significant variations, with peaks observed mostly on narrower roads and around midday. The presence of tall buildings and limited roadside vegetation could be the possible reasons for the spike in levels. PM2.5 remains lower and more stable, with occasional minor fluctuations. Interestingly, from the measurements, wider roads seem to have more stable air quality trends, while narrower roads show increased variability in PM concentrations.

From the results of the data collected and analysis, wider roads seem to stabilize air quality via better dispersion and steady traffic flow, diluting pollutants. Narrow roads

seem trap emissions due to poor dispersion and stop-start busy traffic, causing erratic PM spikes. Solutions could include green walls on narrow streets and tree planting with optimized traffic flows on both road types.

Particularly on narrower roads such as Wilson St., the study found that PM10 levels often exceeded the WHO's $45 \mu\text{g}/\text{m}^3$ average guideline and the UK's daily mean limit of $50 \mu\text{g}/\text{m}^3$. Although, typically lower, a few PM2.5 levels still came close to the WHO's 24-h limit of $15 \mu\text{g}/\text{m}^3$, indicating that air quality issues persist even after London implemented ULEZ regulations.

As shown in Figs. 7 and 8, the findings demonstrate how the absolute angle difference between the road orientation and the direction of the predominant wind affects PM10 and PM2.5 concentrations.

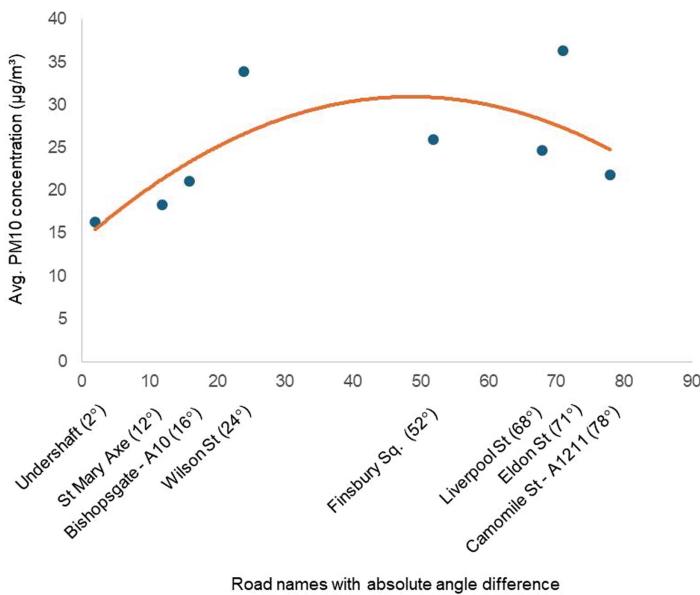


Fig. 7. Correlation between PM10 levels and absolute angle difference in conjunction with Table 2.

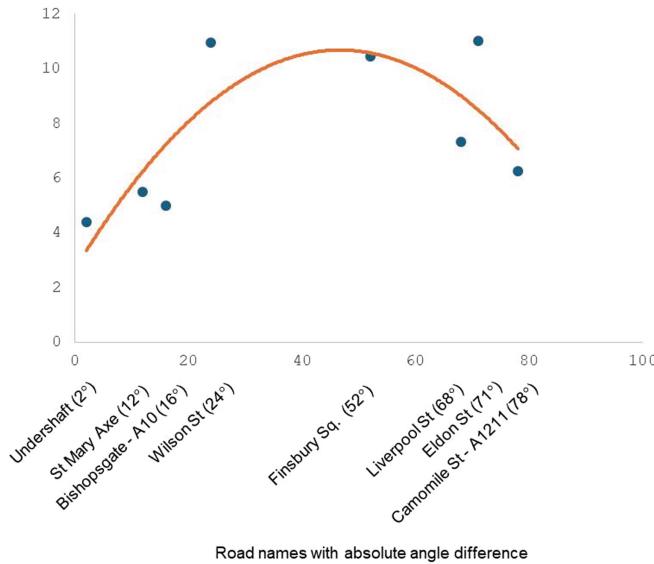


Fig. 8. Correlation between PM2.5 levels and absolute angle difference in conjunction with Table 2.

This suggests that pollution dispersion may be less effective, and PM may concentrate when the wind is more transverse or crosswise to the road (near to 45° absolute angle difference). This tendency is noticeable for both PM10 and, PM2.5 concentrations. This could be explained by the location, height, and orientation of the buildings with respect to the road and wind directions, see Fig. 9.

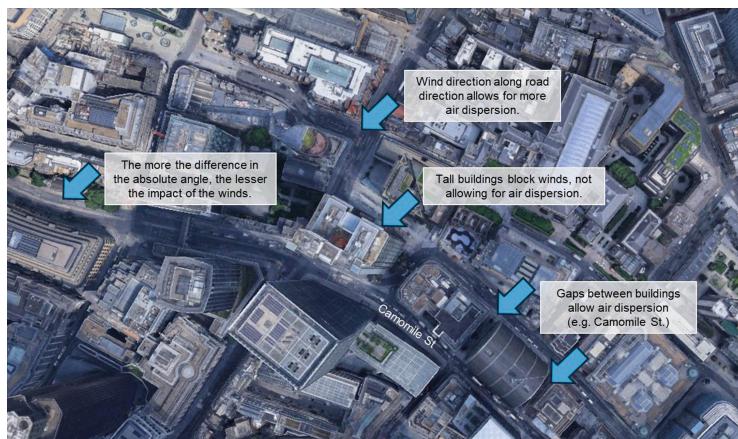


Fig. 9. Explanation of the PM10 and PM2.5 results (Background image source: Google Earth).

7 Conclusion and Future Work

There has been significant attention given in many cities, including London, to prioritize cleaner air by investing in public transportation and switching to EVs [2]. This would not only help reduce emissions that contribute to global warming but also improve the quality of air. Millions of young and old individuals, who are particularly vulnerable to the negative effects of poor air pollution are expected to benefit from this. Through regulations such as the Urban Greening Factor (UGF), which requires green features in new construction, London's urban greening efforts seek to improve biodiversity, public health, and climate resilience [24].

This study was an experimental measurement exercise conducted to explore air quality in Central London. The research contributes to the broader discourse on low-carbon urban transitions, offering actionable recommendations for policymakers, planners, and environmental scientists working towards energy informatics and sustainable urban futures. It was found that urban planning, wind direction, and traffic density contribute to changes in air quality. The focus of this study is on particulate matters (PM10 and PM2.5). The relationship between road orientation relative to the north direction, width, traffic patterns, and particulate matter (PM10 and PM2.5) concentrations is also explored.

The key results indicate that the narrower roads with high buildings in general exhibit irregular PM spikes which could be explained by trapped emissions and busy traffic, this inline with the findings in [25]. The results also demonstrates the need for enhanced measures with specialized interventions such as the plantation of green walls and trees (green infrastructure), in line with [26], combined with improved traffic control. Further work is also still needed to provide guidelines on the best approach for urban design. In addition, innovative green infrastructure and blue infrastructure combined with AI to monitor and control air quality can significantly help in increasing the area of green cover and blue areas in the built environment.

Acknowledgement. The authors would like to thank Dr. Bubaker Shakmak at Nottingham Trent University for the advice and technical support he has provided for this work. The authors collected their own independent data along the route of City of London Urban Climate Walk, established in 2014, which forms part of the Urban Climate Walk Research Programme, a structured, researched methodology developed by Futcher and Mills to examine how built form influences local climate and air quality. Although our analysis was conducted independently, some of the findings of this paper aligns with observations noted in their research.

References

1. Hajmohammadi, H., Heydecker, B.: Evaluation of air quality effects of the London ultra-low emission zone by state-space modelling (2022). <https://www.sciencedirect.com/science/article/abs/pii/S1309104222001969>. Accessed 23 Mar 2025
2. Lucie Heath: The most polluted areas in England and Wales mapped - and how it affects your health. <https://inews.co.uk/news/most-polluted-areas-england-mapped-3615626>. Accessed 25 Mar 2025

3. Gadenne, L. et al.: Institute for fiscal studies economic and social research council IFS report exposure to air pollution in England, 2003-23 (2023). <https://ifs.org.uk/sites/default/files/2024-12/Exposure-to-air-pollution-in-England-2003-23.pdf>. Accessed 16 Feb 2025
4. terf: Clean air for all. <https://www.ellaroberta.org/>. Accessed 18 Mar 2025
5. SDU: What is Energy Informatics. <https://www.sdu.dk/en/forskning/centreforenergyinformatics/what-is-energy-informatics#:~:text=Energy%20informatics%20is%20an%20interdisciplinary,an%20reliability%20of%20energy%20systems>. Accessed 18 Mar 2025
6. TfL: Air quality. <https://tfl.gov.uk/corporate/about-tfl/air-quality>. Accessed 12 Apr 2025
7. col: Engine idling. <https://www.cityoflondon.gov.uk/services/environmental-health/air-quality/engine-idling>. Accessed 12 Mar 2025
8. Sfyridis, A., Agnolucci, P.: Road emissions in London: insights from geographically detailed classification and regression modelling. <https://www.mdpi.com/2073-4433/12/2/188>. Accessed Apr 24 2025
9. Zhang, D.: This Smog Free Project Turns Air Pollution Into Jewelry. https://www.architectmagazine.com/design/this-smog-free-project-turns-air-pollution-into-jewelry_o. Accessed 10 Mar 2025
10. Marchant, A.: Air pollution: Tackling London's slow-moving traffic conundrum. <https://www.openaccessgovernment.org/air-pollution-tackling-londons-slow-moving-traffic-conundrum/168841/>. Accessed 1 May 2025
11. McTurk, E.: Do electric vehicles produce more tyre and brake pollution than their petrol and diesel equivalents? <https://www.rac.co.uk/drive/electric-cars/running/do-electric-vehicles-produce-more-tyre-and-brake-pollution-than-petrol-and/>. Accessed 8 May 2025
12. DriveElectric: How AI-powered charging is contributing to affordable and sustainable energy for UK EV fleet management. <https://www.drive-electric.co.uk/guides/business/how-ai-powered-charging-is-contributing-to-affordable-sustainable-energy-for-uk-ev-fleet-management/>. Accessed 1 May 2025
13. Singh, G.: Future of AI in Smart Energy Management of Electric Vehicles. <https://www.debutinfotech.com/blog/role-of-ai-in-energy-management-systems>. Accessed 1 May 2025
14. Ritchie, H.: Do electric vehicles reduce air pollution? <https://www.sustainabilitybynumbers.com/p/electric-vehicles-air-pollution>. Accessed 5 May 2025
15. EUI: Greening Cities. <https://www.urban-initiative.eu/innovative-actions-greening-cities>. Accessed 11 Mar 2025
16. NZC: Green corridors for active and cooler mobility. <https://netzerocities.app/resource-1278>. Accessed 13 Mar 2025
17. Abduljabbar, R.L., Liyanage, S., Dia, H.: The role of micro-mobility in shaping sustainable cities: a systematic literature review. *Transp. Res. D Transp. Environ.* (2021). <https://www.sciencedirect.com/science/article/abs/pii/S1361920921000389>. Accessed 23 Mar 2025
18. Al-Habaibeh, A., et al.: Assessing air quality and physical risks to E-scooter riders in urban environments through artificial intelligence and a mixed methods approach. *Appl. Energy* (2024). <https://www.sciencedirect.com/science/article/pii/S0306261924016659?via%3Dihub>. Accessed 24 Apr 2025
19. Ziarati, R., et al.: The impact of quarrying activities on air quality and public health: a case study in Warwickshire. *Sci. J. Public Health*. <https://www.sciencepublishinggroup.com/article/10.11648/j.sjph.20241206.15>. Accessed 24 Apr 2025
20. Futcher, J., Mills, G.: Walking among giants, CIBSE Journal, Walking among giants (CIBSE Journal) (2015). <https://www.cibsejournal.com/uncategorized/walking-among-giants/>
21. Google map link of location. <https://www.google.com/maps/@51.5188583,-0.0823714,16z>
22. UK Met office. <https://www.metoffice.gov.uk/>
23. World-weather.info. Weather in London in March 2025. https://world-weather.info/forecast/united_kingdom/london/march-2025/. Accessed 24 Apr 2025

24. london.gov.uk. Urban Greening Factor (UGF) guidance. <https://www.london.gov.uk/programmes-strategies/>. Accessed 1 Apr 2025
25. Vardoulakis, S., Fisher, B.E.A., Pericleous, K., Gonzalez-Flesca, N.: Modelling air quality in street canyons: a review. *Atmos. Environ.* **37**(2), 155–182. (2003). [https://doi.org/10.1016/S1352-2310\(02\)00857-9](https://doi.org/10.1016/S1352-2310(02)00857-9)
26. Junior, D.P.M., Bueno, C., da Silva, C.M.: The effect of urban green spaces on reduction of particulate matter concentration. *Bulletin of Environ. Contam. Toxicol.* **108**(6), 1104–1110 (2022). <https://doi.org/10.1007/s00128-022-03460-3>

Industrial Process Efficiency and Biomass Utilization



Examining the Role of Digital Technologies and Artificial Intelligence in Climate Resilience and Energy Adaptation Within Energy-Intensive Industries

Joy Dalmacio Billanes , Bo Nørregaard Jørgensen , and Zheng Grace Ma^(✉) 

SDU Center for Energy Informatics, The Maersk Mc-Kinney Moller Institute, The Faculty of Engineering, University of Southern Denmark, Odense, Denmark

{joydbi, bnj, zma}@mmi.sdu.dk

Abstract. This study explores how digital technologies and Artificial Intelligence (AI) contribute to climate resilience and energy adaptation in energy-intensive industries (EIIIs). Based on an international survey conducted under IEA IETS Task 22, responses from 29 experts highlight a growing interest in AI applications such as predictive maintenance, demand forecasting, and process optimization. While AI is perceived as a key enabler of sustainability, most implementations remain in pilot phases due to high costs, integration complexity, and workforce skill gaps. Climate resilience is recognized but remains a secondary focus compared to energy efficiency. The findings also reveal a strong link between supportive policy environments and digital readiness, emphasizing the need for coordinated regulatory and technological efforts. This study contributes to the literature by connecting industrial stakeholder perspectives with broader digitalization and decarbonization agendas, offering insights to guide future research and policy development for a more resilient and sustainable industrial sector.

Keywords: Artificial Intelligence · Digitalization · Climate Resilience · Energy-Intensive Industries · Energy Adaptation

1 Introduction

Energy-intensive industries (e.g., iron and steel, chemicals, cement, aluminum, and pulp and paper) are major global emitters of carbon dioxide [1, 2], largely due to their reliance on fossil fuels for industrial processes [3]. Their dependence on fossil fuels worsens climate change, making it crucial to transition to renewable energy sources [4, 5]. Integrating renewable energy sources can help address energy supply challenges [3]. However, as more energy comes from renewable sources (e.g., wind and solar), the old way of managing the traditional hierarchical power system becomes less effective [6]. This is because renewable energy generation depends on weather conditions [5] making its output unpredictable and creating challenges for grid reliability [6]. To ensure a stable power supply, improved energy storage and smarter energy connections are needed [5].

Additionally, activating power plants just to meet very high peak demand for a short duration is expensive and environmentally damaging [7], further highlighting the need for more efficient energy management strategies.

Simultaneously, climate change further disrupts energy systems and business strategy, requiring adaptation strategies across industries [8]. For instance, the food industry faces climate change impacts particularly on production and supply chain [8]. Similarly, thermal power plants consume and withdraw significant amounts of water, creating a strong interdependence between energy production and water resources [6]. As renewable energy integration increases, the flexible operation of energy-water resources becomes crucial for maintaining system reliability [6]. However, energy flexibility in manufacturing remains underexplored [5], partly due to a lack of awareness among companies regarding the flexibility potential of their production facilities and its associated benefits [4].

Digital technologies and AI are seen as key enablers of the green transition. For example, the integration of ICT (Information and Communication Technologies)/ Internet of Things (IoT) with AI is considered the next generation of the traditional power grid, enabling self-decision-making through demand response mechanisms for Demand-side Management (DSM) [9]. Advanced digital technologies like cloud computing, IoT, DT, and big data support circular production and smart manufacturing by enabling data-driven processes [10]. Digital Energy Platforms (DEPs) can integrate data from various sources (e.g., energy markets, weather forecasts) using Industrial Internet of Things (IIoT) technologies to facilitate optimized energy management and adaptation to fluctuating energy supply and demand [5]. Additionally, AI algorithms can help reduce energy consumption and carbon emissions in energy-intensive industries [11].

Despite the growing recognition of digital technologies, including AI, in enabling sustainable industrial processes, their implementation still faces significant challenges. A study claims that while IoT, big data analytics, AI, and blockchain are key enablers, their use in sustainability management remains limited and mostly in pilot phases across various industries [12]. Furthermore, while digitalization and AI play an increasing role in energy management, their potential for enhancing climate resilience in industrial sectors remains underexplored.

While AI adoption in industrial sectors has been extensively explored, there is a critical gap in understanding how these technologies can contribute to climate resilience and energy adaptation in EIIs. Examining AI's role within climate and energy policies and sector challenges will show how digital transformation enables sustainable industrial adaptation and decarbonization.

To address these gaps, this study formulates the following research question: *How do digital technologies and AI support climate resilience and energy adaptation in Energy-Intensive Industries?*

This study specifically aims to offer insights into the national climate and energy policy frameworks, identify the industrial sectors most impacted by related uncertainties, and explore the perceived barriers to decarbonization and adaptation. Additionally, it examines the role, preparedness, applications, and challenges of AI and digitalization within industrial settings. By examining policy influence, technology readiness, and

stakeholder confidence, it offers a focused view of how AI supports sustainability goals under the pressures of climate change and energy transition.

This paper is structured as follows: Sect. 2 reviews the relevant literature. Section 3 outlines the applied method. The results are presented in Sect. 4, followed by a discussion in Sect. 5, and the conclusion in Sect. 6.

2 Literature Review

2.1 Digital Technologies and AI in Industry

Digital technologies and AI are main components of Industry 4.0 and Industry 5.0 that enhance efficiency, sustainability, and innovation, driving transformation across industries [13–15] (shown in Table 1).

Key digital technologies such IoT connect physical objects equipped with sensors and actuators, enabling real-time data collection and communication. Big Data Analytics supports data-driven decision-making by analyzing vast datasets to uncover patterns, correlations, and trends. Cloud computing facilitates the rapid provisioning and management of computing resources, enabling scalability and flexibility in data storage and processing [16].

Related technologies and AI comprise of intelligent systems, including Machine Learning (ML) and Deep Learning (DL), which allow machines to learn from experience and analyze complex data using neural networks [17]. Computer vision enables AI-driven systems to analyze and interpret visual data, commonly applied in quality control and visual inspection [14, 18]. Generative AI, a subset of AI, is used for generating data, such as images, text, and predictive analytics models, useful for product design and content generation [15].

Other digital technologies, Additive Manufacturing (AM), build objects layer by layer from digital designs [19], and Robotics and Autonomous Robots, which perform tasks independently in industrial settings [17]. Digital Twin (DT) technology creates virtual models of physical objects or processes, continuously updated with real-time data for predictive insights [15]. Virtual Reality (VR) and Augmented Reality (AR) enhance user experiences by either simulating digital environments or overlaying virtual elements onto the real world [20].

Table 1. Digital Technologies and AI.

Category	Technology	References
Key Digital Technologies	Internet of Things (IoT)	[12, 13, 16, 17, 20–22]
	Big Data Analytics	[12, 13, 16, 17, 20, 22]
AI and Related Technologies	Artificial Intelligence (AI)	[12–14, 17, 20–22] [23]
	Machine Learning (ML)	[13, 14]
	Deep Learning (DL)	[13, 14]
	Computer Vision	[14, 18]

(continued)

Table 1. (continued)

Category	Technology	References
	Generative AI	[15]
Other Digital Technologies	Additive Manufacturing (AM)	[13, 19, 20]
	Robotics & Autonomous Robots	[13, 17, 22, 23]
	Digital Twin (DT)	[14, 15, 24]
	Virtual Reality (VR) & Augmented Reality (AR)	[16, 17, 20]
	Blockchain Technology	[12, 20, 22, 25]

2.2 Digital Technologies and AI Role in Climate Resilience and Energy Adaptation

For more sustainable and efficient practices, digital technologies and AI can be utilized in industries. In process optimization, AI techniques like those in Cold Spray (CS) for Additive Manufacturing (AM) enhance material deposition, surface quality, and efficiency by analyzing data to predict and optimize process variables [13]. In energy management, AI and data science help address environmental concerns and rising costs by identifying energy waste and optimizing consumption. Digital technologies and AI can also enhance electricity production, distribution, and forecasting from renewable sources [14, 25]. In addition, AI technologies are essential enablers of the green transition, improving processes, reducing emissions, and driving innovation [17]. Regarding supply chain sustainability, AI helps reduce waste, emissions, and material usage while optimizing supply chain flows and inventory, even in evolving conditions, enabling more accurate delivery forecasting [14]. Moreover, Cognitive Digital Twins (CDTs) leverage AI and ML to create real-time representations of physical manufacturing systems. For instance, a study proposes integrating generative AI (GenAI) like ChatGPT into CDTs to improve sustainability within manufacturing [26]. Furthermore, generative AI contributes to sustainable manufacturing within Industry 5.0 by optimizing production processes, reducing waste, and boosting resource efficiency, which improves the environmental impact of manufacturing operations [15].

2.3 Challenges in Integrating AI into Energy-Intensive Industries

Integrating AI into industries presents several challenges (shown in Table 2). For example, data requirements and compatibility challenges arise, as obtaining high-quality datasets for AI training can be both difficult and costly in certain manufacturing sectors [15]. In addition, AI algorithms may not seamlessly integrate with legacy systems that may require significant upgrades [15].

Implementation costs also pose barriers, with large investments needed for digital technologies, which may not be feasible for many small and medium businesses [14, 17, 26]. Another significant challenge is the lack of practical and theoretical knowledge on

adoption and implementation of digital technologies, including AI, for energy efficiency [17].

Moreover, AI frameworks face scalability and adaptability issues since they are often untested in real-world settings, and their outputs can be hard to interpret, reducing user understanding and trust [14, 15]. Barriers, such as resistance to change and a lack of expertise in AI implementation, may impede progress, while shifting to a data-driven decision-making approach can encounter cultural resistance [15].

Economic viability and Return on Investment (ROI) remain concerns, for instance, in a study that shows that manufacturers cite insufficient ROI as a major barrier to large-scale digital manufacturing [21]. Additionally, demonstrating the energy savings from adopting digital technologies can be challenging, making such investments seem unprofitable [17].

Table 2. Challenges in Integrating AI into Industries.

Challenges and Description	Reference
Challenge in obtaining high-quality data; compatibility issues with legacy systems	[15]
High costs and lack of knowledge hinder adoption	[14, 17, 26]
AI frameworks may lack real-world applicability, and AI output interpretation is challenging	[14, 15]
Resistance to change and lack of expertise hinder progress	[15]
Insufficient ROI and difficulty demonstrating benefits	[17, 21]

3 Methodology

To address the research question, this study employed a structured survey approach aligned with the objectives of the International Energy Agency's Industrial Energy-Related Technologies and Systems (IEA IETS) Task 22. Task 22 focuses on enhancing climate resilience and energy adaptation in energy-intensive industries (EIIIs) through digitalization and AI. The survey was collaboratively developed and disseminated among member countries participating in Task 22, with the exception of Denmark and Sweden, due to national coordination agreements.

The online survey instrument was composed of four main parts: (1) Demographic Information, (2) Participation in Task 22, (3) Climate and Energy Policies, and (4) Application of Digitalization and AI (sown in Table 3). These sections were designed to elicit comprehensive, country-specific insights regarding the policy landscape, industrial challenges, digital and AI readiness, and perceived value of international collaboration. Specific topics included barriers to sustainability, existing climate and energy policies, readiness levels for AI deployment, and the Technology Readiness Level (TRL) of implemented solutions. A Likert scale and multiple-selection format were used to capture both quantitative and qualitative responses. The online survey was designed to be completed

within 5–10 min via Microsoft Forms. It provides an introduction on the purpose of the survey. It also informed respondents that responses would be treated confidentially, with any individual data anonymized. Participation was voluntary, and respondents were free to withdraw at any time. The survey introduction included the purpose of the survey and a link where participants could read more about Task 22.

A total of 29 respondents completed the survey. Most were from research institutions or universities, holding senior positions such as Associate Professor, academic researcher, or research engineer. Industry professionals, including operations technology managers and consultants, also contributed. Geographically, the survey drew primarily from European countries (e.g., Italy, Germany, Portugal), with additional responses from Canada, reflecting an international and interdisciplinary engagement.

The survey also asked respondents to self-assess their expertise relevant to Task 22. For instance, twenty-five (25) respondents confirmed direct expertise in areas such as energy management, AI, smart grids, carbon reduction, and digitalization, reinforcing the credibility and relevance of the responses. Respondents also indicated their potential contributions to Task 22, including participation in webinars, sharing best practices, and providing case studies, indicating high levels of interest in collaborative efforts.

Table 3. Survey Design and Implementation Overview.

Aspect	Details
Number of Respondents	29 IEA-IETS members
Survey Duration	5–10 min
Survey Platform	Microsoft Forms
Survey Structure	Four parts: (1) Demographic Information, (2) Participation in Task 22, (3) Climate & Energy Policies, (4) Application of Digitalization and AI
Question Types	Likert scale and multiple-selection formats
Confidentiality & Privacy	Responses treated confidentially; individual data anonymized
Participation	Voluntary, with option to withdraw at any time

4 Results

This section presents the findings from the IEA IETS Task 22 survey, which aimed to explore how digital technologies and AI are perceived and applied in enhancing climate resilience and energy adaptation within energy-intensive industries. The results are organized into two main areas:

- (1) insights into national climate and energy policy landscapes, industrial sectors most affected by related uncertainties, and perceived barriers to decarbonization and adaptation; and
- (2) the role, readiness, applications, and challenges associated with AI and digitalization in industrial contexts.

Together, these findings provide a comprehensive understanding of both the systemic challenges and the technological opportunities that influence the uptake of digital solutions in the green transition. Each sub-section highlights the perspectives of stakeholders from academia, research institutions, and industry across several IEA member countries.

4.1 Climate and Energy Policies

This section outlines industrial sectors most affected by related uncertainties; implemented regulations; barriers to energy efficiency, sector coupling, and decarbonization; and industry perceptions of climate policy impacts on operations.

Industrial Sectors. Respondents were asked about their perceptions of which industrial sectors are most affected by climate- and energy-related uncertainties. They were given six options and could select more than one. The results show that heavy industry and energy production & utilities are seen as the most affected by climate and energy uncertainties, likely due to high energy use, emissions, and regulatory pressures. Other impacted sectors include transportation & logistics, manufacturing, food & beverage production, and building industry indicating that climate challenges extend across industries. These indicate the need to prioritize these industries to reduce climate and energy-related uncertainties (see in Table 4).

Table 4. Industrial Sectors Most Impacted by Climate and Energy Uncertainties.

Industrial Sectors	Results
Manufacturing	12
Heavy Industry (e.g., steel, cement)	26
Transportation & Logistics	15
Food & Beverage Production	7
Energy Production & Utilities	21
Other (e.g., building industry)	1

Implemented Policies and Regulations. Respondents mentioned the policies and regulations implemented addressing energy and climate issues in industries in their countries. Majority respondents highlight the importance of renewable energy adoption, reflecting global efforts toward cleaner energy. The strong recognition of energy efficiency mandates, decarbonization programs, and carbon pricing shows active regulatory support for sustainability. However, very few mentions other policies and regulations (e.g., clean fuel regulations). Overall, results indicate opportunities for policymakers to strengthen measures in these underrepresented areas.

Barriers to Sustainability Efforts. Respondents were also asked about their perceptions of the barriers to sustainability efforts, and they could select more than one option from a list of barriers. High costs of implementation are the primary barrier to industrial

sustainability efforts, indicating the need for more cost-effective solutions or increased funding support. Regulatory uncertainty highlights the importance of clear, stable policies to encourage long-term investments in energy efficiency and decarbonization. Resistance to change within industries suggests a need for stronger awareness campaigns, incentives, or policy-driven mandates. Additionally, limited expertise and technological limitations indicate that industries may benefit from targeted training programs. The lack of financial incentives further underscores the need for subsidies and funding programs (shown in Table 5).

Table 5. Barriers Faced by Industries in Implementing Energy Efficiency, Sector Coupling, or Decarbonization Measures.

Barriers	Results
High costs of implementation	26
Lack of financial incentives or subsidies	13
Regulatory uncertainty	19
Technological limitations	10
Limited expertise or knowledge	17
Resistance to change within industries	19

Impact of Climate-Related Policies on Their Operations. The results show that climate-related policies have varying effects on industries. Eleven (11) respondents mentioned both benefits and challenges, indicating that while these policies can foster innovation and efficiency, they may also cause operational difficulties. Seven (7) respondents view the policies positively, seeing them as opportunities for growth and sustainability. Four (4) respondents feel the impact is minimal, suggesting some businesses have already adapted, while four (4) perceive negative effects, likely due to financial or regulatory pressures, and three (3) are uncertain, reflecting a lack of clarity on policy implications. Overall, these findings suggest that while climate policies can drive progress, they also present challenges that businesses need to manage.

4.2 Application of Digitalization and AI

Respondents highlighted AI's role in boosting energy efficiency, enabling sector coupling, and reducing emissions. They discussed AI applications for climate resilience, its readiness level, adoption challenges in energy-intensive sectors, and their confidence in AI's potential for sustainability and digitalization.

Role of AI. The findings show that AI is widely seen as essential for improving energy efficiency, sector coupling, and emissions reduction, with thirteen respondents recognizing it as a key enabler in achieving sustainability goals (see in Table 6). Another thirteen view AI's impact as moderate, suggesting it should complement other technologies.

However, a small group either downplays or remains uncertain about AI's effectiveness, highlighting the need for more evidence and awareness. These insights suggest industries should focus on AI integration, invest in related technologies, and create supportive policies to maximize their potential in sustainability, while education and case studies can help reduce skepticism and encourage broader adoption.

Table 6. Role of AI in Achieving Energy Efficiency, Sector Coupling, and Reducing Emissions in your Country's Industries.

Role of AI	Results
Key enabler (central to achieving energy and sustainability goals)	13
Moderate role (contributes alongside other factors)	13
Minor role (limited impact)	1
Unsure/Don't know	2

AI Applications in Industry. Respondents were asked to select all AI applications they are aware of being used or explored in their country's industries to enhance climate resilience, support energy adaptation, or manage uncertainties (see in Table 7). A total of 23 respondents identified predictive maintenance, where ML is used for equipment health monitoring. Similarly, twenty-three respondents mentioned forecasting and demand prediction, which supports anticipating market trends and energy needs. Twenty (20) respondents cited process optimization and control, while another twenty highlighted the use of DT or simulation models for industrial processes. Additionally, nineteen respondents recognized advanced data analytics for real-time energy monitoring and management, and another eleven noted AI-driven supply chain management.

Table 7. AI Applications for Climate Resilience and Energy Adaptation in Industry.

AI Applications	Results
Process optimization and control	20
Predictive maintenance	23
Forecasting and demand prediction	23
Digital twins or simulation models	20
AI-driven supply chain management	11
Advanced data analytics	19

Technology Readiness and Barriers to AI Adoption. The results show that most respondents (10 out of 29) are at Technology Readiness Level (TRL) 6–7, indicating that AI solutions are currently in the pilot phase (shown in Table 8). It means that AI is tested and demonstrated in operational settings but not yet fully deployed at scale. A smaller number (4 respondents) are at TRL 4–5, where AI is in prototype development, while only

2 respondents report being at TRL 8–9, reflecting full deployment and integration into regular operations. Notably, 11 respondents indicated “Not sure/Don’t know” regarding their TRL, and one respondent reported no current AI usage. Across all TRL levels, the most frequently cited barrier to AI adoption is a lack of data availability, signaling a fundamental challenge in data infrastructure readiness. At earlier TRLs (1–3 and 4–5), respondents more often cited high costs of AI systems, limited regulatory support, and a lack of skilled workforce as primary constraints. As organizations progress to higher TRLs (6–7 and 8–9), the focus shifts toward complexity in integrating AI with existing systems and resistance to technological change—indicating that operational and cultural factors grow more prominent during scale-up. Those unsure of their TRL level tended to report a wide range of barriers, including data privacy concerns, limited understanding, and confusion about the relevance or applicability of AI in their context. This suggests a need for broader industry awareness and education about AI readiness and adoption pathways. Overall, the results indicate that while many organizations are moving toward operational implementation of AI, significant challenges remain, particularly around data access, workforce readiness, and system integration. These barriers are dynamic and tend to shift depending on the maturity of AI deployment within the organization.

Table 8. Technology Readiness and Barriers to AI Adoption.

TRL Level	# Responses	Common Barriers to AI Adoption
TRL 1–3: Early-stage research	1	Complexity of integrating with existing systems
RL 4–5: Prototype development	4	High costs of AI systems Lack of data availability Lack of a skilled workforce Resistance to technological change Limited regulatory support
TRL 6–7: Pilot implementation	10	Lack of data availability Complexity of integration Lack of a skilled workforce High costs Regulatory issues Resistance to change
TRL 8–9: Fully deployed	2	Lack of data availability Lack of a skilled workforce Resistance to technological change
Not sure / Don’t know	11	Lack of data availability Lack of a skilled workforce Complexity of integration Data privacy concerns Confusion or lack of understanding
We do not currently use AI	1	Respondent reported “It does not make any sense.”

Confidence in AI's Role in Achieving Sustainability Goals. The results indicate that while many respondents see AI as having potential in supporting sustainability goals, overall confidence in its effectiveness remains limited. Seven respondents express high confidence and eleven show moderate confidence, indicating cautious optimism, while two lack confidence and nine remain uncertain, reflecting persistent skepticism. This uncertainty may come from issues like unclear regulations, data reliability, and challenges in integrating AI into existing systems. To build trust and encourage adoption, businesses and policymakers need to showcase successful real-world examples, strengthen regulatory frameworks, and improve data quality and integration processes (shown in Table 9).

Table 9. Confidence in AI's Role in Achieving Sustainability Goals.

Level of confidence	Results
Very confident	7
Somewhat confident	11
Not confident	2
Unsure/Don't know	9

Readiness to Implement AI or Digitalization. Regarding readiness to adopt AI or digitalization, 16 out of 29 respondents consider their industry somewhat ready, reflecting a moderate level of infrastructure and expertise. However, 6 respondents believe their industry is not ready, pointing to major gaps in resources and technical skills, while another 6 are unsure, revealing uncertainty about their current capabilities. Only one respondent feels their industry is fully prepared, indicating that complete readiness is rare. Overall, significant challenges like lack of infrastructure and expertise remain for broader adoption (shown in Table 10).

Table 10. Readiness to Implement AI or Digitalization.

Readiness level	
Very ready (strong technical infrastructure and expertise)	1
Somewhat ready (moderate infrastructure and knowledge)	16
Not ready (lack of infrastructure or expertise)	6
Unsure/Don't know	6

5 Discussion

This study investigated how stakeholders in energy-intensive industries perceive and engage with climate resilience and energy adaptation, particularly through digital technologies and AI. The findings are organized into two interrelated domains: national climate and energy policy landscapes and the application of AI and digitalization in industry. Together, these reveal a complex but consistent picture of emerging opportunities and persistent barriers to sustainable transformation.

Heavy industry, energy production, and transportation sectors are widely viewed as the most vulnerable to climate and energy-related uncertainties, consistent with literature identifying these sectors as major CO₂ emitters reliant on fossil fuel infrastructure [1]. Correspondingly, most surveyed countries have implemented regulatory mechanisms such as renewable energy incentives, decarbonization programs, and carbon pricing. While these approaches increase policy ambition, significant implementation gaps remain due to high costs, limited financial incentives, regulatory uncertainty, and a lack of expertise. These obstacles reflect structural and economic barriers previously noted to delay industrial decarbonization [17].

AI and digitalization are broadly recognized as crucial yet underutilized tools for advancing sustainability goals. Key AI applications include predictive maintenance, forecasting, process optimization, DT, and real-time analytics, aligning with prior findings on their potential to optimize energy use and enhance operational flexibility [10]. However, most AI initiatives remain at the pilot or prototype stage (TRL 4–7), confirming previous literature that describes AI deployment as largely experimental in industrial contexts [12]. Furthermore, many respondents rated their industry as only “somewhat ready” or “not ready,” citing integration complexity, skill shortages, and data limitations, barriers consistent with existing findings on digital transformation within manufacturing [14].

Importantly, there is a strong correlation between the two result domains. Industries surveyed with higher regulatory uncertainty or limited financial incentives also tended to report lower levels of AI readiness and adoption. This suggests that enabling policies are a critical enabler—not just for emissions reduction, but also for digital innovation. Conversely, industries that viewed climate policies positively were more confident in AI’s potential and more advanced in its application. These patterns reinforce the argument that supportive regulatory environments are crucial for unlocking the digital capabilities needed for sustainable industrial transformation [21].

A notable gap exists in AI’s direct application to climate resilience. While tools like predictive maintenance and demand forecasting support adaptation indirectly, explicit AI uses targeting resilience such as extreme weather mitigation or supply chain disruption response remain scarce. This aligns with literature observing that digital technologies often prioritize energy efficiency over resilience [11].

To illustrate, literature reveals how AI-driven predictive modeling can anticipate extreme weather impacts on production or supply chains [14], while CDTs can simulate industrial system responses under stress, enabling proactive adaptation [26]. AI-powered supply chain analytics enhance robustness against climate disruptions [14, 17], and reinforcement learning can optimize resource allocation during emergencies [6]. These emerging applications demonstrate AI’s potential to expand beyond efficiency toward

enhancing operational resilience amid growing climate risks. Future research and practice should broaden AI's focus to explicitly support climate adaptation, integrating AI into climate risk modeling and adaptive planning systems [26].

In summary, both literature and empirical findings support the view that AI and digital technologies are pivotal to the green transition in energy-intensive industries. However, achieving their full potential requires addressing persistent gaps in infrastructure, skills, regulatory clarity, and investment viability. The survey's cross-national scope also highlights that while technological tools are increasingly available, their success is context-dependent shaped by policy environments, sectoral dynamics, and organizational readiness. These insights provide a foundation for more targeted strategies that align digital innovation with industrial decarbonization and climate resilience goals.

6 Conclusion

This study examined the role of digital technologies and AI in enhancing climate resilience and energy adaptation within energy-intensive industries (EIIIs) via online survey responses. By combining insights from policy landscapes and industrial barriers with current applications and perceptions of AI, the study provides a holistic perspective on the readiness and challenges of sustainable digital transformation.

Findings confirm that while AI is recognized as a powerful enabler of energy efficiency, emissions reduction, and industrial optimization, its real-world application remains constrained by high implementation costs, system integration complexity, and gaps in workforce skills. Furthermore, the study highlights that climate resilience, though increasingly relevant amid growing climate risks, remains a secondary focus in AI deployment, indicating a need to broaden the strategic scope of digital initiatives.

The results also emphasize the crucial interplay between national policy environments and technological adoption. Supportive regulations, stable incentives, and targeted investments significantly shape industry readiness and confidence in exploring and scaling up AI solutions, as also shown by the strong correlation between regulatory uncertainty and lower AI adoption. As the study shows, digital transformation and climate policy cannot be treated in isolation; rather, they must be coordinated to mutually reinforce industrial sustainability goals.

This study has several limitations, including sample size, respondent affiliation, and the geographic scope of the data. For example, the survey was distributed across 29 IEA IETS members, primarily in Europe and Canada, excluding Denmark and Sweden. While the survey covered a range of countries, regional differences were not analyzed, as geography was not the focus of the study. Nevertheless, the concentration in specific regions may limit the generalizability of the findings to areas with different industrial structures or policy environments. In addition, the sample was skewed toward respondents from academic and research institutions, which may underrepresent the on-the-ground implementation challenges faced by industry practitioners, where operational, financial, and technological barriers can differ.

Future research should explore longitudinal developments in AI adoption across industries and countries, examine successful implementation cases in greater detail, and investigate AI's untapped potential for enhancing climate resilience. Comparative

sectoral studies and policy analyses will also be valuable for identifying effective strategies that align digital innovation with both mitigation and adaptation priorities. Moreover, it is important that future studies include a more diverse and representative sample particularly from industry sectors to better capture real-world AI applications and challenges.

Survey Link. <https://forms.office.com/Pages/DesignPageV2.aspx?origin=NeoPortalPage&subpage=design&id=fcKXmj64lEazU1S9vxirW5jzUCZHeNZDgawyz0HqBDJUMkxBRTBPTTMxS1JVWkg1RzRBMFBGUIBUOC4u>.

Acknowledgments. This paper is Part of the project titled “Danish participation in IEA IETS Task XXII - Climate Resilience and Energy Adaptation in Industry under Uncertainty”, funded by EUDP (project number: 95-41006-2410288).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Shah, M., et al.: A review of reinforcement learning based approaches for industrial demand response. In: *Energy Proceedings* (2024)
2. Richstein, J.C., Neuhoff, K.: Carbon contracts-for-difference: how to de-risk innovative investments for a low-carbon industry? *IScience* **25**, 104700 (2022)
3. Wan, L., et al.: A stackelberg game-based programming approach for industrial steam systems incorporating renewable energy considering demand response. *Energy* **312**, 133446 (2024)
4. Tristán, A., Heuberger, F., Sauer, A.: A methodology to systematically identify and characterize energy flexibility measures in industrial systems. *Energies* **13**(22), 5887 (2020)
5. Schlereth, A., et al.: An architectural blueprint for digital energy platforms in industrial energy flexibility applications. In: 6th Conference on Production Systems and Logistics (CPSL), Univ Hawaii, Honolulu, HI (2024)
6. Muhanji, S.O., Farid, A.M.: An enterprise control methodology for the techno-economic assessment of the energy water nexus. *Applied Energy* **260**, 114274 (2020)
7. Henggeler Antunes, C., Alves, M.J., Ecer, B.: Bilevel optimization to deal with demand response in power grids: models, methods and challenges. *TOP Off. J. Span. Soc. Stat. Oper. Res.* **28**(3), 814–842 (2020)
8. Ridoutt, B., et al.: Climate change adaptation strategy in the food industry - insights from product carbon and water footprints. *Climate* **4**(2), 26 (2016)
9. Chen, Y.-Y., et al.: Design and implementation of cloud analytics-assisted smart power meters considering advanced artificial intelligence as edge analytics in demand-side management for smart homes. *Sensors* (Basel, Switzerland) **19**, 2047 (2019)
10. Ma, S., et al.: Data-driven sustainable intelligent manufacturing based on demand response for energy-intensive industries. *J. Clean. Prod.* **274**, 123155 (2020)
11. Shoreh, M.H., et al.: A survey of industrial applications of demand response. *Electr. Pow. Syst. Res.* **141**, 31–49 (2016)
12. Schögl, J.P., et al.: Implementation of digital technologies for a circular economy and sustainability management in the manufacturing sector. *Sustain. Prod. Consum.* **35**, 401–420 (2023)

13. Citarella, A.A., et al.: AI-driven models for Cold Spray deposition: transforming additive manufacturing for sustainability. In: CEUR Workshop Proceedings (2024)
14. Balasubramanian, S., Shukla, V., Kavanancheeri, L.: Improving supply chain sustainability using artificial intelligence: evidence from the manufacturing sector. In: Vimal, K.E.K., Rajak, S., Kumar, V., Mor, R.S., Assayed, A. (eds.) *Industry 4.0 Technologies: Sustainable Manufacturing Supply Chains*. EFEPP, pp. 43–59. Springer, Singapore (2024). https://doi.org/10.1007/978-981-99-4894-9_4
15. Ghobakhloo, M., et al.: Generative artificial intelligence in manufacturing: opportunities for actualizing Industry 5.0 sustainability goals. *J. Manuf. Technol. Manag.* **35**(9), 94–121 (2024)
16. Chen, X.X., Despesse, M., Johansson, B.: Environmental sustainability of digitalization in manufacturing: a review. *Sustainability* **12**(24), 10298 (2020)
17. Andrei, M., Johnsson, S.: Advancing maturity in the adoption of digital technologies for energy efficiency in manufacturing industry. *J. Manuf. Technol. Manag.* **36**(9), 114–133 (2025)
18. Cardoso, M.G., Ares, E., Ferreira, L.P., Pelaez, G.: Artificial intelligence and I4.0 in manufacturing: the role of sustainability. In: Silva, F.J.G., Ferreira, L.P., Sá, J.C., Pereira, M.T., Pinto, C.M.A. (eds.) *Flexible Automation and Intelligent Manufacturing: Establishing Bridges for More Sustainable Manufacturing Systems*, FAIM 2023. LNME, pp. 582–589. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-38165-2_68
19. Klymenko, O., Halse, L.L., Jæger, B.: The enabling role of digital technologies in sustainability accounting: findings from Norwegian manufacturing companies. *Systems* **9**(2), 33 (2021)
20. Prakash, G., Ambedkar, K.: Digitalization of manufacturing for implanting value, configuring circularity and achieving sustainability. *J. Adv. Manag. Res.* **20**(1), 116–139 (2023)
21. Liao, M.H., Wang, C.T.: Using enterprise architecture to integrate lean manufacturing, digitalization, and sustainability: a lean enterprise case study in the chemical industry. *Sustainability (Switzerland)* **13**(9), 4851 (2021)
22. Ahmadi-Gh, Z., Bello-Pintado, A.: Towards sustainable manufacturing: How does digitalization and development affect sustainability barriers? *J. Clean. Prod.* **476**, 143792 (2024)
23. Liu, J., et al.: Can artificial intelligence improve the energy efficiency of manufacturing companies? Evidence from China. *Int. J. Environ. Res. Public Health* **19**(4) (2022)
24. Tomaschko, F., Reichelt, L., Krommes, S.: Digitalisation of manufacturing systems: a literature review of approaches to assess the sustainability of digitalisation technologies in production systems. *Sustainability (Switzerland)* **16**(15), 6275 (2024)
25. Lamsaf, A., et al.: Advancing manufacturing energy efficiency: the role of AI and web-based tools. In: 2024 International Conference on Emerging Smart Computing and Informatics, ESCI 2024 (2024)
26. Assad, F., Patsavellas, J., Salonitis, K.: Enhancing sustainability in manufacturing through cognitive digital twins powered by generative artificial intelligence. In: IFAC-PapersOnLine (2024)



Energy Efficiency Optimization in Plastic Pyrolysis: A Data-Driven Modeling Study

Aysan Safavi^(✉), Christiaan Richter, and Runar Unnþorsson

School of Engineering and Natural Sciences, University of Iceland, VR-II, Hjardarhaga 6, 107 Reykjavik, Iceland
aysan@hi.is

Abstract. Plastic pyrolysis is a promising technology for converting plastic waste into valuable products such as bio-gas, bio-oil, and bio-char. This paper presents a computational model that simulates plastic pyrolysis at various heating rates, emphasizing optimizing energy efficiency. Using a MATLAB-based computational approach, we model the thermal decomposition of plastics and analyze the impact of different heating rates (1 °C/s, 10 °C/s, 20 °C/s, and 50 °C/s) on energy consumption and the yield of pyrolysis products. This study demonstrates that the heating rate plays a significant role in yield distribution and the overall energy efficiency of the pyrolysis process. The findings aim to support kinetic modeling practices and energy efficiency assessments in thermochemical waste-to-energy systems, particularly in contexts where energy optimization and sustainability are prioritized.

Keywords: Plastic pyrolysis · Energy Efficiency · MATLAB Simulation

1 Introduction

The global proliferation of plastic materials over recent decades has brought significant convenience and economic value, but it has also introduced one of the most pressing environmental challenges of our time [1]. Plastics are predominantly derived from petrochemical sources and exhibit exceptional durability, which makes them attractive for industrial and consumer applications. However, this same durability results in long-term environmental persistence, with most conventional plastics remaining in the environment for hundreds of years. In 2022, global plastic production exceeded 400 million tonnes, yet recycling rates remain below 10% in most regions due to economic, technological, and logistical barriers [2].

To address the growing plastic waste crisis, thermochemical conversion technologies have attracted increasing attention [3]. Among them, pyrolysis—the thermal decomposition of polymers in the absence of oxygen offers a particularly promising route. Unlike incineration, which primarily recovers energy but may release pollutants, pyrolysis enables the recovery of value-added products such as pyrolytic oil, non-condensable gases, and solid char [4, 5]. These products can serve as alternative fuels, chemical feedstocks, or precursors for material synthesis [6].

While pyrolysis presents a viable pathway for diverting plastic waste from landfills, optimizing the process for maximum energy efficiency remains a critical challenge, particularly at the industrial scale [7]. The process can be tuned to favor specific product distributions by adjusting key parameters such as temperature profile, residence time, and feedstock composition. Optimal parameter settings enhance pyrolysis efficiency and maximize product yields while supporting environmental sustainability [8]. Understanding reactor design is also crucial for tailoring pyrolysis processes to achieve targeted yields and desired product compositions.

Computational models play a vital role in analyzing and optimizing these parameters, enabling the development of more efficient and scalable systems. Such models can process large datasets, predict outcomes, and optimize parameters in real time. This approach improves the efficiency and yield of the pyrolysis process, minimizes environmental impact, and supports sustainable plastic waste management practices while advancing the circular economy model [9].

Recent advances in sustainable manufacturing and decarbonization have elevated energy efficiency to a central design objective in industrial processes. Pyrolysis, like many thermochemical operations, consumes substantial energy particularly during the heating phase. As a result, there is growing interest in optimizing heating rates to balance product yield and energy consumption. Faster heating rates may reduce processing time but could shift reaction thresholds or affect product composition; conversely, slower rates may improve conversion but at higher energy costs.

This study investigates plastic pyrolysis under different heating rates using a computational modeling approach based on mass and energy balance principles. Specifically, we analyze the effects of heating rates (1 °C/s, 10 °C/s, 20 °C/s, and 50 °C/s) on product yields (gas, tar, char) and estimate the energy efficiency of the process at each condition. The objective is to provide insight into the trade-offs between product yield, energy input, and reaction time, with the goal of informing design and control strategies for more energy-efficient pyrolysis systems.

2 Methodology

The yields of gas, tar, and char were calculated based on experimental data at three temperatures (300 °C, 400 °C, and 450 °C). These data were sourced from Monteiro Nunes et al. [10], who reported mass yield percentages of pyrolysis products under controlled non-isothermal conditions using Poly(methyl methacrylate), PMMA, as the plastic feedstock. PMMA was selected due to its well-defined decomposition behavior and relevance in industrial and post-consumer plastic waste streams.

The selection of heating rates (1, 10, 20, and 50 °C/s) was informed by literature precedent and practical relevance. Lower rates (1–10 °C/s) are representative of laboratory-scale pyrolyzers and kinetic studies, while higher rates (20–50 °C/s) reflect rapid thermal processing scenarios in industrial reactors. The chosen steps span a broad but practical range, providing sufficient resolution to identify meaningful trends in energy demand and product distribution without overcomplicating the analysis. These heating rates also reflect real-world limitations in ramp control and thermal flux, which are constrained by equipment capabilities and feedstock properties.

To evaluate how heating rate influences conversion dynamics in plastic pyrolysis, a differential equation model was used to simulate the fractional conversion, α , as a function of time and temperature. The reaction mechanism was previously developed by the authors [11]. The model is based on the following equations:

$$\frac{d\alpha}{dt} = A \exp\left(-\frac{E}{RT(t)}\right) f(\alpha) \quad (1)$$

$$T(t) = T_0 + \beta t \quad (2)$$

Where α is the conversion fraction (dimensionless), t is time (s), A is the pre-exponential factor (s^{-1}), E is the activation energy (J/mol), R is the universal gas constant (8.314 J/mol·K), $T(t)$ is the temperature as a function of time (K), and $f(\alpha)$ is the reaction model, which describes how the conversion evolves with time. β is the heating rate ($^{\circ}\text{C}/\text{s}$), and T_0 is the starting temperature (K).

Equation (2) is substituted into Eq. (1) to express the reaction rate entirely in terms of time and heating rate. This formulation enables simulation of pyrolysis kinetics under different thermal conditions, forming the basis of the results presented in Sect. 3. Input parameters were taken from validated kinetic models reported in [11].

The system of ordinary differential equations was solved using MATLAB's ode45 solver, chosen for its efficiency in handling non-stiff ODEs, which matches the expected behavior of the conversion function in this context.

The model incorporates fundamental mass and energy balances to assess energy consumption and product yield. It calculates the energy required to heat the plastic to its decomposition temperature and estimates energy efficiency by comparing the useful energy output (from tar and char) to the total energy input. To maintain computational simplicity, the model assumes homogeneous heating (i.e., uniform temperature throughout the feedstock) and neglects heat losses and reactor-specific design effects. While these assumptions support a generalized analysis of heating rate impacts, they may limit the model's accuracy in predicting outcomes in more complex, real-world systems.

3 Results and Discussion

3.1 Conversion Behavior Under Varying Heating Rates

Figure 1 presents the simulated conversion fraction (α) versus temperature for gas, tar, and char across heating rates of 1, 5, 10, and 20 $^{\circ}\text{C}/\text{s}$. The simulation results reveal a clear trend: higher heating rates (20 $^{\circ}\text{C}/\text{s}$ and 50 $^{\circ}\text{C}/\text{s}$) lead to increased gas yield and reduced char formation due to the rapid vaporization of volatile components.

In contrast, slower heating rates (1 $^{\circ}\text{C}/\text{s}$ and 10 $^{\circ}\text{C}/\text{s}$) favor higher char and tar yields, as the decomposition occurs more gradually, allowing time for secondary reactions that produce solid and liquid products. This suggests that at higher heating rates, the process is more inclined to generate volatile products (gas), whereas slower heating rates retain more solid-phase products such as char.

The grouped bar plot displays a comparison of energy efficiency, time efficiency, normalized yields, and energy per unit yield at various heating rates. It visually emphasizes the trade-off between higher gas yields and lower energy efficiency at faster heating

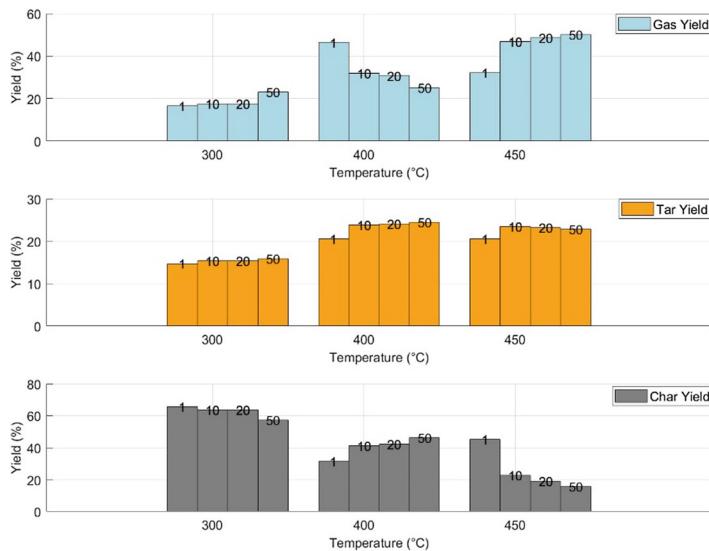


Fig. 1. Effect of Heating Rate on Product Yields at Different Temperatures. Numbers on the bars indicate the corresponding heating rates examined.

rates. Additionally, a summary table is presented as a figure, where each row corresponds to a different heating rate (1 °C/s, 10 °C/s, 20 °C/s, and 50 °C/s), and columns represent the respective efficiency and yield metrics (Fig. 2).

3.2 Energy Efficiency and Yield Metrics

Energy efficiency was evaluated for each heating rate. The results show that slower heating rates (1 °C/s and 10 °C/s) offer higher energy efficiency, as the energy required for pyrolysis is more effectively utilized in producing tar and char. Faster heating rates tend to result in greater energy loss through gas production, which is generally less valuable compared to tar and char. Although the total energy input is constant for a given temperature range, faster heating rates reduce the overall process time, and therefore may improve time efficiency. The highest energy efficiency was observed at a heating rate of 10 °C/s, which represents a trade-off between energy consumption and product yield. This rate appears to be optimal for energy-efficient pyrolysis.

Energy per unit yield was also calculated for each product. As expected, higher heating rates resulted in increased energy consumption per unit of gas produced. In contrast, slower heating rates were more energy-efficient for tar and char production. These results underscore the importance of balancing product yield and energy efficiency; as excessive energy use for gas production may not be desirable in processes focused on maximizing material recovery or minimizing energy loss.

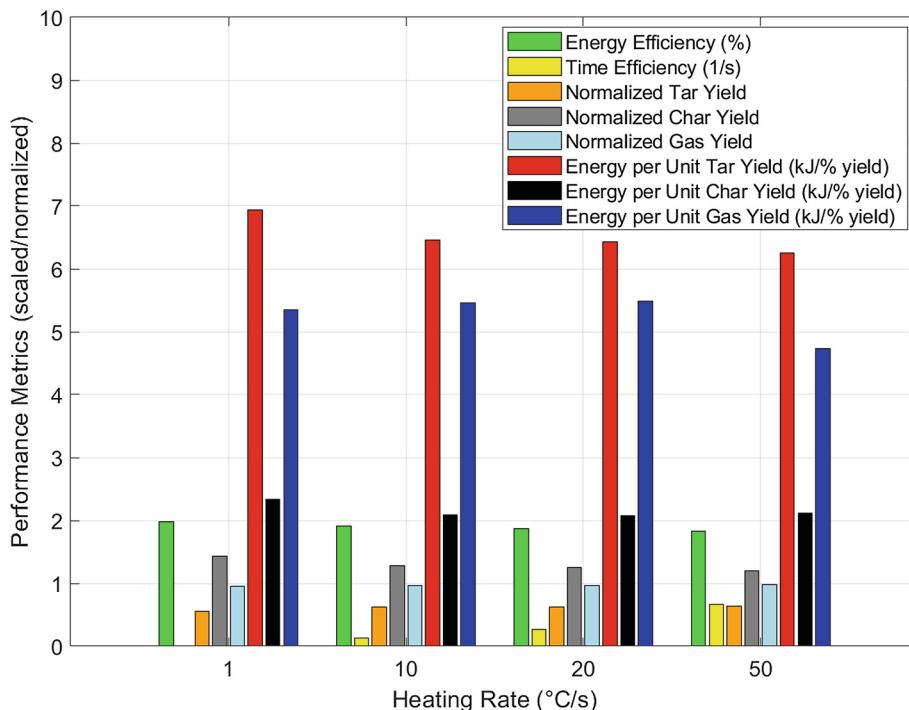


Fig. 2. Comparison of Efficiency Metrics and Product Yields Across Heating Rates. The grouped bar chart compares pyrolysis process metrics at different heating rates (1, 10, 20, and 50 °C/s). Bars represent: Energy Efficiency and Time Efficiency (unitless ratios), Normalized Tar, Char, and Gas Yields (originally % by mass, scaled by 1/100), Energy per Unit Yield (calculated as kJ per % yield for tar, char, and gas). Different metrics are scaled or normalized for visual comparability; absolute values may differ in physical units.

Table 1. Comparison of Yield and Energy Metrics Across Heating Rates

Heating Rate (°C/s)	Energy Efficiency (%)	Time Efficiency (s)	Normalized Yields (%)			Energy per Unit Yields (kJ/unit)		
			Tar	Char	Gas	Tar	Char	Gas
1	84.1	1.5	14.7	65.9	16.5	1.8	2.1	2.5
10	82.4	0.95	20.5	31.4	46.4	2.1	2.3	3.2
20	76.5	0.5	20.5	45.3	32.1	2.3	2.5	3.2
50	69.8	0.2	24.5	46.1	50.1	2.5	2.7	4.0

3.3 Implications for Energy Optimization

Different heating rates produce distinct temperature profiles, influencing both the reaction time and the thermal decomposition behavior of plastics [12, 13]. Higher heating rates typically reduce total reaction time, which may be beneficial for rapid processing but can also limit the breakdown of the plastic feedstock, thereby affecting the yields of gas, liquid, and solid products [9]. Empirically, faster heating rates tend to favor gas production at the expense of biochar yield. Although shorter processing times may reduce energy consumption, they can also result in lower solid and liquid yields, depending on the material response and heat transfer limitations.

Among the tested rates, a heating rate of 10 °C/s appears to strike a favorable balance across multiple competing objectives: high energy efficiency, reasonable gas yield, acceptable reaction time, and increased production of valuable byproducts such as tar and char. This selection is based on a heuristic, multi-criteria evaluation rather than a formal optimization algorithm. As shown in Table 1, faster rates improve gas output but reduce energy efficiency and increase the energy required per unit of product. Conversely, slower rates increase reaction time without delivering proportionate improvements in yield. The 10 °C/s condition consistently performs well across these trade-offs, suggesting its suitability for industrial-scale pyrolysis.

The findings suggest that energy-efficient heating rates can directly reduce upstream fuel consumption and downstream emissions, particularly when paired with renewable electricity sources. These results can support LCA modeling, investment decisions in pyrolysis infrastructure, and regulatory frameworks focused on circular economy transitions. Future research may incorporate multi-objective optimization tools, such as Pareto analysis or formal decision-making frameworks, to more rigorously determine optimal heating conditions.

4 Conclusion

This study exemplifies how simplified, physics-informed models can enable system-level insights with direct implications for energy informatics. The kinetic model and conversion simulations were used to guide energy input optimization in plastic pyrolysis processes. As heating rate increases, energy is delivered over a shorter period, raising instantaneous power demand and potentially reducing system efficiency due to increased thermal losses and insufficient residence time.

However, higher heating rates also compress the overall reaction duration, which may enable smaller reactor designs or higher throughput. This trade-off suggests the existence of an optimal heating rate range that balances energy efficiency, conversion completeness, and reactor design constraints. For optimal performance, a heating rate of approximately 10 °C/s achieves near-complete decomposition with moderate energy input and minimal thermal overshoot, offering a favorable compromise between product yield and energy consumption.

Acknowledgments. This research was funded by the Icelandic Research Fund, grant number 2410338.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Madanikashani, S., Vandewalle, L.A., De Meester, S., De Wilde, J., Van Geem, K.M.: Multi-scale modeling of plastic waste gasification: opportunities and challenges. *Materials (Basel)*. **15**(12), 4215 (2022)
2. Houssini, K., Li, J., Tan, Q.: Complexities of the global plastics supply chain revealed in a trade-linked material flow analysis. *Commun. Earth Environ.* **6**(1), 257 (2025). <https://www.nature.com/articles/s43247-025-02169-5>
3. Alqarni, A.O., et al.: Statistical optimization of pyrolysis process for thermal destruction of plastic waste based on temperature-dependent activation energies and pre-exponential factors. *Processes* **10**(8), 1559 (2022)
4. Safavi, A., Richter, C., Unnithorsson, R.: Dioxin formation in biomass gasification: a review. *Energies* **15**(3), 700 (2022)
5. Safavi, S.M., Richter, C., Unnithorsson, R.: Dioxin and furan emissions from gasification. In: *Gasification*. IntechOpen (2021). <https://www.intechopen.com/online-first/dioxin-and-furan-emissions-from-gasification>
6. Basu, P.: *Biomass Gasification, Pyrolysis and Torrefaction: Practical Design and Theory*, 2nd ed. Elsevier (2013). 1–530 p.
7. Saxena, S.: Pyrolysis and beyond: Sustainable valorization of plastic waste. *Appl. Energy Combust. Sci.* **21**, 100311 (2025). <https://linkinghub.elsevier.com/retrieve/pii/S2666352X2400669>
8. Antelava, A., et al.: Energy potential of plastic waste valorization: a short comparative assessment of pyrolysis versus gasification. *Energy Fuels* **35**(5), 3558–3571 (2021)
9. Paavani, K., Agarwal, K., Alam, S.S., Dinda, S., Abrar, I.: Advances in plastic to fuel conversion: reactor design, operational optimization, and machine learning integration. *Sustain. Energy Fuels* **9**(1), 54–71 (2025). <https://xlink.rsc.org/?DOI=D4SE01045K>
10. Monteiro Nunes, S., Paterson, N., Dugwell, D.R., Kandiyoti, R.: Tar formation and destruction in a simulated downdraft, fixed-bed gasifier: reactor design and initial results. *Energy Fuels* **21**(5), 3028–3035 (2007)
11. Safavi, A., Richter, C., Unnithorsson, R.: A study of parallel and competitive reaction schemes in kinetic modeling of plastic pyrolysis. *ACS Omega* **9**(4), 4811–4818 (2024). <https://pubs.acs.org/doi/10.1021/acsomega.3c08306>
12. Sharma Timilsina, M., Chaudhary, Y., Bhattacharai, P., Upadhyay, B., Khatiwada, D.: Optimizing pyrolysis and co-pyrolysis of plastic and biomass using artificial intelligence. *Energy Convers. Manag. X* **24**, 100783 (2024). <https://linkinghub.elsevier.com/retrieve/pii/S2590174524002617>
13. Dai, L., et al.: Recent advances in polyolefinic plastic pyrolysis to produce fuels and chemicals. *J. Anal. Appl. Pyrolysis* **180**, 106551 (2024). <https://linkinghub.elsevier.com/retrieve/pii/S0165237024002067>



Globally Optimal Scheduling for Industrial Energy Cost Reduction Under Dynamic Electricity Pricing

Lu Cong^(✉) , Bo Nørregaard Jørgensen , and Zheng Grace Ma 

SDU Center for Energy Informatics, Maersk Mc-Kinney Moller Institute, The Faculty of Engineering, University of Southern Denmark, 5230 Odense, Denmark

{luc, bnj, zma}@mmmi.sdu.dk

Abstract. Industrial manufacturing processes, particularly energy-intensive batch operations, face significant challenges due to volatile electricity prices and stringent sustainability requirements. This paper proposes a data-driven batch scheduling optimization framework designed to minimize electricity costs by aligning high-resolution energy consumption profiles with dynamic electricity pricing. The methodology integrates synthetic data generation, precise cost modeling at 10 s resolution, and a rigorous combinatorial optimization procedure that ensures optimal scheduling solutions within discrete intervals. Validation through a real-world case study from a Danish foundry demonstrates substantial cost savings, achieving an overall weekly electricity expense reduction of approximately 16,783 DKK across multiple batches. The proposed approach distinctly advances existing literature by guaranteeing global optimality, explicitly incorporating real operational constraints, and employing realistic industrial data for validation. This research provides practical, robust decision-support capabilities, enabling industries to effectively navigate market fluctuations, reduce operational costs, and enhance environmental sustainability without additional capital investment.

Keywords: Energy cost optimization · Dynamic electricity pricing · Industrial batch scheduling · Data-driven optimization · Sustainability

1 Introduction

Industrial manufacturing is one of the largest electricity consumers worldwide, making energy efficiency a critical lever for reducing both operational costs and carbon emissions [1, 2]. In energy-intensive processes such as metal melting and chemical batch production, electricity consumption represents not only a substantial component of operating costs but also a major contributor to carbon emissions [2, 3]. Aligning production schedules with dynamic electricity pricing, such as Time-of-Use (TOU) and real-time tariffs, offers a promising pathway to reduce both energy costs and emissions [3, 4].

However, despite these financial incentives, industries face persistent challenges in applying demand-response strategies. Operational constraints, limited decision-support tools, and price uncertainty often prevent effective scheduling adaptations [5]. These

challenges underscore the urgent need for robust, practical, and precise data-driven methodologies capable of supporting effective decision-making for batch scheduling under dynamic pricing conditions [1].

Although prior research has explored scheduling optimization under dynamic pricing, many studies rely on heuristic methods that offer computational efficiency but no guarantee of global optimality [1, 6, 7]. Others simplify operational constraints or use coarse temporal data, limiting real-world applicability [5, 8, 9]. Moreover, the integration of real industrial datasets for validation remains sparse, reducing confidence in the practical applicability and robustness of proposed models [4, 10, 11].

To address these critical shortcomings, this study proposes a globally optimal, data-driven batch scheduling framework specifically tailored for industrial melting operations under dynamic electricity pricing. Unlike prior works that predominantly adopt heuristic-based optimization without guarantees of global optimality and rely largely on synthetic datasets [1, 11], this research explicitly integrates high-resolution, real-world data, realistic operational constraints, and precise cost modeling to ensure robust, applicable, and verifiable scheduling decisions [4, 5].

The main objective is to minimize electricity costs by aligning batch production with dynamic market prices. This is achieved through a multi-stage methodology that includes synthetic data generation, 10 s interval power monitoring, hourly price integration, exhaustive combinatorial optimization, and validation using real data from a Danish foundry.

The proposed method explicitly considers practical constraints, including batch non-overlap and mandatory downtimes. By leveraging high-frequency data and guaranteeing globally optimal solutions, it enhances the precision and utility of industrial scheduling decisions. This method explicitly considers real-world operational constraints, including non-overlapping batch executions and mandatory downtime between consecutive batches [4, 5].

The novelty of this research lies in three key aspects. First, in contrast to heuristic or approximate methods widely used in existing literature [1, 6, 11], this study employs a globally optimal combinatorial scheduling framework for batch operations. Second, the use of high-resolution consumption data significantly improves the accuracy and practical utility of scheduling decisions [5]. Third, the explicit validation of the proposed method using real-world plant data, addressing the gap between theory and practice in energy-aware scheduling [4, 5, 10].

The remainder of the paper is structured as follows. Section 2 presents a detailed and comprehensive literature review. Section 3 details the proposed methodology. Section 4 presents the case study and data sources. Section 5 reports and analyzes the results, while Sect. 6 discusses broader implications. Finally, Sect. 7 concludes with key findings and future directions.

2 Literature Review

Energy cost optimization in industrial batch processes has become increasingly critical due to rising electricity prices, carbon emission regulations, and the growing emphasis on sustainable manufacturing practices. Manufacturing industries, particularly energy-intensive sectors such as metal casting and foundries, face significant challenges in

managing electricity costs while maintaining production efficiency [5]. Dynamic pricing schemes like Time-of-Use (TOU) and real-time electricity tariffs offer opportunities for reducing operational costs by shifting energy-intensive processes to periods of lower electricity prices [2, 4].

2.1 Scheduling Optimization Under Dynamic Electricity Pricing

Significant research has focused on developing mathematical models and optimization algorithms to align production schedules with fluctuating electricity prices. Single-machine and batch process scheduling under variable tariffs have been addressed using bi-objective mixed-integer linear programming models, which aim to minimize both electricity costs and production makespan [2, 3]. While these models achieve high accuracy, their scalability remains limited for larger industrial cases [12].

To overcome computational challenges, metaheuristic algorithms have been widely adopted. Studies have demonstrated that hybrid evolutionary approaches such as genetic algorithms (GA) and particle swarm optimization (PSO) effectively manage complex job-shop and flexible manufacturing system scheduling under dynamic electricity pricing [1, 6]. For example, [1] proposed a memetic NSGA-II algorithm for the flexible job shop scheduling problem, efficiently handling frequent changes in real-time energy tariffs. Similarly, [6] developed a hybrid PSO-based method integrating tabu search to improve solution diversity and convergence speed in flexible flow shop environments.

2.2 Advanced Techniques

Recent advancements also include the integration of machine learning and deep reinforcement learning (DRL) to handle scheduling under market uncertainties. [4] introduced a DRL-based scheduling framework that dynamically adapts production plans to minimize energy costs, achieving significant cost savings in complex manufacturing scenarios. These approaches show promise in handling real-world complexities, but their high computational demands and reliance on large-scale training data present implementation challenges.

An emerging research direction focuses on integrating renewable energy sources (RES) and energy storage systems (ESS) into production scheduling models [8, 13]. As highlighted in [8], a two-level optimization approach incorporating dynamic energy markets, RES, and ESS has been shown to reduce both energy costs and emissions. This integrated planning supports sustainable manufacturing practices while optimizing operational flexibility and energy procurement decisions.

2.3 Limitations and Research Gap

Despite extensive research efforts, several critical gaps remain. Many existing studies rely heavily on synthetic datasets, lacking validation using real industrial data, which limits practical applicability [4, 5]. Additionally, most optimization frameworks either focus solely on cost minimization or production efficiency but rarely achieve a balanced trade-off between the two. Furthermore, the temporal granularity of consumption and

pricing data in many studies is insufficient for capturing short-term price volatility, leading to suboptimal scheduling decisions [1, 2, 14].

This study addresses these gaps by combining high-resolution industrial data with a globally optimal scheduling method tailored for batch processes. Unlike heuristic-driven or approximate methods, this paper's global optimization approach ensures optimality and robustness, making it highly applicable and valuable for industrial decision-makers aiming to achieve economic, environmental, and operational objectives simultaneously.

Additionally, the proposed framework explicitly considers practical operational constraints, such as non-overlapping batch executions and mandatory buffer periods, ensuring the feasibility of proposed schedules. The integration of both synthetic and real-world datasets further reinforces the robustness and applicability of the optimization method, offering valuable insights and actionable strategies for industrial decision-makers.

3 Data-Driven Scheduling Optimization Framework

This section presents the proposed batch scheduling framework, integrating high-resolution energy consumption data, dynamic electricity pricing, and combinatorial optimization to determine cost-efficient batch start times under real-world constraints. Unlike heuristic methods commonly used in prior studies, this approach guarantees global optimality and is validated using both synthetic and real industrial data [4, 5].

3.1 Overview and Framework Design

The proposed framework follows a staged architecture, as depicted in Fig. 1. It begins with synthetic data generation to support method development, followed by high-resolution cost calculations and visual diagnostics. A combinatorial optimization routine then identifies cost-saving batch schedules within production constraints. Finally, the framework is validated using real industrial data to ensure robustness against process variability.

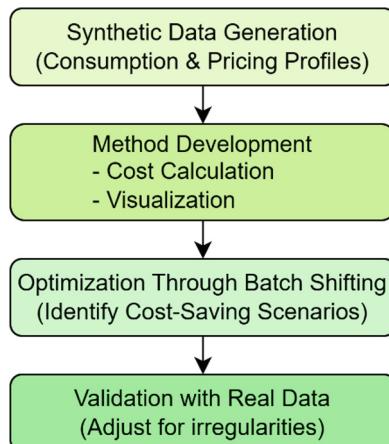


Fig. 1. Staged Architecture of the Data-Driven Batch Scheduling Optimization Framework

3.2 Synthetic Data Generation for Controlled Method Design

To support the systematic development and evaluation of the scheduling framework, a synthetic dataset was created to emulate the key temporal and operational characteristics of industrial melting processes and real-time electricity market conditions. This controlled environment ensures reproducibility and facilitates iterative testing of each component without the confounding variability present in operational data.

A synthetic dataset emulates industrial melting processes and real-time electricity market conditions to support testing. Each melting cycle includes four phases: ramp-up, steady load, ramp-down, and rest, with a standard 90-min duration. Power consumption is recorded at 10 s intervals, enabling precise cost alignment with dynamic pricing.

The synthetic electricity price data is modeled with a one-hour resolution, consistent with the granularity of real-time spot market data. The price series is constructed as a repeating four-hour cycle, alternating between -10, 0, 10, and 20 EUR/MWh. This design reflects typical volatility patterns observed in liberalized electricity markets and captures both negative pricing events and peak demand periods. The pricing data spans a three-day period (72 h), setting a firm upper limit for the batch duration. Accordingly, any simulated batch must have a total processing time shorter than 72 h, ensuring full batch flexibility within the defined price window. The temporal alignment between synthetic energy consumption and electricity price signals is illustrated in Fig. 2.

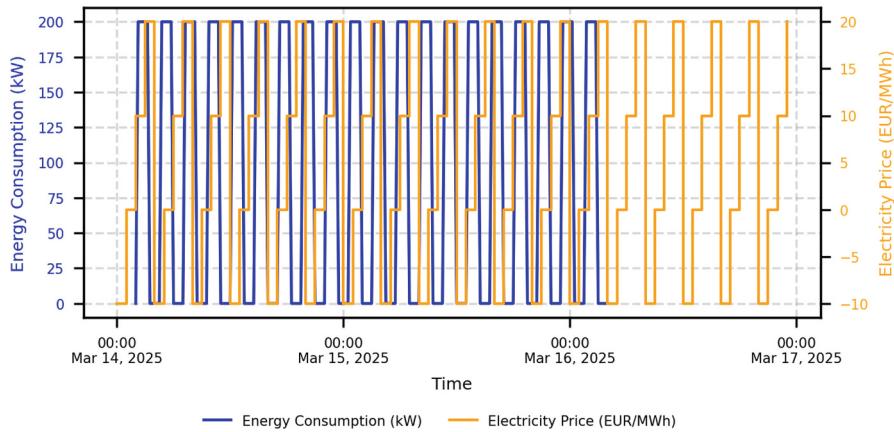


Fig. 2. Temporal Alignment of Synthetic Energy Consumption Profile with Electricity Price Fluctuations Over a Three-Day Horizon

The complete configuration of the synthetic dataset is summarized in Table 1, which outlines the resolution, structure, and periodicity of both the consumption and pricing components.

Table 1. Configuration Parameters for Synthetic Dataset

Parameter	Consumption Profile	Electricity Pricing Profile
Data Resolution	10 s	1 h
Consumption Pattern	Rise: 15 min Constant: 60 min Fall: 15 min Rest: 60 min	Repeating cycle: –10, 0, 10, 20 EUR/MWh
Duration of Cycle	90 min	4 h
Total Data Period	=72 h	72 h

3.3 Cost Calculation Procedure

The optimization evaluation procedure comprises two main components: cost calculation and visualization. For each melting cycle, the high-resolution consumption data is aligned with the hourly electricity price profile by replicating hourly prices at 10 s intervals. The total electricity cost C of a batch operating over a time interval $[t_0, t_1]$ is computed by integrating the instantaneous product of power consumption and electricity price:

$$C = \int_{t_0}^{t_1} P(t) \cdot \pi(t) dt \quad (1)$$

Where $P(t)$ is the power consumption in kW and $\pi(t)$ is the electricity price in DKK/kWh at time t . This enables accurate cost estimation even in scenarios with fine-grained consumption fluctuations.

To support exploratory analysis and validation, the results are visualized using both interactive and static plotting libraries. Plotly is used to enable dynamic exploration of cost structures and temporal alignment, while Matplotlib generates high-resolution static plots for documentation and reporting.

3.4 Batch Scheduling Optimization Procedure

The goal of the optimization procedure is to minimize total electricity cost by optimizing the start times of melting batches within a given planning window. Each batch is treated as a fixed-duration time block derived from either synthetic or real consumption data.

A full exhaustive search is performed over all valid start times at one-hour intervals. For each candidate's start time, the batch cost is calculated using the method described in Sect. 3.3. The configuration with the lowest cost is identified as the optimal solution. When extending the optimization to multiple batches, two critical operational constraints are enforced: (i) no overlap between batch executions, and (ii) a minimum one-hour buffer between consecutive batches to reflect production setup or cooling time.

Two distinct optimization strategies are implemented:

- **Independent Batch Optimization:** Each batch is optimized in isolation, assuming that scheduling of one batch does not constrain others.
- **Global Batch Optimization:** All batches are scheduled jointly, enforcing non-overlapping execution and a minimum one-hour buffer between batches.

This combinatorial optimization approach ensures rigorous and globally optimal scheduling solutions within the discretized scheduling space, providing tangible cost reductions in practice.

3.5 Validation with Real Industrial Data

To assess the real-world applicability and robustness of the proposed scheduling optimization framework, the full scheduling method was validated using actual manufacturing data from an industrial melting facility. Unlike the synthetic dataset, the real operational data does not exhibit uniform cycles or fixed patterns. Therefore, a data-driven approach was employed to preprocess and structure the raw inputs.

Power consumption data was collected at a 10 s resolution over a continuous six-day production window. In the absence of annotated batch identifiers, a heuristic segmentation rule was used to extract operational batches: a new batch was assumed to begin when the time gap between two successive power consumption events exceeded one hour. This approach allowed consistent batch identification across the unstructured data stream.

Once batches were identified, each was subjected to the full optimization pipeline. For each batch, a feasible shifting window was established beginning at the next full hour after its actual start time and extending until 16:00 on Friday, April 19—the latest operationally viable time. Both single-batch and multi-batch optimization routines were executed, with the latter enforcing non-overlapping constraints and a one-hour buffer between batch start times. Validation confirmed the framework’s effectiveness under real operational conditions, demonstrating its robustness and practical applicability.

4 Case Study: Energy Cost Optimization in a Danish Foundry

To demonstrate the application of the optimization framework in a real industrial setting, a case study was conducted using data from the Vald. Birn A/S foundry, located in the Jutland region of Denmark, which operates electric melting furnaces.

The dataset includes high-frequency (10 s resolution) power consumption data and hourly electricity spot prices for the DK1 market zone. DK1 encompasses Western Denmark, including Jutland and Funen, and is characterized by substantial intraday price variation. The spot price data was retrieved from the ENERGIDATA platform maintained by Energinet [15].

The selected case study period spans from 00:00 on Sunday, April 14, to 00:00 on Saturday, April 20, 2024. While data is available for the entire week, only the six working days (Sunday to Friday) were analyzed, as no furnace activity occurred on Saturday.

Figure 3 shows the temporal alignment between furnace power consumption and hourly electricity prices throughout the week. This visualization highlights the energy-intensive nature of the process and the variability in electricity costs that motivates optimization.

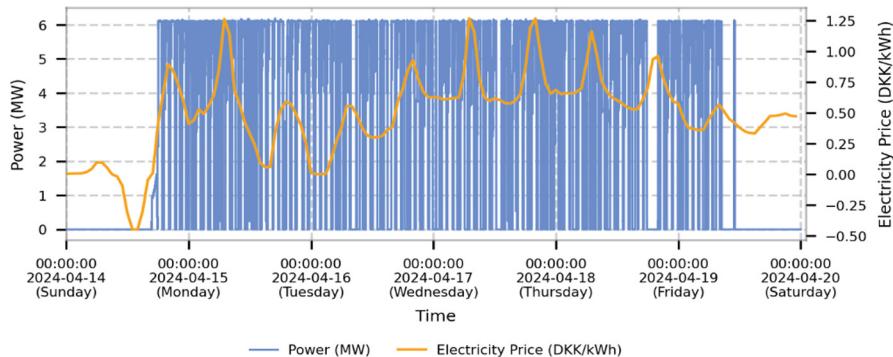


Fig. 3. Real Furnace Power Consumption and Electricity Price Over One Week in April 2024

5 Results

This section presents the outcomes of the proposed batch scheduling optimization framework, which integrates dynamic electricity pricing with high-resolution energy consumption data to minimize operational energy costs. The analysis proceeds from controlled experiments using synthetic datasets to real-world validation based on operational data from an industrial melting facility in Denmark. The goal is to demonstrate the feasibility and effectiveness of shifting batch start times to align energy-intensive operations with periods of lower electricity prices. Results are validated using both synthetic and real-world datasets from a Danish foundry, highlighting the achievable cost savings under practical constraints.

5.1 Cost Analysis Based on Synthetic Data

The synthetic experiments systematically assessed cost reduction potential under controlled conditions. The electricity price dataset covers a continuous 72-h period starting at 00:00 on March 14, 2025. Each synthetic batch includes multiple melting cycles, reflecting real industrial operations. In this experiment, the batch runs for 50 h, enabling full flexibility to explore all feasible start times within the 72-h price window. Figure 4 illustrates the total electricity cost associated with each possible shifted start time for the batch within the price window. A total of 22 discrete shifting options were evaluated at hourly increments. The figure clearly shows that the batch start time has a significant impact on total electricity costs, with cost variations driven by the alignment of energy-intensive production phases with periods of low or negative electricity prices.

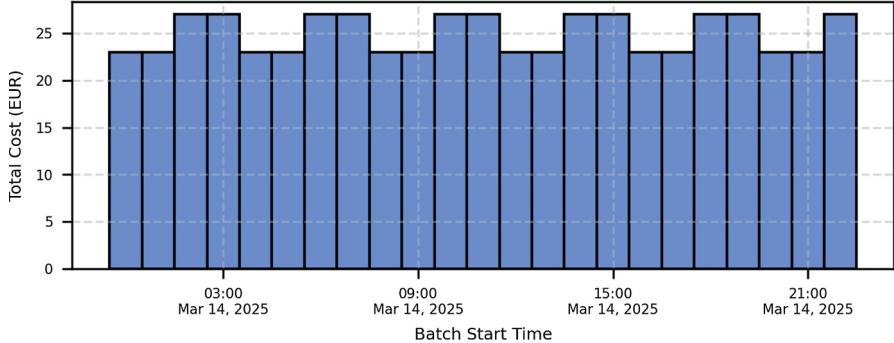


Fig. 4. Total Batch Cost for a Single Batch at Different Start Times Over a 72 h Synthetic Price Window

5.2 Batch Identification and Cost Analysis from Operational Data

To quantify the cost implications of real-time electricity pricing on industrial melting processes, furnace power consumption data recorded at 10 s intervals was analyzed over a six-day operational window. Using a heuristic temporal separation criterion—specifically, a minimum one-hour gap between consecutive consumption events—four distinct operational batches were identified. These batches are visually depicted in Fig. 5, where each is color-coded and overlaid on the corresponding electricity price profile in the DK1 zone of Denmark. The figure illustrates the temporal alignment between batch activity and hourly electricity price fluctuations. Each batch was delineated by its precise start and end timestamps and subsequently analyzed to compute the total electricity cost by synchronizing the high-resolution power data with the hourly pricing signals.

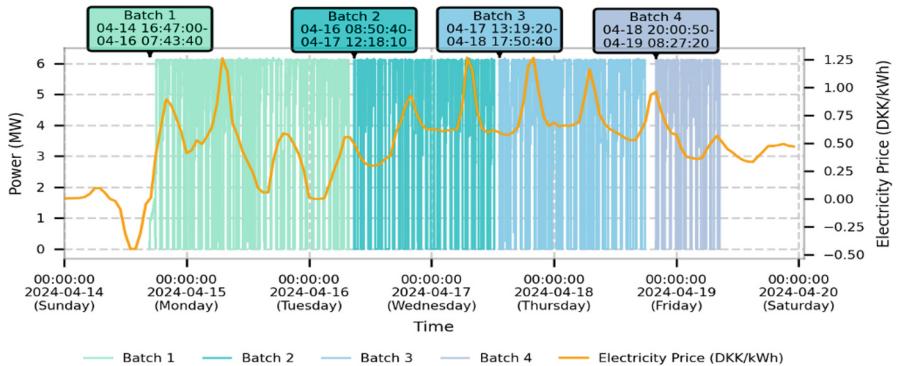


Fig. 5. Identified Batches Within the Week Aligned with Electricity Price

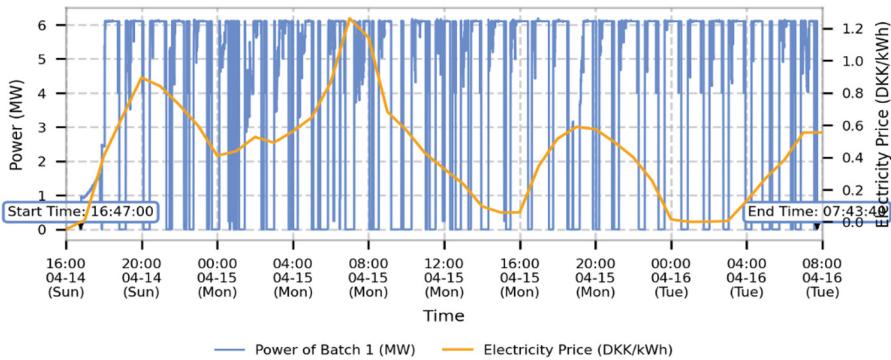
Table 2 presents the temporal attributes and electricity costs associated with the four identified operational batches. The durations of these batches ranged from approximately 12.44 to 38.94 h. Correspondingly, electricity expenditures varied significantly, reflecting differences in both duration and the timing of energy consumption relative to hourly price fluctuations.

Table 2. Summary of Identified Batches

Batch	Start time to End time	Duration [h]	Cost [DKK]
1	2024-04-14 16:47:00 (Sunday) 2024-04-16 07:43:40 (Tuesday)	38.94	67,228.97
2	2024-04-16 08:50:40 (Tuesday) 2024-04-17 12:18:10 (Wednesday)	27.46	60,150.09
3	2024-04-17 13:19:20 (Wednesday) 2024-04-18 17:50:40 (Thursday)	28.52	75,601.48
4	2024-04-18 20:00:50 (Thursday) 2024-04-19 08:27:20 (Friday)	12.44	25,209.68

5.3 Single Batch Optimization

To evaluate the potential cost savings achievable through temporal realignment, Batch 1 was selected for detailed optimization analysis. The original consumption profile of Batch 1, spanning from 16:47 on Sunday, April 14 to 07:43 on Tuesday, April 16, is illustrated in Fig. 6, alongside the corresponding hourly electricity price. The power profile shows multiple melting cycles executed during variable pricing periods, indicating an opportunity for cost reduction through strategic rescheduling.

**Fig. 6.** Original Power Consumption Profile of Batch 1 and Corresponding Electricity Price

The optimization window was defined from 17:00 on Sunday—the next full hour after the original batch start—through to 16:00 on Friday, April 19, marking the latest operationally feasible end time within the week. A grid search with one-hour increments was performed to evaluate the cost impact of all possible shifted batch schedules within this window.

Figure 7 depicts the total electricity cost associated with each shifted version of Batch 1. The analysis identified the minimum-cost schedule occurring when the batch commenced at 09:00 on Monday, April 15 and concluded at 23:56 on Tuesday, April 16.

This rescheduling resulted in a total electricity cost of 60,609.81 DKK, compared to the original 67,228.97 DKK—a reduction of 6,619.16 DKK, or approximately 9.8%. The cost curve shows that significant variation arises even within a narrow shifting window, reinforcing the importance of price-aware scheduling.

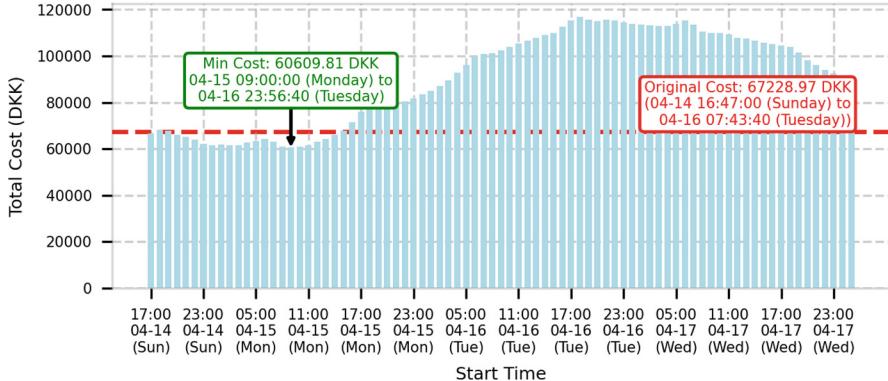


Fig. 7. Electricity Cost of Batch 1 Across All Possible Shifted Start Times

To further elucidate the benefit of this shift, Fig. 8 contrasts the original and optimized power consumption profiles overlaid on the electricity price curve. The optimized batch execution clearly avoids peak pricing periods while maintaining the same process duration and sequence, thereby validating the effectiveness of the proposed optimization strategy in reducing operational energy costs without compromising production integrity.

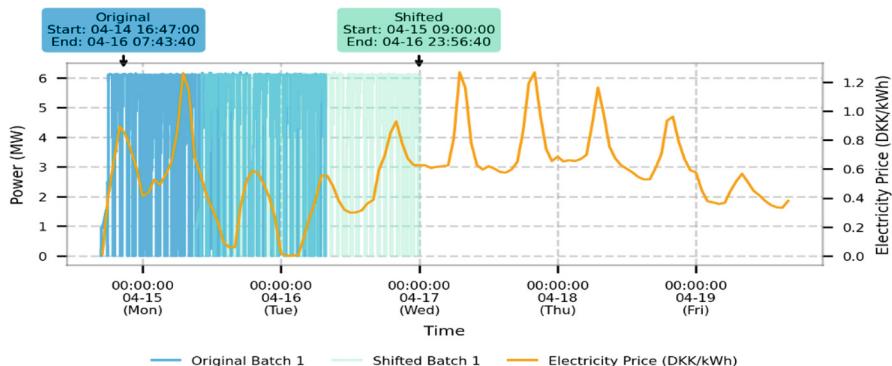


Fig. 8. Comparison of Original and Optimized Batch 1 Profiles Aligned with Electricity Price

The same optimization procedure was independently applied to all four batches, yielding notable cost reductions in each case. Optimal start times were identified through exhaustive testing of feasible shifts, assuming non-overlapping execution and a minimum one-hour interval between batches. As summarized in Table 3, electricity cost savings

ranged from 6,619.16 DKK to over 45,000 DKK. While this approach does not consider interactions between batches, it highlights the maximum potential for cost reduction when scheduling is performed independently.

Table 3. Optimized Schedules and Savings for All Batches

Batch	Start_time to End_time	Optimal Cost [DKK]	Saved Cost [DKK]
1	2024-04-15 09:00:00 (Sunday) 2024-04-16 23:56:40 (Tuesday)	60,609.81	6,619.16
2	2024-04-15 13:00:00 (Monday) 2024-04-16 16:27:30 (Tuesday)	28,600.68	31,549.41
3	2024-04-15 12:00:00 (Monday) 2024-04-16 16:31:20 (Tuesday)	30,301.95	45,299.53
4	2024-04-15 22:00:00 (Monday) 2024-04-16 10:26:30 (Tuesday)	11,068.80	14,140.88

5.4 Global Schedule Optimization

Building upon the single-batch optimization results, the method was extended to generate a globally optimized schedule for all four batches within the same operational week. The objective was to minimize the total electricity cost while satisfying key operational constraints—namely, non-overlapping batch execution and a minimum one-hour buffer between consecutive batches.

Table 4 summarizes the globally optimized start and end times for each batch, along with the corresponding electricity cost and the savings achieved relative to the original schedules. Although Batch 2 incurred a higher cost than in its original schedule, the overall optimization led to a substantial reduction in total weekly energy expenses—from 228,190.21 DKK to 211,407.09 DKK—resulting in net savings of 16,783.12 DKK.

Table 4. Optimized Weekly Schedule and Electricity Cost Summary

Batch	Start_time to End_time	Optimal Cost [DKK]	Saved Cost [DKK]
1	2024-04-15 09:00:00 (Monday) 2024-04-16 23:56:40 (Tuesday)	60,453.70	6775.27
2	2024-04-17 03:00:00 (Wednesday) 2024-04-18 06:27:30 (Thursday)	71,860.63	-11,710.54
3	2024-04-18 11:00:00 (Thursday) 2024-04-19 15:31:20 (Friday)	54,468.37	21,133.11
4	2024-04-14 17:00:00 (Sunday) 2024-04-15 05:26:30 (Monday)	24,624.39	585.29

The optimized schedule is visualized in Fig. 9, which illustrates the temporal alignment of batch operations with the hourly electricity price profile. The figure highlights

how the rescheduling preferentially aligns batch activity with lower-price periods across the week.

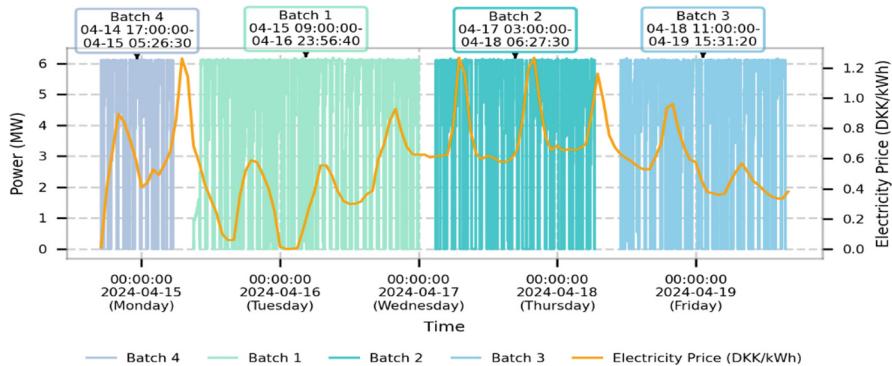


Fig. 9. Globally Optimized Weekly Batch Schedule Aligned with Electricity Price Profile

6 Discussion

This study confirms the practical potential of data-driven scheduling optimization to reduce electricity costs in energy-intensive industrial batch processes under dynamic pricing. Experimental results using both synthetic and real industrial data demonstrate that strategically shifting batch start times—without altering production durations or output volumes—yields substantial economic benefits, reinforcing and extending previous research [1, 3, 12, 16].

In synthetic experiments, batch processes modeled over a 72-h price window showed that simply shifting batch start times significantly reduced electricity costs. In the single-batch scenario, moving Batch 1 from Sunday 16:47 to Monday 09:00 saved approximately 6,619 DKK, consistent with prior studies such as [3] and [12], which demonstrated cost benefits from aligning energy-intensive operations with off-peak tariffs. However, unlike these studies that employed heuristic or multi-objective approaches, the present method guarantees global cost optimality through exhaustive combinatorial search.

The multi-batch (global) optimization analysis further underscores the superiority of holistic scheduling strategies over independent adjustments. Across the six-day production schedule involving four batches, global optimization achieved cumulative weekly savings of 16,783.12 DKK, representing a 7.4% reduction in total electricity costs. This result remained robust despite slight cost increments observed for individual batches, notably Batch 2. Similar observations were highlighted by recent studies using combinatorial and evolutionary approaches [1, 16], where global optimization consistently outperformed isolated batch adjustments. Moreover, the implementation effort required is relatively modest, further enhancing the practical viability of this optimization framework for industrial settings.

Compared to heuristic methods such as genetic algorithms or PSO, the exhaustive search used in this study guarantees global optimality but incurs higher computational

cost. While many prior models rely on synthetic data or simplified constraints, this framework integrates real industrial data, high-resolution power profiles, and operational constraints such as non-overlapping batches and mandatory downtimes [2, 4, 5, 16, 17]. In contrast to recent stochastic or machine learning-based scheduling approaches [4], it uses deterministic inputs and avoids the need for large-scale training data. Although not algorithmically novel, the method demonstrates strong alignment with real production conditions and serves as a practical roadmap for industrial adoption. It also offers a robust benchmark against which heuristic and data-driven alternatives can be evaluated [1, 16, 18].

Nevertheless, two caveats warrant emphasis. First, although power is measured every 10 s, prices are only hourly, so cost estimates—despite sub-minute load fidelity—cannot exceed the one-hour price granularity without true intra-hour market data. Second, our exhaustive grid search grows factorially with batch count and time-step resolution, becoming impractical for large or finer-grained problems; in such cases, meta-heuristics, decomposition methods, or parallelized searches can deliver near-optimal schedules within reasonable runtimes.

Optimization runtimes on a test system equipped with a 12-core, 3.7 GHz processor and 64 GB RAM averaged less than 10 s for single-batch scheduling and approximately 4 min for scheduling four batches over a six-day planning horizon. While these runtimes are acceptable for offline planning, the factorial growth in computational complexity associated with increasing batch numbers or finer discretization poses scalability challenges. Potential enhancements include pruning strategies such as branch-and-bound or employing heuristic algorithms to obtain rapid, near-optimal solutions. Furthermore, by incorporating multi-machine configurations or inter-batch dependency constraints, this optimization framework can readily extend to more complex, networked industrial production scenarios.

Detailed power consumption profiles revealed within each batch multiple melting cycles with intermittent feeding. This fine-grained operational characteristic, explicitly observed in real data analyses, suggests untapped potential for further optimization. Future research should explore intra-batch scheduling strategies, including temporary pauses during peak prices, to unlock additional flexibility and cost reductions.

7 Conclusion

This study proposed a data-driven scheduling optimization framework for industrial batch processes under dynamic electricity pricing. By integrating high-resolution power consumption profiles with market price signals, the method identifies optimal batch start times that substantially reduce electricity costs without disrupting operations. Validation using synthetic and real industrial data from a Danish foundry demonstrated cost savings of 6,619 DKK for a single batch and 16,783 DKK across four batches—equivalent to a 7.4% weekly reduction in electricity expenditure—highlighting the benefits of holistic over isolated scheduling.

The research advances the field by providing a globally optimal combinatorial optimization method, overcoming the limitations of heuristic approaches [12, 16, 19]. Incorporating both synthetic and real-world data ensures robust validation, while the use of 10

s consumption intervals enhances cost modeling precision and responsiveness to market fluctuations [3–5].

Practically, the framework offers a ready-to-implement tool for energy managers to achieve significant savings without additional investments. Its alignment with real-world constraints ensures easy integration into existing workflows, enabling industries to leverage dynamic pricing for both economic and sustainability gains.

Despite its contributions, this study has certain limitations. Limitations include the use of discretized hourly intervals, reliance on deterministic price forecasts, and focus on single-machine scenarios. Addressing these through finer time resolutions, stochastic price modeling [4, 5], and extension to multi-machine environments can further enhance applicability and savings potential. Exploring intra-batch optimization also presents a promising future direction.

This work contributes to the field of energy-aware production scheduling by demonstrating that significant, verifiable cost savings can be achieved through data-driven batch scheduling under dynamic electricity pricing, with strong alignment to both operational feasibility and sustainability goals.

Acknowledgements. This paper is part of project “Data-driven best-practice for energy-efficient operation of industrial processes - A system integration approach to reduce the CO₂ emissions of industrial processes” (Case no.64020-2108) by the Danish funding agency, the Energy Technology Development and Demonstration (EUDP) program, Denmark; part of the project titled “Enabling Sector Coupling of Energy-Intensive Industrial Processes in Energy Systems (Electricity and District Heating)” (jr. Nr. FRO-23-0026; Erhvervsfyrtårn Syd Fase 2), The European Regional Development Fund; part of the project titled “Danish participation in IEA IETS Task XXII - Climate Resilience and Energy Adaptation in Industry under Uncertainty”, funded by EUDP (project number: 95-41006-2410288).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Burmeister, S.C., Guericke, D., Schryen, G.: A memetic NSGA-II for the multi-objective flexible job shop scheduling problem with real-time energy tariffs. *Flex. Serv. Manuf. J.* **36**, 1530–1570 (2024)
2. Cheng, J., et al.: Bi-criteria single-machine batch scheduling with machine on/off switching under time-of-use tariffs. *Comput. Ind. Eng.* **112**, 721–734 (2017)
3. Wang, S., et al.: Bi-objective optimization of a single machine batch scheduling problem with energy cost consideration. *J. Clean. Prod.* **137**, 1205–1215 (2016)
4. Wang, C., et al.: Deep reinforcement learning-based energy-conscious scheduling under time-of-use electricity price. *IEEE Trans. Autom. Sci. Eng.* **22**, 11490–11504 (2025)
5. Felder, M., et al.: Energy-Flexible Job-Shop Scheduling Using Deep Reinforcement Learning (2023)
6. Ding, J., et al.: Energy aware scheduling in flexible flow shops with hybrid particle swarm optimization. *Comput. Oper. Res.* **125**, 105088 (2021)
7. Guo, J., et al.: A GAPN Approach for the flexible job-shop scheduling problem with indirect energy and time-of-use electricity pricing. *PROCESSES ** **10**(5) (2022)

8. Burmeister, S.C., Guericke, D., Schryen, G.: A two-level approach for multi-objective flexible job shop scheduling and energy procurement. *Cleaner Energy Syst.* **10**, 100178 (2025)
9. Wang, Z., Man, Y.: Machine learning-based intermittent equipment scheduling model for flexible production process. *Appl. Artif. Intell. Process Syst. Eng.* **23**, 473–495 (2021)
10. Xiao, H., et al.: Optimization of Cold Rolling Jobs under Time-of-Use Electricity Prices Based on Multi-Objective Genetic Algorithm and Power Prediction (2024)
11. Qian, S.y., Jia, Z.h., Li, K.: A multi-objective evolutionary algorithm based on adaptive clustering for energy-aware batch scheduling problem. *Future Gener. Comput. Syst.* **113**, 441–453 (2020)
12. Zhang, S., et al.: Improved mixed-integer linear programming model and heuristics for bi-objective single-machine batch scheduling with energy cost consideration. *Eng. Optim.* **50**, 1380–1394 (2018)
13. Xie, S., Li, Y., Wang, P.: Optimal scheduling of off-site industrial production in the context of distributed photovoltaics. *Energies* **17**(9), 2156 (2024)
14. Abikarram, J.B., McConky, K., Proano, R.: Energy cost minimization for unrelated parallel machine scheduling under real time and demand charge pricing. *J. Clean. Prod.* **208**, 232–242 (2019)
15. Energinet: Elspot Prices. <https://www.energidataservice.dk/tso-electricity/Elspotprices>. Accessed May 2025
16. Wu, P., Cheng, J., Chu, F.: Large-scale energy-conscious bi-objective single-machine batch scheduling under time-of-use electricity tariffs via effective iterative heuristics. *Ann. Oper. Res.* **296**, 471–494 (2021)
17. Shrouf, F., et al.: Optimizing the production scheduling of a single machine to minimize total energy consumption costs. *J. Clean. Prod.* **67**, 197–207 (2014)
18. Jiang, E.-D., Wang, L.: Multi-objective optimization based on decomposition for flexible job shop scheduling under time-of-use electricity prices. *Knowl.-Based Syst.* **204**, 106177 (2020)
19. Tian, Z., Zheng, L.: Single machine parallel-batch scheduling under time-of-use electricity prices: new formulations and optimisation approaches. *Eur. J. Oper. Res.* **312**, 512–524 (2024)

Energy Informatics for Electric Vehicles and Mobility Systems



The Irish Highway Network: A Novel Test Instance for the Charging Station Location Problem

Jingyu Xiang^(✉), Paula Carroll, and Annunziata Esposito Amideo

School of Business, University College Dublin, Dublin, Ireland

jingyu.xiang@ucdconnect.ie

Abstract. Electric Vehicles (EVs) are one of the most promising sustainable transportation solutions but require efficient refueling strategies to complete their long-distance journeys. When addressing the charging station location problem (CSLP), a well-developed network enables operators to strategically locate charging stations, efficiently mitigating range anxiety among EV drivers, and supporting the wider adoption of EVs. This study provides a description of the methodology for creating a real-world CSLP instance based on the Irish highway network and traffic count data. We capture key settlements and intersections from the Irish national road map, and construct the Irish highway network as a graph, made up of motorways, national primary roads, and national secondary roads. Applying the Lincoln MPO Travel Model, we estimate the number of EV traffic flows between all origin-destination pairs based on traffic counts on each road segment. This study evaluates the Irish highway network test instance by analyzing graph-based metrics and compares its characteristics with other widely-used CSLP instances from the literature. The results of the comparative analysis indicate that the proposed Irish highway network is a well-connected and favourable test instance to evaluate CSLP optimization models.

Keywords: OR in Energy · Sustainable Transportation · Charging Station Location Problem · Irish Highway Network

1 Introduction

Severe climate change has attracted worldwide attention, prompting urgent discussions among governments, researchers, and policymakers. A total of 194 parties, including 193 states and the European Union, signed the Paris Agreement, targeting to achieve carbon neutrality by the mid-21st century. As the second-largest contributor to global greenhouse gas emissions among all energy-related sectors, the emissions from transport collectively accounts for approximately one-quarter of all emissions worldwide [10]. Consequently, the transition to electric vehicles (EVs) or other sustainable transport solutions is a key strategy for reducing greenhouse gas emissions and achieving carbon neutrality. As the adoption

of EVs increases, the availability and accessibility of the charging infrastructure become critical to ensuring an efficient transportation system that supports EV charging demands. The charging station location problem (CSLP) aims to strategically locate charging stations to allow EVs with limited driving ranges to refuel, so that they can complete their journey without running out of charge [11].

In Ireland, the Electric Vehicles Charging Infrastructure Strategy 2022-2025 [4] launched a 100 million euro investment on public charging infrastructure to assure car users the feasibility of switching to an EV, as part of the national strategy to have one million EVs on the road by 2030. Existing research on the CSLP relies heavily on network-based optimization models to determine optimal locations of charging stations. Therefore, to optimize the strategy for charging station deployment in Ireland, a realistic and accurate representation of the Irish highway network is needed, including nodes and highway links, candidate locations for charging stations, and traffic volumes in origin-destination (O-D) pairs. Without a realistic data-driven test instance, advanced CSLP optimization models may yield non-optimal and impractical solutions.

In existing CSLP literature, there is a lack of a realistic data-driven Irish highway network, which limits the applicability of CSLP solutions to actual infrastructure planning. How to construct a realistic representation of the Irish highway network with a reliable traffic flow dataset for solving the CSLP emerges as our research question. To address this research question, we demonstrate the approach to constructing the Irish highway network instance by integrating geographical information and traffic count data that researchers can adapt to develop realistic networks for solving the CSLP in their own case studies.

This paper is structured as follows. Section 2 presents the background literature on CSLP optimization models and the test networks used to solve the CSLP. Section 3 describes the process for building the Irish highway network instance and generating the traffic count data for each O-D pair. Section 4 presents the evaluation of our developed Irish highway network along with a comparison with other commonly used test networks. Section 5 summarizes the notable results and concludes our study.

2 Literature Review

A wide range of modeling approaches has been employed in CSLP studies, while available test instances used in these studies are very limited, as summarized in Table 1.

Simchi-levi and Berman [19] build a sample 25-node network to test their heuristic algorithm for the Traveling Salesman Location Problem (TSLP). The Flow-Capturing Location Model (FCLM) [7], developed to address facility location problems with flow-based demands, first employs the 25-node network from Simchi-levi and Berman [19]. The FCLM utilizes population weights attributed to each node and applies a gravity model [6] to generate traffic flow data of O-D pairs. The model assumes that demand flows can only be captured if at

least one facility is located on the shortest path and maximizes the number of flows that a fixed number of facilities can capture. Then, the same 25-node network is widely used in flow-based CSLP models, such as the first Flow Refueling Location Model (FRLM), which incorporates limited vehicle driving range in the FCLM [13]; the modified FRLM with dispersion of candidate sites on arcs [14]; the Deviation Flow Refueling Location Model (DFRLM), which allows EVs to deviate from their shortest paths [12].

The Sioux Falls network is another classic test network, quite often used in continuous network design problems. Originally constructed to validate efficient methods for solving road network equilibrium traffic assignment problems [16], the Sioux Falls network is also employed in many CSLP models. For example, it is used in the Multipath Refueling Location Model (MPRLM) [9], which allows electric vehicles to take multiple deviation paths between all O-D pairs.

Several network instances employed for the CSLP are developed based on real road networks. Considering the capacity of charging stations, Upchurch et al. [22] test the capacitated Flow Refueling Location Model (CFRLM) on a simplified version of the Arizona state highway network. The network consists of Arizona's 25 largest cities and links between them, including major US state and interstate highways. Candidate sites for charging stations in the CFRLM include the 25 nodes, as well as another 25 nodes distributed using the Added Node Dispersion Problem method (ANDP) [15] to allow for station locations between the major nodes. The flow volumes between 25 cities are estimated using a gravity model, based on their population weights.

Due to the limited available test networks, some CSLP studies also employ randomly generated network instances to evaluate proposed optimization models [1, 8]. When studying a capacitated charging station location problem while accounting for multi-period travel demand, Hosseini et al. [8] generate test networks by randomly placing nodes in a 100×100 grid using a uniform distribution. Then, they generate a Minimum Spanning Tree (MST) to ensure that all nodes are connected and add extra links for better connectivity. The traffic flows of randomly selected O-D pairs are estimated using the same gravity model with randomly assigned node weights.

Drawing from the discussion, CSLP studies still lack available, realistic and large-scale test instances developed specifically for solving the CSLP, and most test instances used in CSLP studies estimate traffic flows in O-D pairs by using the gravity model based on node population weights, which may result in notable deviations from real-world travel demands.

3 Highway Network Construction Methodology for CSLP

This section describes the methodology for building the CSLP test instance of the Irish highway network, ensuring accurate modeling of network topology and traffic flow count data, adapted from approaches in previous studies [1, 8, 18].

Table 1. Key Features of Available Test Instances Used in CSLP Studies

Authors	Networks	Nodes	Edges	Candidate Sites	O-D pairs	Traffic Input
Hodgson (1990) [7]; Kuby and Lim (2005) [13]; Kim and Kuby (2012) [12]	25-node network by Simchi-Levi & Berman (1988) [19]	25	43 undirected	25 nodes	600 directed	Estimated via gravity model
Kuby and Lim (2006) [14]	25-node network by Simchi-Levi & Berman (1988) [19]	25	43 undirected	25 nodes & added sites	600 directed	Estimated via gravity model
Upchurch et al. (2009) [22]	Arizona state highway network	25	39 undirected	25 nodes & added sites	600 directed	Estimated via gravity model
Huang et al. (2015) [9]	Sioux Falls network [16]	24	552 directed	24 nodes	176 directed	Given O-D pairs
Capar et al. (2013) [1]; Hosseini et al. (2017) [8]	Randomly generated instances	25–250	/	All nodes generated	Randomly selected	Estimated via gravity model

3.1 Node Placement and Highway Connections

The Irish highway test instance for CSLP is composed of a set of nodes representing settlements and highway intersections, and a set of edges representing the connecting highway links. We first determine the locations of these nodes based on the topology of the Irish highway network.

Obtaining the Irish national road network from Transport Infrastructure Ireland [20], we import the downloaded KML file, which contains the geographic information for motorways, national primary roads, and national secondary roads, into QGIS for spatial analysis. Integrated with geographic information, we capture nodes along the highways, including key settlements, highway intersections, and junctions. The placement of nodes is determined to ensure a comprehensive representation of the structure and connectivity of the highway network. We then classify these selected nodes as “centers” that can generate and attract traffic flows or “connections”. “Connections” do not generate travel demands but serve to maintain network connectivity and enable more different route options for each O-D pair. To determine the classification, we collect the population of settlements near these nodes from the 2022 Ireland Census report [2]. The subset of “centers” is selected based on the following criteria, including population size and spatial distribution:

- As defined by the Central Statistics Office of Ireland, a settlement is classified as a city if it has a population exceeding 50,000, and those with populations between 1,500 and 49,999 are classified as urban towns, while those with fewer than 1,500 inhabitants are considered rural areas. In our study, all nodes near the settlements with population exceeding 5,000 are collectively referred to as “centers”, which can generate and attract observable number of inter-city traveling flows;
- In remote or sparsely connected regions, the nodes with relatively small populations are also classified as “centers”, to prevent isolated areas from being excluded in the analysis of travel demand patterns. This consideration helps ensure generated traffic flows are geographically well-distributed across the entire network;
- Nodes located near the border with Northern Ireland are also classified as “centers” to account for cross-border traffic flows.

We then connect these nodes with Irish highway links. Using the geometry-based field calculator in QGIS, the lengths of all highway links between these identified nodes are measured in meters, ensuring accurate measurements of existing highway connection segments. The final test instance of the Irish highway network is represented as a graph $G = (V, E)$, where:

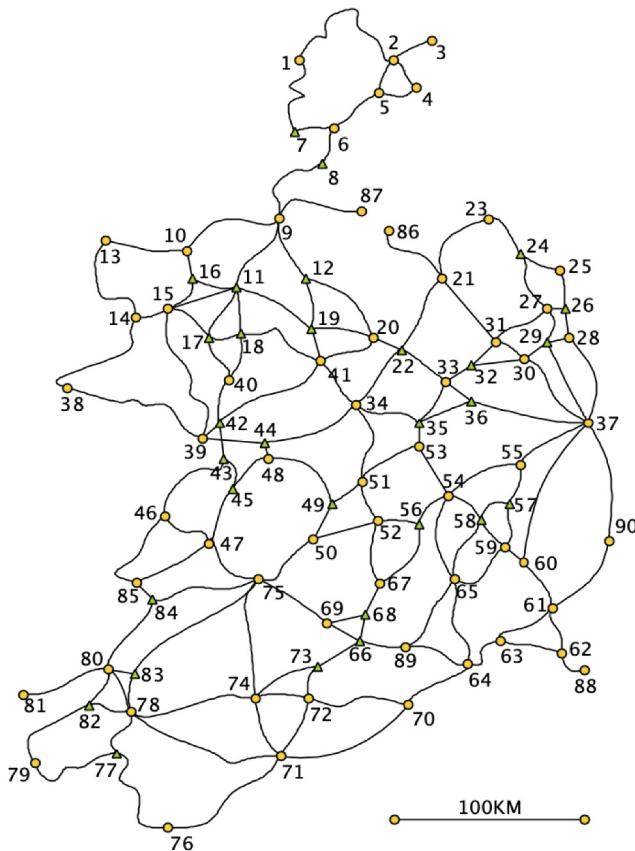
- V represents the set of total 90 nodes, including 60 “centers” that generate and attract inter-city travel demands in traffic flow analysis and 30 “connections” serving as highway intersections and junctions. All these nodes are candidate locations for charging stations in the CSLP;
- E represents the set of edges, consisting of 152 undirected links that correspond to road segments between all nodes.

Figure 1 illustrates the test instance of the Irish national highway network, visualised using the yED Graph Editor based on the determined nodes and links. The round nodes represent “centers” that generate and attract inter-city travel demands, while the triangle “connections” represent other highway intersections and junctions.

In addition, the existing charging stations on Irish highways are also included in our proposed instance, based on the Ireland car charging points map provided by the Electricity Supply Board [5]. Given that every charging station is placed at one of those nodes in the Irish highway instance, existing charging points, which are in reality either on a highway node or midpoint link, are associated to the nearest node location.

3.2 EV Traffic Flow Estimation

To effectively address the CSLP on the Irish highway network, a realistic dataset of traffic flow on the highway network instance is essential. Instead of using population weights of nodes, this study utilizes traffic counts on the highway road to estimate EV traffic flows in O-D pairs on the Irish highway network.



Source: Author's own

Fig. 1. The Irish Highway Network for Charging Station Location Problems

Transport Infrastructure Ireland provides traffic count data [21] on road segments, with a dynamic mapping interface to detail the traffic volume by hour of day and vehicle class from traffic counters in a variety of formats. However, most CSLP optimization models require traffic flow data specifying the number of vehicles traveling between O-D pairs. To bridge this gap, we adopted the gravity-based Lincoln MPO Travel Model [17] to estimate O-D flows between center nodes in our constructed instance using available traffic count data.

The gravity model assumes that the trips produced at an origin and attracted to a destination are directly proportional to the total trip productions at the origin and the total attractions at the destination, and inversely proportional to the total attractions of all zones, which is represented as:

$$T_{ij} = \frac{A_j F_{ij} K_{ij}}{\sum_{j=1}^n A_j F_{ij} K_{ij}} \times P_i, \quad \forall i, j \quad (1)$$

where:

- T_{ij} is the number of trips produced at i and attracted to j ;
- P_i is the total trip production at i ;
- A_j is the total trip attraction at j ;
- F_{ij} is a calibration term for interchange ij , representing a friction factor or travel time factor;
- K_{ij} is a socioeconomic adjustment factor for interchange ij ;
- i is a origin zone;
- j is a destination zone;
- n is number of zones.

The equation distributes the number of trips produced at each origin i (P_i) across all potential destinations j , proportionally based on the product of total attractiveness of the destination (A_j), the impedance or friction factor between zones (F_{ij}) and any socioeconomic adjustment factors (K_{ij}). It ensures that the sum of trips from zone i to all destinations equals the total production P_i .

To apply the MPO Travel Model to estimate the number of O-D pairs between center nodes in the Irish highway instance, we collect the daily traffic volume from the nearest traffic counters of center nodes. For each center node, the trip production (P_i) is calculated by aggregating traffic volumes on all highway segments leading out of the node. Conversely, the trip attraction (A_j) is determined by aggregating the traffic volumes from all highway segments entering the node. Moreover, when calculating the trip production and attraction of nodes adjacent to Northern Ireland, we take the number of trips entering and leaving Northern Ireland into account. The friction factors (F_{ij}) in the Lincoln MPO Travel Model are estimated using the travel time between O-D pairs, which inversely correlates with the shortest path length. In addition, the K-factors (K_{ij}) are not considered in our estimation.

Using this approach, we transform traffic count data on road segments into a structured O-D matrix of traveling demands, ensuring the compatibility of the available collected data and the input dataset of CSLP models for further analysis and optimization. According to Climate Action Plan [3], the number of EVs registered in Ireland is targeted to exceed one million by 2030, with additional 125,000 sustainable journeys. Based on this projected growth, we assume that private EVs will account for approximately 30% of the total highway traffic volume by 2030. By applying this proportion, the estimated traffic flow data of O-D pairs for private EVs is generated.

The data sets for our Irish highway network instance are available on Github¹, including CSV files of nodes, edges, current charging points and EV traffic

¹ [Click here to download](#).

flow O-D pairs, a visualisation of the Irish highway network, and the Python import script. A README file is provided within the repository to document the instance metadata and explain the structure and contents of the dataset.

4 Results and Analysis

The goal of this study is to construct and validate a realistic test instance of the Irish highway network for solving the CSLP. We conduct a series of measurements of graph-based metrics that capture key topological and structural features of the network, such as size and scale, connectivity, diameter, average path length and average edge length, as summarized in Table 2. By comparing evaluation metrics with those instances used in previous CSLP studies, the results demonstrate that our developed Irish highway instance serves as a realistic and empirically validated test network for EV infrastructure planning.

Table 2. Comparison of Test Networks Used in CSLP Models

Network	Nodes	Edges	Avg Degree	Diameter	Avg Path Length	Avg Edge Length
25-node Network (1988) [19]	25	43 undirected	3.44	38.00	14.23 (0.374)	4.60 (0.121)
Sioux Falls Network (1975) [16]	24	76 directed	3.16	730	300.2 (0.411)	104.02 (0.142)
Arizona State Network (2009) [22]	25	39 undirected	3.12	431.00	205.10 (0.476)	74.85 (0.174)
Ireland Highway Network	90	152 undirected	3.38	555.10 km	202.35 km (0.365)	36.24 km (0.065)

Legend. **Avg Degree:** The average number of edges per node in the graph;

Diameter: The longest shortest path length between any two nodes;

Avg Path Length: The average shortest path lengths between all pairs of nodes;

Avg Edge Length: The average length of all edges;

4.1 Size and Scale

Compared to benchmark test networks such as the 25-node network [19], the Sioux Falls network [16] and the Arizona state network [22], our constructed Irish highway network instance provides a large-scale structure, consisting of 90 nodes and 152 edges, and totaling over 3,500 traffic flow O-D pairs. With a higher number of nodes, edges, and potential charging station locations, the number of variables and constraints in CSLP optimization models grows substantially, making it more computationally difficult to find an optimal solution. This prompts the application and development of advanced heuristics to tackle real-world CSLP case studies.

4.2 Connectivity

Figure 2 shows the degree distribution in the Irish highway network. Most nodes have a degree of 3 or 4, indicating typical highway intersections and junctions. A few nodes have very high degrees (e.g. 6 or 7), likely representing major highway junctions or big cities. Notably, several nodes have a degree of 1, indicating that they are endpoints within the highway network, often in remote or sparsely connected regions. While such nodes do not significantly contribute to overall network connectivity or serve as high-traffic points, they are critical for ensuring the accessibility of the Irish highway network, especially for settlements and attraction points in rural areas.

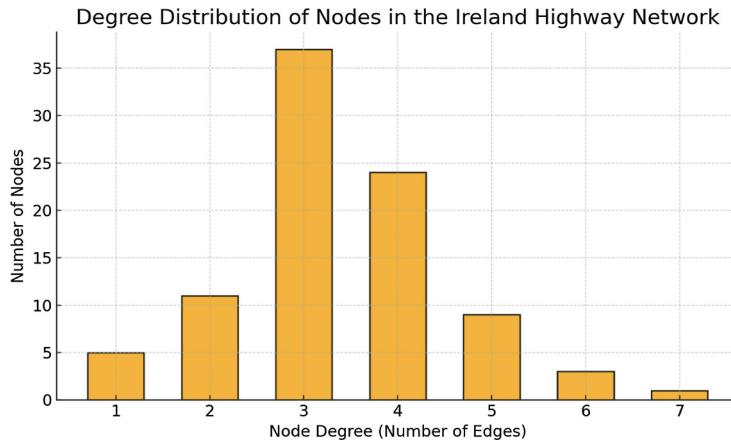


Fig. 2. Degree distribution of nodes in the Ireland highway network

To compare the connectivity of different networks, average degrees are measured, presented in Table 2, simply indicating the average number of edges per node in the graph. When compared to smaller benchmark test networks in CSLP studies, the average node degree of 3.38 in our network indicates a well-connected graph structure. This level of connectivity is particularly notable, considering the inclusion of remote nodes that, despite their weak connections, can enable EV traffic flows throughout the whole highway network.

4.3 Diameter, Average Path Length, and Average Edge Length

In the Irish highway network, the diameter defined as the longest shortest path between any two nodes is 555.10 km. The average shortest path lengths between all pairs of nodes equals 202.35 km and the average length of all edges in the network is only about 36 km. To enable comparisons across different networks, the values of the average path length and average edge length are scaled by dividing them by the network diameter. The resulting scaled values are shown

in parentheses in Table 2. When comparing the scaled values of average path length and average edge length with the same graph diameter, the Irish highway network shows the smallest values among the networks analyzed. The shorter edges and path lengths relative to the diameter indicates that nodes in our test instance are more easily reachable on average and the network supports a larger number of alternative deviation paths for EVs between O-D pairs with limited vehicle driving ranges. This suggests the Irish highway instance exhibits a higher level of accessibility and connectivity, enabling EV users to have more route options, thereby enhancing route flexibility in CSLP applications.

This evaluation validates that the instance not only reflects the spatial layout of the Irish highway network but also demonstrates strong connectivity, higher accessibility, and with a larger scale compared to other networks. It proves that the Irish instance is a well-suited and favourable choice as a realistic test instance developed specifically for solving the CSLP.

5 Conclusion and Findings

This study aims to develop a realistic test instance of the Irish highway network for solving the CSLP, while providing a detailed approach that can be replicated in similar case studies using available geographical and traffic data.

We have shown that by integrating the national road map, spatial data, population distribution, and traffic count data, we can produce a realistic CSLP instance that effectively captures real-world conditions and travel demands. This study first locates nodes along the highways, and classifies them into two categories: “centers” and “connections”. Using the GIS approach, we then obtain the measurements of highway connections as lengths of edges in the Irish highway network. After determining the Irish highway network structure, we apply the Lincoln MPO Travel Model to estimate the number of travel demands generated between selected O-D pairs, based on daily traffic count data collected on highway segments. Subsequently, we evaluate the Irish highway network, analyze graph-based metrics and compare them with those of benchmark test networks. The results demonstrate that the developed Ireland highway network is well connected and exhibits stronger accessibility, making it a practical test instance for solving the CSLP.

Even though the Irish highway network instance performs well in terms of scale, connectivity, and accessibility, further enhancements can be explored beyond the scope of this paper. For example, we can expand the highway network to include regional roads and incorporate more candidate locations to enhance the accuracy of optimal solutions for charging station locations in the whole transportation system. To better accommodate real-world conditions, temporal variations in traffic and charging demands, as well as elevation variations along highway roads, can be taken into consideration. In addition, varying charging behaviours and multimodal networks can be more explored for more advanced approaches to charging station placement strategies.

Acknowledgements. Jingyu Xiang is a PhD student at the University College Dublin College of Business, funded by the China Scholarship Council (CSC).

References

1. Capar, I., Kuby, M., Leon, V.J., Tsai, Y.-J.: An arc cover-path-cover formulation and strategic analysis of alternative-fuel station locations. *Eur. J. Oper. Res.* **227**(1), 142–151 (2013)
2. Central Statistics Office: Census 2022 Summary Results. Government of Ireland (2022). <https://data.cso.ie/product/C2022SR>
3. Department of the Environment, Climate and Communications: Climate Action Plan 2024. Government of Ireland (2023). <https://www.gov.ie/en/publication/79659-climate-action-plan-2024/>
4. Department of Transport and ZEVI: Electric Vehicle Charging Infrastructure Strategy 2022–2025. Government of Ireland (2023). <https://www.zevi.ie/publications/ev-infrastructure-strategy-2022-2025>
5. Electricity Supply Board: Charge Point Map. Government of Ireland (2025). <https://esb.ie/what-we-do/ecars/charge-point-map#>
6. Fotheringham, A., O’Kelly, M.: Spatial Interaction Models: Formulations and Applications. Kluwer Academic Publishers (1989)
7. Hodgson, M.: A flow-capturing location-allocation model. *Geogr. Anal.* **22**(3), 270–279 (1990)
8. Hosseini, M., MirHassani, S.A., Hooshmand, F.: Deviation-flow refueling location problem with capacitated facilities: model and algorithm. *Transp. Res. Part D: Transp. Environ.* **54**, 269–281 (2017)
9. Huang, Y., Li, S., Qian, S.: Optimal deployment of alternative fueling stations on transportation networks considering deviation paths. *Networks Spatial Econ.* **15** (2015)
10. IEA: Global energy-related CO₂ emissions by sector. IEA, Paris. <https://www.iea.org/data-and-statistics/charts/global-energy-related-co2-emissions-by-sector>. Licence: CC BY 4.0
11. Kchaou-Boujelben, M.: Charging station location problem: a comprehensive review on models and solution approaches. *Transp. Res. Part C: Emerging Technol.* **132**, 103376 (2021)
12. Kim, J.-G., Kuby, M.: The deviation-flow refueling location model for optimizing a network of refueling stations. *Int. J. Hydrogen Energy* **37**(6), 5406–5420 (2012)
13. Kuby, M., Lim, S.: The flow-refueling location problem for alternative-fuel vehicles. *Socioecon. Plann. Sci.* **39**(2), 125–145 (2005)
14. Kuby, M., Lim, S.: Location of alternative-fuel stations using the flow-refueling location model and dispersion of candidate sites on arcs. *Netw. Spat. Econ.* **7**(2), 129–152 (2006)
15. Kuby, M., Lim, S., Upchurch, C.: Dispersion of nodes added to a network. *Geogr. Anal.* **37**(4), 383–409 (2005). <https://doi.org/10.1111/j.0016-7363.2005.03704002.x>
16. LeBlanc, L.J., Morlok, E.K., Pierskalla, W.P.: An efficient approach to solving the road network equilibrium traffic assignment problem. *Transp. Res.* **9**(5), 309–318 (1975). [https://doi.org/10.1016/0041-1647\(75\)90030-1](https://doi.org/10.1016/0041-1647(75)90030-1)

17. Lima & Associates: Lincoln Travel Demand Model. Lincoln Metropolitan Planning Organization (2011). <https://www.lincoln.ne.gov/files/sharedassets/public/planning/mpo/projects-amp-reports/tdm11.pdf>
18. MirHassani, S.A., Ebrazi, R.: A flexible reformulation of the refueling station location problem. *Transp. Sci.* **47**(4), 617–628 (2013)
19. Simchi-Levi, D., Berman, O.: A heuristic algorithm for the traveling salesman location problem on networks. *Oper. Res.* **36**(3), 478–484 (1988)
20. Transport Infrastructure Ireland: National Road Network (2013). <https://data.gov.ie/dataset/national-road-network-2013>
21. Transport Infrastructure Ireland: Traffic Count Data. <https://trafficdata.tii.ie/publicmultinodemap.asp>
22. Upchurch, C., Kuby, M., Lim, S.: A model for location of capacitated alternative-fuel stations. *Geogr. Anal.* **41**(1), 85–106 (2009)
23. Wang, Y.-W.: An optimal location choice model for recreation-oriented scooter recharge stations. *Transp. Res. Part D: Transport Environ.* **12**, 231–237 (2007). <https://doi.org/10.1016/j.trd.2007.02.002>
24. Wang, Y.-W., Lin, C.-C.: Locating multiple types of recharging stations for battery-powered electric vehicle transport. *Transp. Res. Part E: Logist. Transp. Rev.* **58**, 76–87 (2013)



Enhancing EVRP Benchmark Instances with Energy Estimates

Clíodhna Ní Shé^(✉) and Paula Carroll

School of Business, University College Dublin, Dublin, D04 V1W8, Ireland
cliodhna.nishe@ucdconnect.ie

Abstract. The increased use of electric vehicles (EVs) to deliver goods in the last mile of the delivery process means that there is a need for more realistic energy consumption models. Benchmark instances of the electric vehicle routing problem (EVRP) currently use an energy consumption model of one unit of energy for every one unit of distance travelled. The underestimation of energy consumption can lead to problems in the routes, in the life cycle and cost benefit analysis of EVs, and their demand and impact on the grid. We use a structured approach to enrich current benchmark instances, following the CRISP-DM data mining framework. This paper provides insights into explanatory variables that impact energy consumption, and provides a replicable methodology to create realistic benchmark instances using real world data.

Keywords: Energy Consumption · EVRP · Sustainable Transport · Data Enrichment

1 Introduction

A key environmental challenge is the need for cleaner and more sustainable energy sources to minimise emissions and pollution globally. A promising solution in the transport sector is the electrification of transportation. The Electric Vehicle Routing Problem (EVRP) is an extension of the Capacitated Vehicle Routing Problem (CVRP), a commonly studied problem in combinatorial optimisation. The EVRP consists of determining the optimal set of routes for a given number of Electric Vehicles (EVs) to service a given set of customers and their demand for goods or services. The objective of the optimisation can vary, for example it may be minimising distance, operational cost or energy consumption. Key constraints of the problem include the limited battery range of the EVs, the limited number of charging stations available and the capacity of the vehicles. The goal is to find routes that start and end at the depot, ensure all customers are visited exactly once, their demand serviced, and that the constraints are not violated, while also optimising the chosen objective. The EVRP is particularly relevant in last mile logistics, which refers to the last stretch of the supply chain of goods, from the last distribution centre to the recipients' preferred destination [16]. The underestimation of energy consumption can lead

to problems in the routes, and in the life cycle and cost benefit analysis of EVs and their demand and impact on the grid.

EVRP Benchmark instances are standardised datasets used to evaluate and compare the performance of algorithms. They provide a consistent framework to test models and algorithms, allowing for evaluation and benchmarking. Variants and extensions of the EVRP focus on different aspects such as time windows, battery degradation, nonlinear charging or battery swapping [9, 29]. However, for the purpose of this study, we will focus on the original variant, the classic EVRP. Benchmark instances provide the number of vehicles, the number of customers, the number of charging stations (CS), the capacity of the vehicles, the energy capacity of the battery, the energy consumption coefficient, the Euclidean coordinates of each node (customer, depot and CS) and the demand of each customer. The instance network is the set of nodes. Which are a set of locations given by two dimensional Euclidean coordinates, along with the edges (roads) between each pair of nodes. The weight of each edge is defined as the energy consumption along that edge. The current benchmark EVRP instances consider the energy consumed to depend solely on the distance travelled and the weight of the goods to be delivered. The energy consumption is modelled as a one-to-one relationship, such that one unit of energy is consumed for every one unit of distance travelled. This is a simplification that makes the problems easier and quicker to solve, but this is not a realistic assumption. The research gap is a lack of test instances that incorporate a realistic energy consumption model. We address the following research questions:

1. How realistic are current EVRP test instances?
2. How do ambient temperature, vehicle speed and road gradient affect EVRP energy consumption?
3. How can we adapt and enrich current EVRP test instances to take a realistic energy consumption model into account?

We transform the current instances from 2-dimensional coordinates to 3-dimensional coordinates and enrich the instances to include ambient temperature, speed and road gradient components. Our methodology is replicable. We provide realistic test datasets that can be used to evaluate the performance of EVRP models and algorithms, taking energy consumption into account. This leads to more accurate EV energy consumption, leading to more accurate analysis of EV routes, EV life cycle and cost benefit analysis and their demand and impact on the grid.

We create 12 different scenarios for each benchmark instance. We use Ireland as a case study, and focus on two locations in Ireland to enrich the test instances. We overlay them on a geographical location to project the 2D instances onto 3D surface where gradient impact can be explored. The two locations are Cork and Dublin, as they are the two most populated cities in the Republic of Ireland [26]. Dublin is more densely populated than Cork, and Cork is hillier compared to Dublin. For both locations, we take four different temperature scenarios, three different speed scenarios and we add gradient data. We evaluate the impact that this adaption has on the weights of the edges. We show that adding the gradient

has the biggest impact on the edge weight distribution, compared to the other explanatory variables, ambient temperature and speed. Our main contribution is the framework to enhance and enrich current benchmark instances.

The remainder of the paper is structured as follows, Sect. 2 provides a brief literature review on EVRP energy consumption models. In Sect. 3 we provide a comprehensive description of our methodology, how we gathered, prepared and modelled the data. In Sect. 4 we explore the characteristics of the distribution of the energy consumption of the new instances. Section 5 concludes the paper.

2 Literature Review

EVs need to be charged at charging stations due to their limited driving ranges [11]. Decisions about when and where to charge depend on how energy is consumed along a journey. Erdogan and Miller-Hooks [6] and Schneider et al. [22] assume that energy consumption is a linear function of distance travelled. This assumption is common in EVRP literature but is not a realistic assumption. Goeke and Schneider [8] derive a mechanical power energy consumption model as a function of the road gradient, the speed of the vehicle and the vehicle parameters. This mechanical power model is used by many, for example [10, 12, 13, 25, 27]. We show the description of each parameter in Table 1. The battery energy b_{ij} consumed between node i and j is a function of the power P_{ij} , the time taken t_{ij} to travel from node i to j , and the efficiency parameters φ and ϕ and is seen in Eq. (1):

$$b_{ij} = \begin{cases} \phi^d \cdot \varphi^d \cdot P_{ij} \cdot t_{ij} & \text{if } P_{ij} \geq 0 \text{ kW,} \\ \phi^r \cdot \varphi^r \cdot P_{ij} \cdot t_{ij} & \text{if } P_{ij} < 0 \text{ kW.} \end{cases} \quad (1)$$

where

$$P_{ij} = \left(\frac{1}{2} \cdot C_d \cdot \rho \cdot A \cdot v_{ij}^2 + m \cdot g \cdot (\sin(\alpha_{ij}) + C_r \cdot \cos(\alpha_{ij})) \right) \cdot v_{ij} \quad (2)$$

The air density is assumed to be constant for simplicity. Acceleration phases are disregarded in this model and the average speed across the distance is used. Other explanatory variables that affect the driving range are the use of air conditioning or heating systems, which, according to Restrepo et al. [20] can reduce the original energy available by up to 30%. Low temperatures decrease the battery efficiency and cause performance losses. Yuksel and Michalek [28] show that, compared to mild climate regions, energy consumption of EVs can rise in warmer or lower temperature climates, which results in up to a 41% decrease in the driving range. Focusing on ambient temperature alone, Rastani et al. [18] show that energy consumption in an EVRP can increase by up to 68% due to ambient temperature. Rastani et al. [18] derive the ambient temperature coefficient function from empirical data from a study performed on a Nissan Leaf. The ambient temperature coefficient function for the Nissan Leaf that depends on the ambient temperature T (in $^{\circ}\text{C}$) is in Eq. (3)

$$h_{LEAF}(T) = \begin{cases} 0.3392 - 0.005238T - 0.0001078T^2 + 1.047 \times 10^{-5}T^3 + \\ 3.955 \times 10^{-7}T^4 - 1.362 \times 10^{-8}T^5 - 3.109 \times 10^{-10}T^6 & \text{if } T < 22^{\circ}\text{C,} \\ 0.4211 - 0.01627T + 0.0004229T^2 & \text{if } T \geq 22^{\circ}\text{C.} \end{cases} \quad (3)$$

Table 1. Parameters in the Energy Consumption Model

Notation	Value	Description	Unit	Source
g	9.81	Gravitational constant	m/s^2	Demir et al. [4]
ρ	1.2041	Air density	kg/m^3	Demir et al. [4]
A	3.467	Frontal surface area	m^2	Renault [19]
m	1622	Curb mass	kg	Renault [19]
C_r	0.01	Coefficient of rolling resistance	—	Demir et al. [4]
C_d	0.7	Coefficient of aerodynamic drag	—	Demir et al. [4]
ϕ_d	1.184692	Efficiency parameter (motor mode)	—	Goeke & Schneider [8]
ϕ_r	0.846055	Efficiency parameter (generator mode)	—	Goeke & Schneider [8]
φ_d	1.112434	Efficiency parameter (discharging)	—	Goeke & Schneider [8]
φ_r	0.928465	Efficiency parameter (recuperation)	—	Goeke & Schneider [8]
t_{ij}	—	Time to traverse edge e_{ij}	s	—
v_{ij}	—	Velocity on edge e_{ij}	m/s	—
α_{ij}	—	Gradient of edge e_{ij}	—	—

They normalise the function as follows:

$$AT(T) = \frac{h_{LEAF}(T)}{h_{LEAF}(22^\circ\text{C})} \quad (4)$$

The instances that algorithms are currently widely tested on, are Schneider et al. [22] and Mavrovouniotis et al. [14]. The instances by Schneider et al. (2014) are comprised of 56 what they call large instances, each with 100 customers and 21 recharging stations and 36 small instances, each with 5, 10 or 15 customers per instance. Mavrovouniotis et al. [14] convert classic Capacitated VRP (CVRP) instances to EVRP instances by introducing information about the EVs and charging stations in the form of EV energy capacity, charging stations and an energy consumption coefficient, with instances ranging from 29 to 1006 nodes. There are four sets of instances, generated from the popular CVRP instances of Christofides and Eilon [2], Christofides et al. [3], and Fisher [7], and the more recent instances of Uchua et al. [24]. The instances with more than 150 customers are considered large instances, in line with the literature surrounding the generation of the instances [2,3,7,24].

Our literature review shows the gap between popular EVRP test instances and realistic energy consumption. We contribute a realistic energy consumption model by combining the mechanical power energy model with the ambient temperature model.

A measure of the distance between distributions is needed to evaluate the effect of the energy consumption explanatory variables. The Kullback–Leibler (KL) divergence measure quantifies the relative entropy or difference in information represented by two distributions [1]. It is a widely used method for measuring the fit of two distributions. A KL divergence of 0 indicates that the distributions are very similar, $0 < KL < 0.5$ indicates that the distributions are quite close with some slight differences, $0.5 < KL < 1$ indicates that there are moderate differences and $KL > 1$ signifies significant differences between the distributions.

3 Methodology

We use the CRISP-DM framework to implement a workflow to enrich the current EVRP benchmark instances by Mavrovouniotis et al. [14]. This framework breaks the workflow into these major phases; business understanding, data understanding, data preparation, modelling, evaluation and deployment [23]. The business understanding goals are the research questions outlined in Sect. 1. We use the mechanical power model in Eq. (1), along with the ambient temperature coefficient function in Eq. (3) for the energy consumption model. Thus, the final energy consumption (EC) function, as a function of the batter energy b_{ij} , the road gradient α_{ij} , the distance d_{ij} and the ambient temperature T is:

$$EC_{ij} = b_{ij}(v_{ij}, \alpha_{ij}, d_{ij}) \cdot AT(T) \quad (5)$$

3.1 Data Understanding

We collect and analyse the data in Table 1. This data is vehicle parameters, which we extracted from the literature and from the Renault website [19]; temperature data from Met Éireann (the Irish National Meteorological Service) [15]; speed data from the RSA (Irish Road Safety Authority) [21]; and the road elevation data, from the Open-Elevation API [17].

Vehicle Parameters. For simplicity, we assume that every instance has a fleet of the same vehicle. These vehicle parameters can be easily modified, to tailor the instances if a specific case study is being carried out. We take the parameters from Goeke and Schneider [8], in Table 1. We use the Renault Kangoo, that is used by An Post (Ireland's postal service). The parameters for the Renault Kangoo are the curb mass 1622 kg, the frontal surface area and the battery capacity 44 kWh [19]. We will calculate the frontal surface area in Sect. 3.2.

Temperature. We obtain temperature data from the Met Éireann website, the full hourly data series from the Pheonix Park location in Dublin, and from Cork Airport in Cork [15]. The Pheonix Park weather station has weather data from 16/08/2003 until 01/02/2025. The Cork airport weather station has data from 01/01/1962 until 01/02/2025.

Speed. We build upon the benchmark instances to incorporate a speed limit to each edge (road) in the network. We divide the edges in the network into three categories based on the Irish road network: regional roads, local roads and urban area roads. The speed limits on these roads are 80 km/h, 60 km/h and 50 km/h respectively [21]. The average speed across edges for the energy consumption model depends on the state of the traffic. We assume three scenarios: rush hour, daytime and nighttime. The assignment of road categories for each instance will be discussed in Sect. 3.2.

Elevation. To transform the instances from 2-dimensions to 3-dimensions, we source elevation data from the Open-Elevation Application Programming Interface (API) [17]. We overlay the points in the benchmark instances on real world locations. For the small instances, we use our case study locations Cork and Dublin. Figure 1 shows the nodes from the instance E-n60-k5-s9 (60 nodes, 5 EVs and 9 charging stations) and overlaid on Cork City. For the larger instances, we choose a point in mainland Europe as Ireland is an island and the instance span is larger than the island. We use the elevation data to calculate the net estimate road gradient between nodes in the instance, as explained in Sect. 3.2.

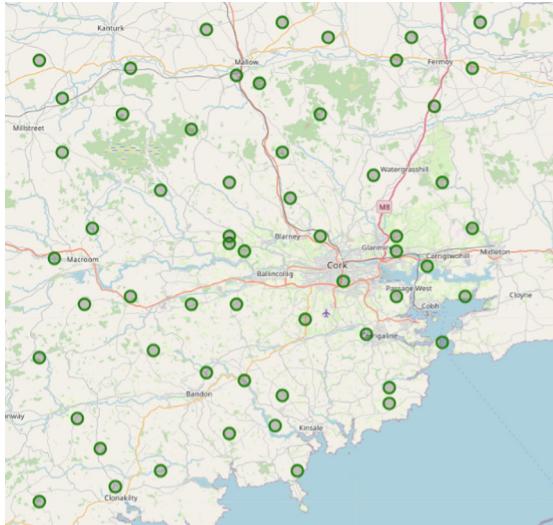


Fig. 1. Visualization of the EVRP instance E-n60-k5-s9 overlaid on the geographic area of Cork City, Ireland.

3.2 Data Preparation

Vehicle Parameters. The vehicle parameters are taken from Table 1 or from the Renault website, except for the frontal surface area. The frontal surface can be computed as frontal Area = Overall Width \times Overall Height. The width and height is available on the Renault website [19]. Therefore the frontal area of the Renault Kangoo = $1.86 \times 1.864 = 3.467 \text{ m}^2$.

Temperature. We clean both the Cork and Dublin datasets by removing any empty datapoints and cut both datasets to contain the five most recent years of temperature data from 2019–2024, to account for the upward trend in temperatures driven by climate change [5]. We define four different average temperature scenarios by the time of day (night/day) and the season (Winter, Summer). The

histograms in Fig. 2 illustrate the differences between Summer and Winter temperatures, as well as between night and day, justifying our scenarios. To obtain a temperature value for the four different scenarios, we calculate the mean temperature across 2019–2024 for Dublin and Cork, as shown in Table 2.

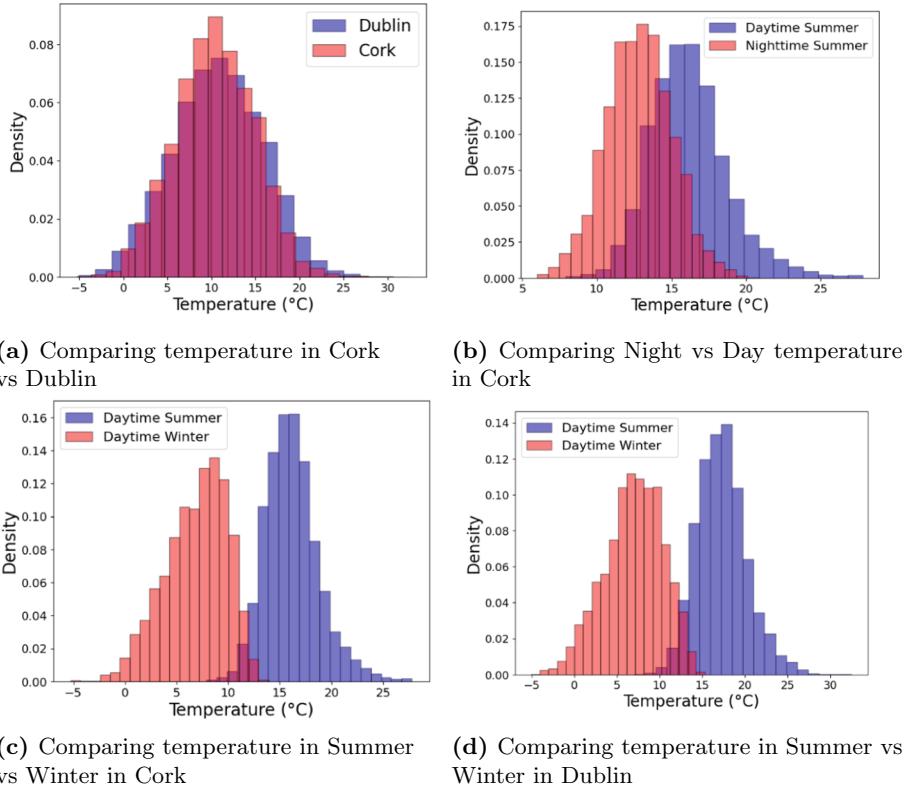


Fig. 2. Comparison of Temperature distributions for Cork and Dublin in Winter and Summer and Night and Day

Speed. As mentioned in Sect. 3.1, we enrich the edges from the benchmark instances to include a speed scenario. The edges in the network are assigned a road category depending on the distance the furthest endpoint of the edge (road) is from the centroid, as seen in Algorithm 1. The centroid of the nodes is the geometric centre, the average position of all of the nodes. A visual example of the strategy for a small instance (E-n60-k5-s9) is shown in Fig. 3. The average speed across the edges depends on external factors such as traffic, junctions and traffic lights. Therefore, we define three typical speed scenarios: rush hour, day (not rush hour), and night (only traffic lights and junctions slowing the vehicles down).

Table 2. Mean temperature for different seasons and times of day in Dublin and Cork (°C)

City	Season	Winter (°C)	Summer (°C)
Dublin	Daytime	6.93	17.26
	Night	5.68	13.06
Cork	Daytime	6.86	16.14
	Night	5.89	12.81

Table 3. Speed scenarios Based on Time Period

Time Period	Average Speed (percentage of speed limit)
Rush Hour	50%
Day (Not Rush Hour)	70%
Night	90%

Gradient. Using the elevation data we calculate and store the gradient (elevation change per unit distance) between each i and j as follows: $\text{gradient}_{ij} = \frac{\text{elevation}_j - \text{elevation}_i}{\text{distance}_{ij}}$.

3.3 Modelling

We transform the weight of each edge in the network to represent the energy consumed along that edge, using the formula EC_{ij} in Eq. (5). With the adapted edge weight value, we can view the problem in terms of energy consumption, as opposed to distance travelled. We set up the scenarios in Table 4. The key will be used as a label in figures to indicate the corresponding scenario. For example, the key C_SDR represents the scenario in Cork during Summer daytime rush hour.

Algorithm 1. Road Category Assignment

```

1: Compute centroid  $c$ 
2: Compute  $rad \leftarrow \max_i \text{dist}(c, i)/3$ 
3: for each node pair  $(i, j)$  do
4:    $d \leftarrow \max(\text{dist}(c, i), \text{dist}(c, j))$ 
5:   if  $d < rad$  then
6:     Assign local road
7:   else if  $d < 2 \cdot rad$  then
8:     Assign secondary road
9:   else
10:    Assign regional road
11:   end if
12: end for

```

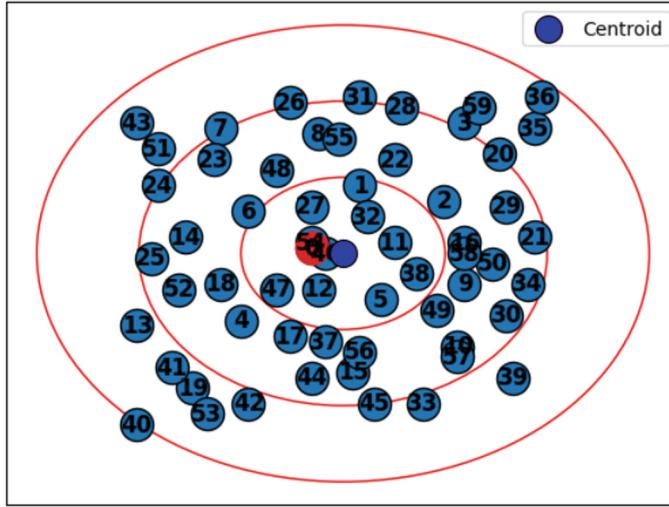


Fig. 3. Visual representation of the road category assignment (instance E-n60-k5-s9). The numbers represent customers and charging stations and the red circle is the depot. Urban Area Road within the smallest circle, Local Road within the second circle, and Regional Road within the largest circle. (Color figure online)

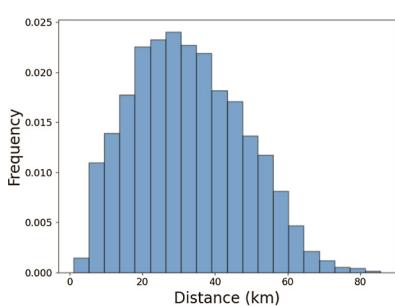
3.4 Evaluation

We evaluate the enriched benchmark instances by comparing the edge weight values in the network. We demonstrate the analysis on two benchmark instances, a small 60 node instance with 5 EVs and 9 charging stations (E-n60-k5-s9) and a large 577 node instance with 30 EVs and 4 charging stations (X-n577-k30-s4), for brevity in this paper. We chose these instances as they represent a small instance and a large instance, and we found similar results in the rest of the instances. The distribution of the edge weights for the two original instances, where EC_{ij} is dependent only on distance travelled are shown in Fig. 4. For the large instance (Fig. 4b), the distribution is more right skewed than for the small instance (Fig. 4a). The large instances (with the prefix X) derive from the CVRP instances created by Uchoa et al. [24]. These Uchoa instances are more computationally challenging as CVRP instances, compared to the other CVRP instances due to the problem size.

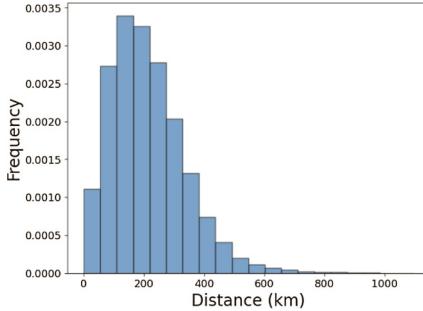
We explore the distribution of the edge weights under varying ambient temperature conditions, keeping the road gradient uniform (0 rad) and the speed situation uniform (during the day - 70% of the speed limit). We compare the four Dublin ambient temperature scenarios from Table 2, (Summer, Winter, night time and daytime) by plotting the distributions of the energy consumption. Comparing these distributions visually, for both the small and large instances, there is a slight difference in the histograms in Fig. 5. The ambient temperature factor of the energy consumption is a multiplicative factor, and therefore will only change the magnitude of the distribution, not the shape.

Table 4. Scenarios based on location, season, time of day, and speed category.

Location	Season	Time of Day	Speed Category	Key
Cork	Summer	Day	Rush Hour	C_SDR
Cork	Summer	Day	Day	C_SDD
Cork	Summer	Night	Night	C_SNN
Cork	Winter	Day	Rush Hour	C_WDR
Cork	Winter	Day	Day	C_WDD
Cork	Winter	Night	Night	C_WNN
Dublin	Summer	Day	Rush Hour	D_SDR
Dublin	Summer	Day	Day	D_SDD
Dublin	Summer	Night	Night	D_SNN
Dublin	Winter	Day	Rush Hour	D_WDR
Dublin	Winter	Day	Day	D_WDD
Dublin	Winter	Night	Night	D_WNN



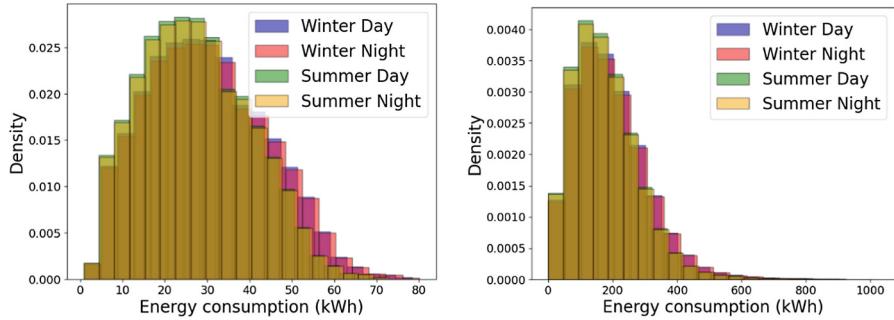
(a) Edge weights for small instance (E-n60-k5-s9).



(b) Edge weights for large instance (X-n577-k30-s4).

Fig. 4. Edge weight distributions - distance (km)

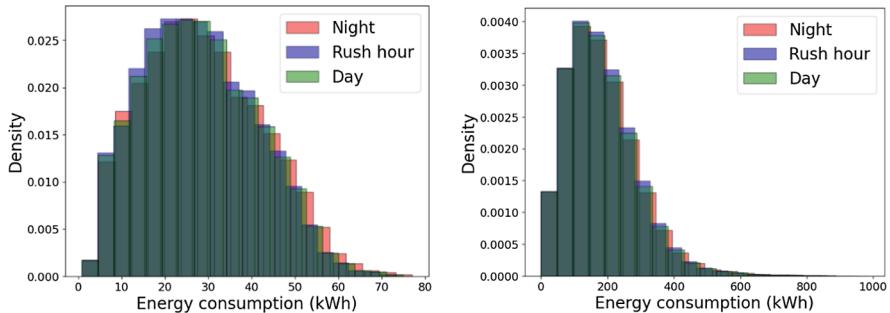
We compare the distribution of the edge weights under varying speed conditions, keeping the road gradient uniform (0 rad) and the ambient temperature uniform (yearly mean temperature of 10.23 ° C). We represent the edge weights, which correspond to energy consumption, for both the small and the large instance in the histograms in Fig. 6. We compare the impact of the gradient by comparing the distribution of edge weights when the explanatory variable gradient is incorporated compared to when there is a fixed gradient. We keep speed setting (during the day - 70 % of the speed limit) and the ambient temperature (yearly mean temperature in Dublin of 10.23 ° C) uniform. We show the energy adapted edge weights, for both the small and the large instance in the histograms in Fig. 7. In both figures, for the small instance and the large instance, the shape of the distribution of the edges is differ when there is a real



(a) Edge weights for small instance (E-n60-k5-s9). (b) Edge weights for large instance (X-n577-k30-s4).

Fig. 5. Energy Edge weight distributions - Varying temperatures.

world gradient included, compared to a fixed gradient. The edges with negative energy consumption is due to regenerative braking occurring when the road gradient is less than 0 rad and the EV is travelling downhill.

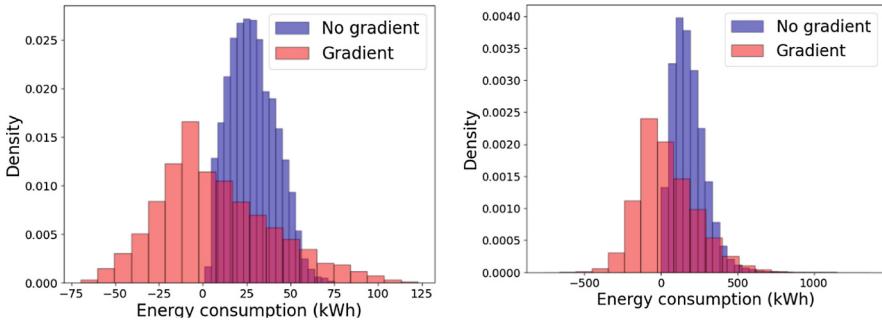


(a) Edge weights for small instance (E-n60-k5-s9). (b) Edge weights for large instance (X-n577-k30-s4).

Fig. 6. Energy Edge weight distributions - Varying speeds.

3.5 Deployment

To apply this methodology, we associate each instance with a specific location, and corresponding speed and ambient temperature values. Existing benchmark instances can be reused across different gradient locations, speeds, and ambient temperatures.



(a) Edge weights for small instance (E-n60-k5-s9). (b) Edge weights for large instance (X-n577-k30-s4)

Fig. 7. Energy Edge weight distributions - Gradient vs no Gradient.

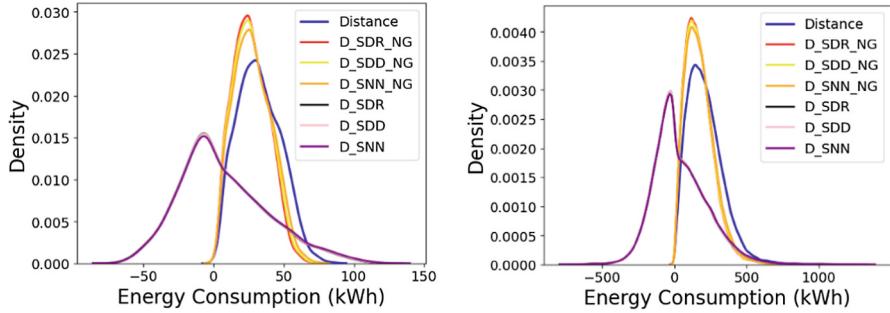
4 Results

In this section, we analyse the results by fitting curves to the original and transformed edge weights for the different scenarios. We compare them both visually and using KL to quantify the difference between the curves. The KL divergence measure between the curves for ambient temperature in Fig. 6 and for speed Fig. 5 are ≈ 0 which confirms that there is no significant effect both explanatory variables.

The curves comparing the gradient cases in both Fig. 7a and Fig. 7b, have greater differences than the previous comparisons. In particular, for the large instance, the KL value $KL(G||NG) = 7.76$ (where G is the distribution with the gradient and NG is without the gradient). This is a measure of how much information is lost when using distribution NG to approximate G. A KL value of 7.76 shows that there is a significant difference between distributions.

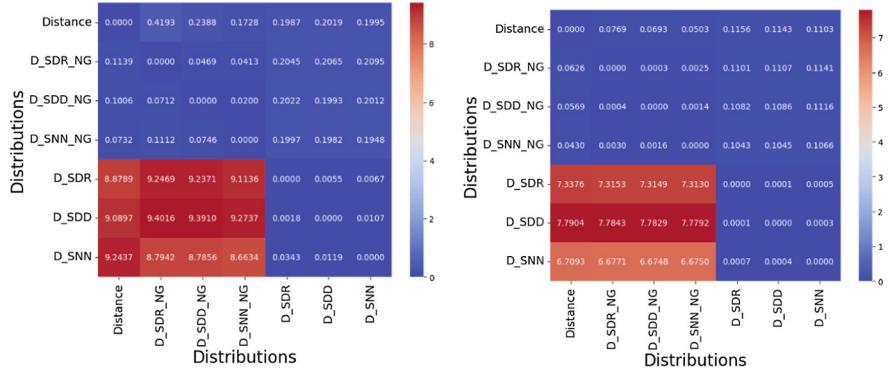
As a final example we show the scenarios for the Summer in Dublin from Table 4, with the key `_NG` added when we left the gradient as constant (0.5 rad) for the two benchmark instances. We plot the curves that are fit to the histograms of the different scenarios in Fig. 8 to compare the distributions for both the small and large instances. The distributions are visually different for the edge weights when road gradient is taken into account. A heat coloured KL divergence matrix for each instance is shown for both the small instance in Fig. 9a and the large instance in Fig. 9b. These matrices clearly show that the distribution of the energy consumption for the edges in the benchmark instances changes when the road gradient is taken into account. The KL divergence measures are larger for the smaller benchmark instance (E-n60-k5-s9) compared to the larger instance (X-n577-k30-s4). We have performed this comparison for all of the enriched Mavrovounotis et al. [14] instances. The results are shown in Table 5.

The introduction of the road gradient to the energy consumption model transforms the distribution of the edges in a similar way across all of the instances. The distribution becomes more right skewed, and there is a higher variance.



(a) Edge weights distributions for small instance (E-n60-k5-s9). (b) Edge weights distributions for large instance (X-n577-k30-s4).

Fig. 8. Comparing Energy Edge weight distributions.



(a) KL divergence for distributions for small instance (E-n60-k5-s9). (b) KL divergence for distributions for large instance (X-n577-k30-s4).

Fig. 9. Comparing Energy Edge weight distributions - KL Divergence.

5 Discussion and Conclusion

This paper provides answers to the three research questions that we introduced in Sect. 1. We show that the distribution of energy consumption across edge weights in the enriched benchmark instances is significantly different than the original unmodified benchmark instances. We have shown that the road gradient is the explanatory variable that has the greatest effect on the energy consumption. We have successfully adapted and enriched the current EVRP benchmark instances with real world data.

A limitation of this study is that the values in Table 3 are based on simplified scenarios. Future work will include more realistic speed scenarios derived from real world traffic data, and more realistic road category assignment.

Table 5. KL values for each EVRP instance when comparing the energy consumption of the edges without the road gradient and the edges with road gradient.

Instance	KL Value
X-n147-k7-s4	7.543
X-n221-k11-s9	7.226
X-n360-k40-s9	7.082
X-n469-k26-s10	7.253
X-n577-k30-s4	7.023
X-n698-k75-s13	7.571
X-n759-k98-s10	7.613
X-n830-k171-s11	7.942
X-n920-k207-s4	7.691
X-n1006-k43-s5	7.806

We provide a replicable framework that can be followed to transform current benchmark instances without a realistic energy consumption model, to benchmark instances with an energy consumption model that depends on speed, ambient temperature and road gradient. This enables better estimation of the energy consumption of EVs in EVRP problems. The change in edge weights, where the weight now represents real-world energy consumption, will impact the current solutions to EVRPs and potentially impact life cycle assessment and cost benefit analysis of EVs. This could potentially have impacts on the electrification of transport and opens up opportunities to explore the intersection of electrification of transport and the electricity grid. These enriched benchmark instances can be used as realistic scenarios for evaluating and comparing the performance of algorithms and heuristics designed to solve the EVRPs.

These adapted instances can also be used to compare solutions to EVRPs in different scenarios, for decision making on when and where to deliver goods. Companies may make fleet planning decisions based on the routes they obtain from solving the instances for different vehicles, speed scenarios, ambient temperature or gradient. More realistic energy consumption estimation leads to more accurate predictions in electrical needs, which is imperative as we move towards electrifying transport. These enriched instances are useful for the successful electrification of transport, helping to mitigate climate change.

Acknowledgements. The research conducted in this publication was funded by Research Ireland and co-funding partners under grant number 21/SPP/3756 through the NexSys Strategic Partnership Programme.

References

1. Belov, D.I., Armstrong, R.D.: Distributions of the Kullback–Leibler divergence with applications. *British J. Math. Stat. Psychol.* **64**(2), 291–309 (2011). <https://bpspsychub.onlinelibrary.wiley.com/doi/abs/10.1348/000711010X522227>
2. Christofides, N., Eilon, S.: An algorithm for the vehicle-dispatching problem. *OR* **20**(3), 309–318 (1969). <http://www.jstor.org/stable/3008733>
3. Christofides, N., Mingozzi, A., Toth, P.: Exact algorithms for the vehicle routing problem, based on spanning tree and shortest path relaxations. *Math. Program.* **20**(1), 255–282 (1981). <https://doi.org/10.1007/bf01589353>
4. Demir, E., Bektaş, T., Laporte, G.: An adaptive large neighborhood search heuristic for the pollution-routing problem. *Eur. J. Oper. Res.* **223**(2), 346–359 (2012). <https://doi.org/10.1016/j.ejor.2012.06.044>. <https://www.sciencedirect.com/science/article/pii/S0377221712004997>
5. Environmental Protection Agency: What impact will climate change have on Ireland? <https://www.epa.ie/environment-and-you/climate-change/what-impact-will-climate-change-have-for-ireland/>
6. Erdogan, S., Miller-Hooks, E.: A green vehicle routing problem. *Transpor. Res. Part E Logist. Transp. Rev.* **48**(1), 100–114 (2012). Select Papers from the 19th International Symposium on Transportation and Traffic Theory. <https://doi.org/10.1016/j.tre.2011.08.001>
7. Fisher, M.L.: Optimal solution of vehicle routing problems using minimum K-trees. *Oper. Res.* **42**(4), 626–642 (1994)
8. Goeke, D., Schneider, M.: Routing a mixed fleet of electric and conventional vehicles. *Eur. J. Oper. Res.* **245**(1), 81–99 (2015). <https://doi.org/10.1016/j.ejor.2015.01.049>
9. Jie, W., Yang, J., Zhang, M., Huang, Y.: The two-echelon capacitated electric vehicle routing problem with battery swapping stations: formulation and efficient methodology. *Eur. J. Oper. Res.* **272**(3), 879–904 (2019). <https://doi.org/10.1016/j.ejor.2018.07.002>
10. Kancharla, S.R., Ramadurai, G.: Electric vehicle routing problem with non-linear charging and load-dependent discharging. *Exp. Syst. Appl.* **160**, 113714 (2020). <https://doi.org/10.1016/j.eswa.2020.113714>
11. Keskin, M., Laporte, G., Çatay, B.: Electric vehicle routing problem with time-dependent waiting times at recharging stations. *Comput. Oper. Res.* **107**, 77–94 (2019). <https://doi.org/10.1016/j.cor.2019.02.014>
12. Lera-Romero, G., Miranda Bront, J.J., Soulignac, F.J.: A branch-cut-and-price algorithm for the time-dependent electric vehicle routing problem with time windows. *Eur. J. Oper. Res.* **312**(3), 978–995 (2024). <https://doi.org/10.1016/j.ejor.2023.06.037>
13. Longhitano, P.D., Bérenguer, C., Echard, B.: Joint electric vehicle routing and battery health management integrating an explicit state of charge model. *Comput. Ind. Eng.* **188**, 109892 (2024). <https://doi.org/10.1016/j.cie.2024.109892>
14. Mavrovouniotis, M., Menelaou, C., Timotheou, S., Ellinas, G., Panayiotou, C., Polycarpou, M.: A benchmark test suite for the electric capacitated vehicle routing problem. In: 2020 IEEE Congress on Evolutionary Computation (CEC), pp. 1–8. IEEE Press (2020). <https://doi.org/10.1109/CEC48606.2020.9185753>
15. Met Éireann: Met Éireann - The Irish Meteorological Service (2025). <https://www.met.ie/climate/available-data/>. Accessed 09 Jun 2025

16. Olsson, J., Hellström, D., Pålsson, H.: Framework of last mile logistics research: a systematic review of the literature. *Sustainability* **11**(24) (2019). <https://doi.org/10.3390/su11247131>
17. Open Elevation Project: Open Elevation API (2025). <https://open-elevation.com/>. Accessed 10 May 2025
18. Rastani, S., Yüksel, T., Çatay, B.: Effects of ambient temperature on the route planning of electric freight vehicles. *Transp. Res. Part D: Transp. Environ.* **74**, 124–141 (2019). <https://doi.org/10.1016/j.trd.2019.07.025>
19. Renault: Renault Kangoo Electric for Sale | E-Tech | London — loadsofvans.com (2025). <https://www.loadsofvans.com/offers/renault-kangoo-e-tech-mwb>. Accessed 09 May 2025
20. Restrepo, J., Rosero, J., Tellez, S.: Performance testing of electric vehicles on operating conditions in Bogotá DC, Colombia. In: 2014 IEEE PES Transmission & Distribution Conference and Exposition - Latin America (PES T&D-LA), pp. 1–8 (2014). <https://doi.org/10.1109/TDC-LA.2014.6955276>
21. RSA: Advice for using the roads in Ireland (2025). <https://www.rsa.ie/road-safety/roadusers/tourists>. Accessed 4 May 2025
22. Schneider, M., Stenger, A., Goeke, D.: The electric vehicle-routing problem with time windows and recharging stations. *Transp. Sci.* **48**(4), 500–520 (2014). <http://www.jstor.org/stable/43666939>
23. Shearer, C.: The CRISP-DM model: the new blueprint for data mining. *J. Data Warehouse*. **5**(4), 13–22 (2000)
24. Uchoa, E., Pecin, D., Pessoa, A., Poggi, M., Vidal, T., Subramanian, A.: New benchmark instances for the capacitated vehicle routing problem. *Eur. J. Oper. Res.* **257**(3), 845–858 (2017). <https://www.sciencedirect.com/science/article/pii/S0377221716306270>
25. Wang, Y., Zhou, J., Sun, Y., Fan, J., Wang, Z., Wang, H.: Collaborative multidepot electric vehicle routing problem with time windows and shared charging stations. *Exp. Syst. Appl.* **219**, 119654 (2023). <https://doi.org/10.1016/j.eswa.2023.119654>
26. World Population Review: Ireland population 2024 (2024). <https://worldpopulationreview.com/countries/ireland#density-by-city>. Accessed 11 June 2025
27. Yu, V.F., Jodiawan, P., Gunawan, A.: An adaptive large neighborhood search for the green mixed fleet vehicle routing problem with realistic energy consumption and partial recharges. *Appl. Soft Comput.* **105**, 107251 (2021). <https://doi.org/10.1016/j.asoc.2021.107251>
28. Yuksel, T., Michalek, J.J.: Effects of regional temperature on electric vehicle efficiency, range, and emissions in the united states. *Environ. Sci. Technol.* **49**(6), 3974–3980 (2015). pMID: 25671586. <https://doi.org/10.1021/es505621s>
29. Zhang, S., Chen, M., Zhang, W.: A novel location-routing problem in electric vehicle transportation with stochastic demands. *J. Clean. Prod.* **221**, 567–581 (2019). <https://doi.org/10.1016/j.jclepro.2019.02.167>



Economic and Environmental Benefits of Centralized MILP Optimization of EV Fleet Charging

Lucija Hajsok^(✉)  and Tea Žakula 

Faculty of Mechanical Engineering and Naval Architecture, University of Zagreb,
Ivana Lucića 5, 10000 Zagreb, Croatia
lhajsok@fsb.hr

Abstract. With the goal of achieving climate neutrality, many European countries are accelerating the electrification of the transport sector—one of the major contributors to CO₂ emissions. To reduce costs and improve sustainability, they are increasingly integrating electric vehicle (EV) fleets into the power grid. However, this integration introduces challenges such as increased pressure on grid stability, making coordinated charging strategies essential to fully realize the benefits of EV deployment. To address the EV fleet charging scheduling problem, this study develops a centralized optimization model based on Mixed-Integer Linear Programming (MILP), a widely adopted and flexible approach. While existing research typically focuses either on cost or emissions, and often within the scope of a single-country case study, this paper presents a comparative analysis of both cost and sustainability outcomes across five European countries: Croatia, France, Germany, Sweden, and Poland. By evaluating smart charging potential under uniform operational assumptions, the study highlights how the effectiveness of EV integration varies significantly depending on national electricity market structures and energy mixes. The results provide valuable insights for policymakers, energy planners, and fleet operators, demonstrating the importance of tailoring EV strategies to specific national contexts.

Keywords: EV fleet management · EV integration · EV charging optimization · demand response · cross-country analysis · cost analysis · sustainability analysis

1 Introduction

The global electric vehicle (EV) stock is projected to reach just over one-fourth of all vehicles on the road by 2035 [1]. This rapid growth is driven by a combination of supportive government policies, declining battery costs, improved charging infrastructure, and a growing consumer preference for sustainable transportation. The European Union (EU) plays a central role in driving this transformation. Under the European Green Deal and the “Fit for 55” legislative package, the EU aims to reduce greenhouse gas (GHG) emissions by at least 55% by 2030 compared to 1990 levels, and achieve climate neutrality by 2050 [2]. In the transport sector, the EU has mandated a 100% reduction in

CO₂ emissions from new cars and vans by 2035, effectively signaling the end of internal combustion engine vehicle sales across member states [3]. These goals highlight the strategic importance of electrification and the decarbonization of both vehicle fleets and the electricity generation sector. Additionally, the integration of smart technologies and coordinated EV charging strategies can significantly contribute to energy system sustainability and cost savings for consumers, while reducing strain on the grid and supporting the integration of renewable energy sources.

The path to large-scale electrification and decarbonization often begins with small, everyday decisions—such as how and when electric vehicles are charged at home. At the household level, there is a growing awareness among consumers about the importance of energy efficiency and cost savings [4–7]. In many countries, consumers benefit from dual-tariff electricity pricing, where lower electricity rates apply during nighttime. While EV owners may not always be fully aware of this, charging vehicles during off-peak hours contributes to grid stability and more efficient energy market participation. When this concept is scaled up—for example, assuming an owner manages a fleet of 1 000 EVs—the opportunity becomes even more significant. Moreover, in scenarios in which electricity prices fluctuate, e.g. on hourly or even shorter time scale, fleet owners can adopt strategic charging plans to minimize costs and maximize economic benefits. However, it is important to distinguish between electricity pricing and the carbon intensity of the power supply, as the two do not always align. In some countries, the lowest electricity prices coincide with periods of high renewable energy generation, resulting in both low costs and low CO₂ emissions. In other regions, however, low-cost electricity may be driven by fossil-fuel-based baseload generation, leading to higher emissions even during low-price periods. As a result, while cost-optimized charging can bring economic gains, its environmental impact is highly context-dependent. Fleet owners aiming to reduce both costs and carbon emissions must therefore consider the specific electricity market dynamics and generation mix of their region.

Besides pricing and emissions, the EV optimization framework and the choice of optimization method significantly influence both cost-effectiveness and environmental outcomes of electric vehicle integration into the energy system. Various optimization frameworks have been proposed to manage electric vehicle (EV) charging, each with different levels of control and scalability. These include centralized models, in which a single controller optimizes charging for the entire fleet [8–11]; decentralized approaches, which distribute control among individual units or local agents [12–16] and hybrid methods, which combine elements of both centralized and decentralized control to enable multi-stage or hierarchical coordination [17–21]. This study adopts a centralized optimization framework, which is well-suited for small- to medium-scale systems. It provides globally optimal solutions with low implementation complexity, making it an effective and practical choice for comparative analysis across electricity markets.

Once an optimization framework is selected, an appropriate optimization method is required to determine the most efficient charging strategy. In this study, Mixed-Integer Linear Programming (MILP) is selected due to its flexibility and proven effectiveness in EV fleet scheduling. While other techniques such as Linear Programming (LP), Model Predictive Control (MPC), and heuristic methods exist, MILP offers a balance between modeling precision and computational feasibility, making it ideal for the comparative,

multi-country analysis conducted in this paper. MILP is particularly well-suited for centralized models, as it allows simultaneous handling of continuous variables (e.g., energy delivered per interval) and binary decisions (e.g., charging ON/OFF status) [8, 22–24]. The optimization is performed over a one-day horizon using a 15-min time resolution, enabling detailed modeling of charging behavior and energy market dynamics.

Despite the growing importance of sustainability, most existing studies prioritize total cost minimization as the primary objective. As a result, CO₂ emissions are frequently overlooked or treated as a secondary objective. This study aims to emphasize that a single-objective, cost-focused approach may not yield the best results. To address this, and to highlight the potential need for multi-objective optimization frameworks that balance both cost and sustainability considerations, a two-step methodological approach is adopted. First, a cost-only optimization is performed to determine charging schedules with the lowest daily cost. In the second step, the resulting CO₂ emissions are analyzed to evaluate the environmental consequences of the cost-driven decisions.

Another research gap in the literature is the predominant focus on country-specific studies, which reduces the generalizability of findings. Given that each energy market has distinct characteristics—such as energy mix, regulatory frameworks, and price dynamics—results from one national context may not be directly applicable to others. This study addresses that gap by evaluating the potential cost savings and sustainability impact of EV charging optimization across five European countries—Croatia, France, Germany, Poland, and Sweden—under identical operational assumptions. The countries were selected based on their distinct energy profiles: Croatia, a small country with a mixed energy supply and dependence on imports; Germany, with significant share of renewables but continued reliance on fossil fuels; France, dominated by nuclear energy; Poland, with one of the highest carbon intensities in electricity generation among European countries; and Sweden, featuring one of the cleanest electricity mixes in Europe (Fig. 1).

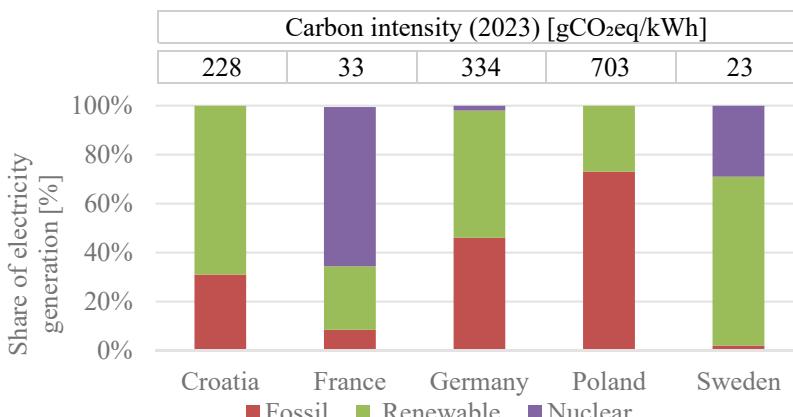


Fig. 1. Electricity generation from fossil fuels, nuclear and renewables [25] and associated carbon intensity in selected European countries for 2023 [26]

The results highlight how differences in energy market structures and carbon intensities influence the feasibility and sustainability of EV fleet integration across Europe. Moreover, the observed mismatch between cost savings and emission reductions highlights the need for future research to incorporate both cost efficiency and sustainability directly into the optimization process.

The structure of the paper is as follows: Sect. 2 presents the input data used for the optimization algorithm, along with a detailed description of the MILP formulation. It also introduces the two analyzed scenarios: a baseline scenario, representing unoptimized or typical charging behavior, and an optimal scenario, generated by the proposed optimization model. Section 3 provides a comparative analysis of both scenarios, focusing on charging costs and sustainability metrics. The analysis is conducted both within each country—comparing baseline and optimal strategies—and across countries, highlighting how differences in electricity markets, particularly price volatility and grid carbon intensity, influence the results. Finally, Sect. 4 summarizes the main conclusions and outlines directions for future research.

Indices:

$c \in C = \{1, 2, \dots, N_{ev}\}$	index of electric vehicles
$t \in T = \{1, 2, \dots, N_t\}$	index of time intervals
p_t	electricity price at time interval t
$a_{c,t} \in \{0, 1\}$	availability of vehicle c for charging at time t (1 if present, 0 if away)
$e_{c,t}$	energy consumed due to driving by vehicle c during interval t
$SOC_{initial,c}$	initial battery state of charge for vehicle c
$SOC_{target,c}$	required final SOC for vehicle c
SOC_{min}	minimum allowable SOC
SOC_{max}	battery capacity (maximum SOC)
$N_{chargers}$	number of available charging stations
$x_{c,t} \in R \geq 0$	energy charged to vehicle c at time t (in kWh)
$y_{c,t} \in \{0, 1\}$	binary variable; 1 if vehicle c is charging at time t , 0 otherwise
$SOC_{c,t} \in [0, SOC_{max}]$	SOC for vehicle c at the end of time t

2 Case Study Overview

Addressing the challenges associated with electric vehicle (EV) charging requires consideration of two main perspectives: the supply side, representing the power system, and the demand side, which encompasses the charging infrastructure. The supply side involves the energy market, where electricity is traded, and participants may engage in capacity markets or provide ancillary services. On the demand side, this study focuses specifically on charging stations, which shape electricity consumption patterns and offer opportunities for cost optimization. While the demand side can also include components such as EV users, on-site renewable generation, and energy storage systems, these are excluded from the presented analysis. Instead, this study focuses on charging stations, examining how cross-country variations in electricity prices and carbon intensity influence financial and environmental outcomes.

In this study, a two-step methodological approach is developed: cost optimization in the first step, followed by a post-optimization sustainability analysis. In the first step, cost

optimization is performed using a centralized control approach, where a single control entity collects input data and simultaneously optimizes the charging of an entire EV fleet, consisting of 10 vehicles and 2 chargers, using Mixed-Integer Linear Programming (MILP) for real-time optimization. The developed optimization algorithm relies on three key inputs: electricity prices (varies in time), driving schedule predictions (varies in time), and EV technical characteristics. The output of the model is an optimal charging schedule that minimizes total charging costs over a one-day planning horizon, with 15-min time intervals. In the second step, the resulting charging schedules, together with hourly carbon intensity values, serve as inputs for the post-optimization sustainability analysis. This analysis evaluates the environmental impact of the cost-optimized charging schedules. Figure 2 shows a schematic overview of the system, illustrating the connection between the real-world EV charging infrastructure and the optimization algorithm. It highlights key components on both the supply side (energy sources and the grid) and demand side (charging stations and EV fleet), as well as the data exchange between the physical environment and the centralized optimization framework. In addition, the flowchart illustrating this two-step methodology is presented in Fig. 3.

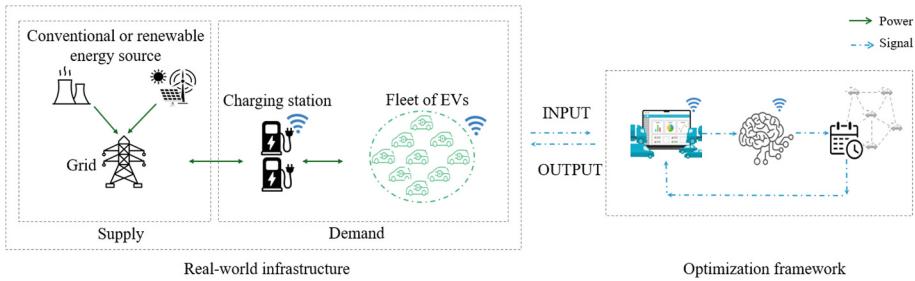


Fig. 2. Schematic overview of the electric vehicle (EV) fleet charging system

2.1 Optimization Inputs and Constraints

The optimal scheduling of electric vehicle (EV) charging represents a critical component in minimizing operational costs while ensuring grid stability and vehicle readiness. However, the problem is complex due to the interplay of technical constraints, user behavior, and dynamic electricity pricing. This section presents the input data used in the development of the centralized charging optimization algorithm, which is formulated as a Mixed-Integer Linear Program (MILP).

The algorithm incorporates day-ahead electricity prices for 5 countries: France, Germany, Poland, Sweden, from NordPool [27], and Croatia from the CROPEX [28]. For the purposes of this study, data from April 23, 2025, are used as a representative example, as shown in Fig. 4. This date was selected randomly to illustrate the modeling approach in a single-day case study. These prices are integrated into the cost optimization model as a time-series input that varies on an hourly basis. In addition, the most important input for the sustainability analysis is the hourly carbon intensity values, which reflect the time-varying CO₂ emissions per unit of electricity generated. Figure 5 illustrates the

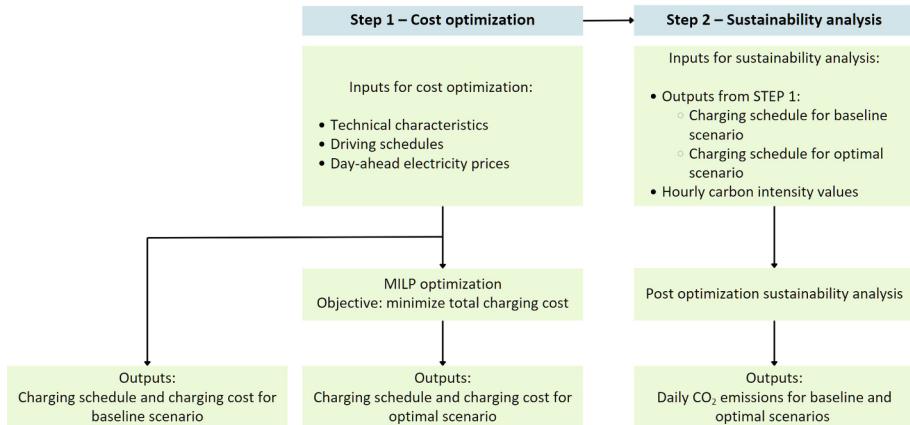


Fig. 3. Overview of the two-step methodological framework: cost optimization followed by post-optimization sustainability analysis

hourly variations in carbon intensity on the same date [26], highlighting the differences in grid emission profiles across the selected countries.

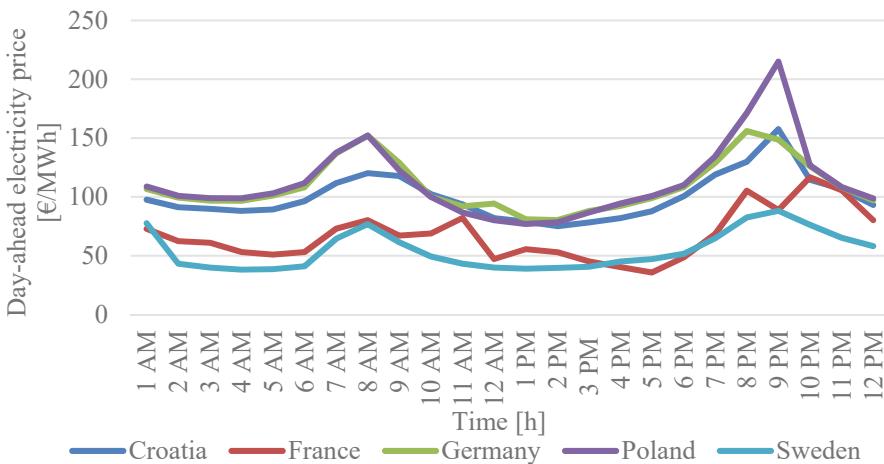


Fig. 4. Day-ahead electricity prices for selected countries on 23rd April 2025

The technical specifications of the EV fleet and charging infrastructure define the physical constraints of the cost optimization model. These include parameters such as battery capacity, minimum and target state of charge (SOC), energy consumption, and charging availability. Table 1 summarizes the key technical inputs used for simulation.

Among all input categories, driving behavior introduces the greatest uncertainty, significantly affecting the availability of EVs for charging. To address this, the algorithm incorporates flexible, statistically generated driving schedules using the Monte Carlo method instead of relying on static, predefined patterns. For each EV, the maximum

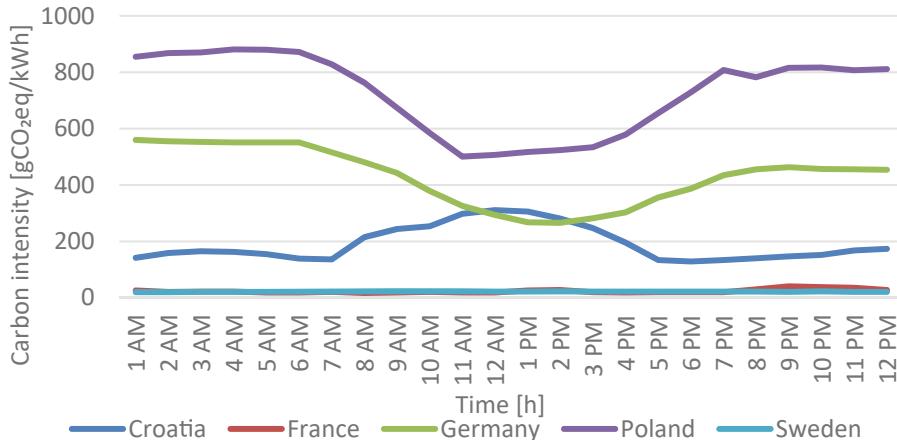


Fig. 5. Hourly carbon intensities for selected countries on 23rd April 2025

Table 1. Technical parameters used in the EV charging optimization model

Input	Value
Battery capacity [kWh]	60
Minimum state-of-charge (SOC) [%]	20
Target SOC [%]	100
Average consumption [Wh/km]	180
Average speed [km/h]	50
Number of cars [-]	10
Number of chargers [-]	2
Charging power [kW]	unlimited

number of possible driving intervals per day was first calculated based on its battery capacity and energy consumption characteristics, ensuring that each vehicle could complete a full day of operation without requiring mid-day charging. These driving intervals were then randomly distributed across the day, except for a fixed non-driving period from 9:00 PM to 12:00 AM the following day, during which all EVs were assumed to be parked and available for charging. This stochastic method enables the simulation of diverse and dynamic usage profiles across the fleet, better capturing real-world variability in driving behavior. The driving schedule profiles for 10 EVs over a 24-h period are presented in Fig. 6.

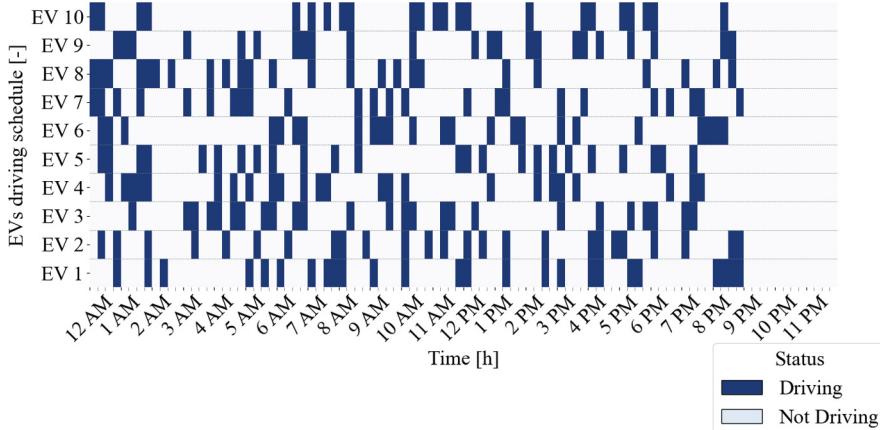


Fig. 6. Driving schedule profiles of 10 EVs over a 24-h period

2.2 Optimization Method

This section presents the mathematical formulation of the charging optimization problem. The objective is to determine an optimal charging schedule for a fleet of EVs that minimizes total electricity costs while satisfying system constraints such as charger availability, battery limitations, and vehicle availability.

The problem is formulated as a Mixed-Integer Linear Program (MILP), with continuous variables representing charging amounts and binary variables indicating whether a vehicle is charging at a given time step. In this analysis, the time horizon is discretized into 15-min intervals, providing sufficient resolution for capturing variations in electricity prices and vehicle usage patterns.

Objective Function. The primary objective is to minimize the total cost of charging all vehicles over the time horizon:

$$\min \sum_{c \in C} \sum_{t \in T} x_{c,t} \bullet p_t \quad (1)$$

Simultaneous Charging Constraint. At any given time, the number of vehicles charging must not exceed the number of available chargers:

$$\sum_{c \in C} y_{c,t} \leq N_{chargers}, \forall t \in T \quad (2)$$

Vehicle Availability Constraint. Vehicles can only be charged if they are present at the charging location:

$$y_{c,t} \leq a_{c,t}, \forall c \in C, \forall t \in T \quad (3)$$

Charging Logic. Energy can only be delivered to a vehicle if it is actively charging:

$$x_{c,t} \leq SOC_{max} \bullet y_{c,t}, \forall c, t \quad (4)$$

Battery Dynamics Calculation. Each vehicle’s battery SOC is updated based on the previous SOC, energy consumed during driving, and energy received from charging:

$$SOC_{c,t} = SOC_{c,t-1} - e_{c,t} + ex_{c,t}, \forall c, t \quad (5)$$

The initial condition is given by:

$$SOC_{c,0} = SOC_{initial,c}, \forall c \quad (6)$$

Each vehicle must reach or exceed its target SOC by the end of the planning horizon:

$$SOC_{c,H} \geq SOC_{target,c}, \forall c \quad (7)$$

SOC Bounds. The SOC must always remain within safety and capacity limits:

$$SOC_{min} \leq SOC_{c,t} \leq SOC_{max}, \forall c, t \quad (8)$$

2.3 Scenario Definition for Comparative Analysis

To evaluate the potential savings within each country, two scenarios are developed: a baseline case and an optimal case. The baseline scenario represents a *business-as-usual* approach, where all vehicles begin charging at the end of the day—when their state of charge (SOC) is at its lowest. In contrast, the optimal case reflects the outcome of a centralized model that utilizes a Mixed-Integer Linear Programming (MILP) optimization method, aiming to minimize the total charging cost. Both scenarios are based on identical input data, including the availability of 10 electric vehicles (EVs) and two chargers, allowing for a consistent comparison of cost efficiency and sustainability.

3 Results and Discussion

3.1 Economic Performance: Cost Comparison

This section presents the results of potential cost savings within each country and compares outcomes across different national electricity markets. The baseline case reflects a business-as-usual charging strategy, where 10 EVs are charged using 2 chargers only when their SOC reaches a critical level. In this study, the critical level is assumed to occur at the end of the day, and this charging behavior is predefined based on periods when the vehicles are not in use. Since this scenario does not utilize optimization, it typically results in higher total charging costs. The total baseline charging cost is calculated for all selected countries, providing a benchmark for comparison. In addition, Fig. 7 illustrates the SOC trajectories of each EV for Croatia as a representative example of the baseline case.

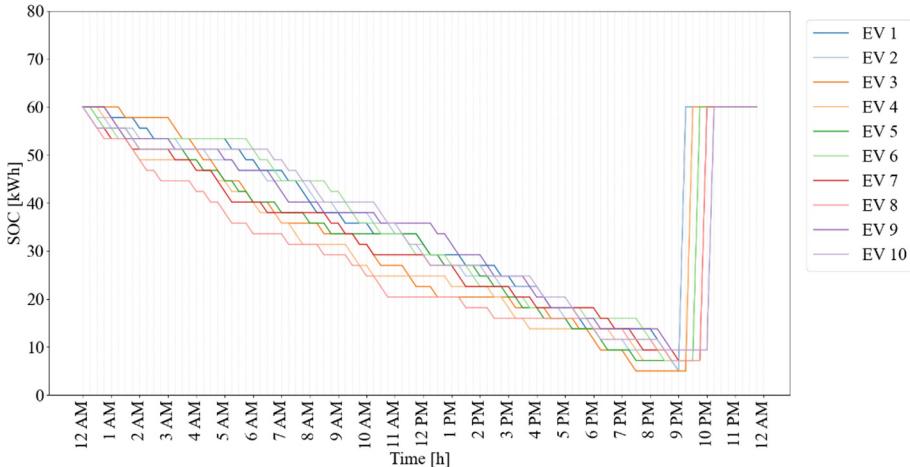


Fig. 7. Variations in daily SOC for the baseline case for Croatia

In contrast to the baseline, the optimal case represents the best-case scenario, where an optimization algorithm is used to minimize the total daily charging cost. Vehicles are still required to reach their target SOC by the end of the day, but the optimization model identifies more cost-effective time slots for charging—shifting charging activity to periods with the lowest electricity prices. Figure 8 illustrates the results for Croatia under the optimal charging scenario, showing the adjusted SOC trajectories of all EVs. Compared to the baseline case, this strategy achieves a 28% reduction in total charging costs.

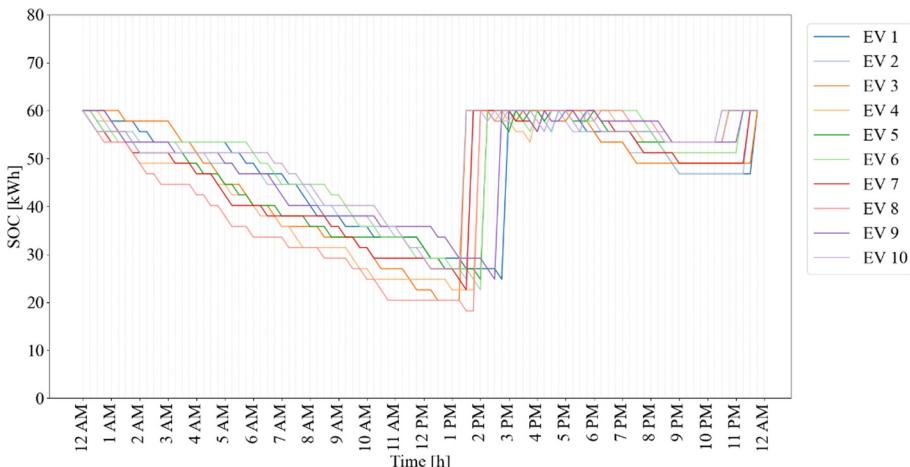


Fig. 8. Variations in daily SOC for the optimal case for Croatia

Figure 9 presents the daily charging costs for 10 EVs and 2 available chargers under both the baseline and optimal scenarios for all analyzed countries. As the results suggest, the variations in daily charging costs between countries are significant in both the baseline and optimal scenarios, reflecting differences in electricity prices and market structures. Germany and Poland report the highest total costs in both scenarios, while Sweden, benefiting from a more stable electricity market, shows the lowest total baseline cost—as well as the smallest percentage reduction in the optimal scenario. Cost savings range from 28% (Croatia) to 62% (France). Notably, even a 28% reduction is substantial—when scaled to large EV fleets, such savings can result in meaningful economic benefits, especially considering that improvements as small as 5% can be impactful at scale. It is important to note that this analysis is based on a specific day, and results may vary depending on daily electricity price fluctuations and driving patterns.

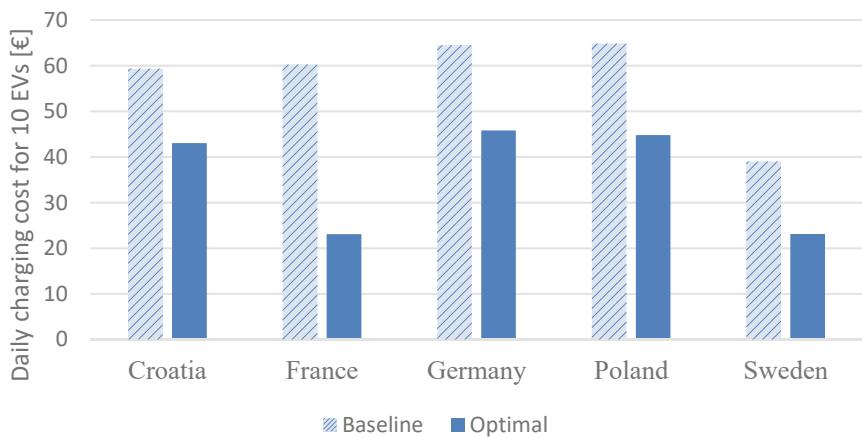


Fig. 9. Daily charging cost for the baseline and optimal scenarios across selected countries

3.2 Environmental Impact: CO₂ Emission and Sustainability

The previous section demonstrated the cost-saving potential achievable through real-time optimization of electric vehicle (EV) charging. This section shifts the focus to environmental impact, aiming to explore how minimizing daily charging costs—as defined by the objective function—and the resulting charging patterns influence sustainability outcomes, particularly in terms of carbon emissions, across different countries. To provide a reference point, we consider the same EV fleet characteristics—each vehicle having a 60 kWh battery and an average consumption of 180 Wh/km. Under these conditions, a single charge allows for a daily driving distance of approximately 333 km per vehicle, resulting in a total fleet distance of 3 333 km per day. If this same distance were covered by conventional internal combustion engine (ICE) vehicles, assuming an average emission factor of 170 g CO₂/km [29], the fleet would emit approximately 500 kg of CO₂ daily. This value serves as a worst-case baseline for emissions, enabling a more meaningful comparison of EV fleet sustainability across different electricity generation mixes. Figure 10 illustrates the resulting daily CO₂ emissions across selected countries, comparing both the baseline and the optimized charging scenarios.

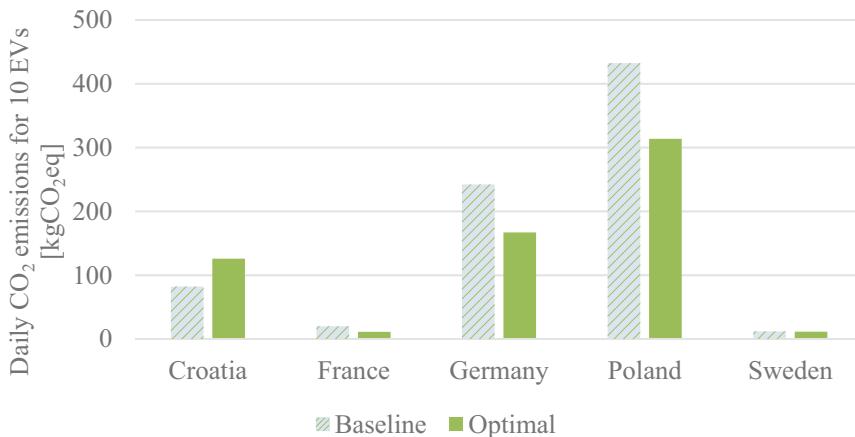


Fig. 10. Daily CO₂ emissions for baseline and optimal scenarios across selected countries

The results indicate that the implementation of EVs significantly reduces daily CO₂ emissions compared to internal combustion engine (ICE) vehicles. In most countries, further reductions are achieved through optimized charging strategies that align with periods of lower carbon intensity. However, there are some outliers—for instance, in the case of Croatia, the optimized charging scenario actually results in higher CO₂ emissions due to the specific characteristics of the local electricity generation mix. These national variations highlight that cost-effective charging does not always lead to more sustainable outcomes. Therefore, it is essential to incorporate both economic and environmental considerations into the early stages of EV fleet planning to ensure that charging strategies support both cost efficiency and long-term sustainability goals.

3.3 Discussion

In this section, a comparative analysis of cost savings and sustainability outcomes is presented for each country. The following Fig. 11 illustrates both the charging cost savings and the change in CO₂ emissions achieved in the optimal case, relative to the baseline scenario. Electricity price trends often correlate with carbon intensity, meaning that lower electricity prices—typically driven by increased renewable energy generation—are generally associated with lower CO₂ emissions. This relationship is observed in countries like France, Germany, Poland, and Sweden. However, Croatia exhibits the opposite pattern, with optimized charging leading to an increase in CO₂ emissions. This is largely due to the fact that Croatia is not yet energy independent and must rely on electricity imports. The rise in emissions can be attributed to imports from neighbouring countries such as Bosnia and Herzegovina and Serbia, where coal remains a dominant energy source [26]. It is important to emphasize that these results are based on a single reference day, and the situation may vary depending on the daily electricity mix and market conditions.

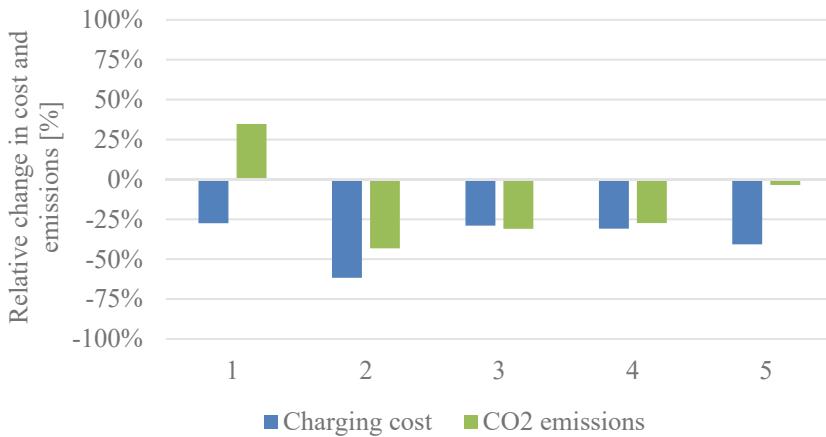


Fig. 11. Relative change in charging cost and CO₂ emissions across selected countries

France achieves the greatest cost reduction (−62%) and a substantial decrease in emissions (−43%). Its electricity system is primarily based on nuclear power. With policies targeting a phase-out of fossil-fuel vehicles by 2035 and strong investments in EV infrastructure, France is well-positioned to scale electrification in a sustainable way.

Germany's energy mix still includes a significant share, nearly 50%, from fossil sources, yet it also benefits from a strong presence of renewables, particularly solar and wind. The results of this study show a 29% reduction in charging costs and a 31% decrease in CO₂ emissions, indicating that EV integration can already deliver meaningful economic and environmental benefits under current conditions. As Germany advances toward its goal of phasing out fossil fuels, these benefits are expected to become even more substantial.

Sweden, on the other hand, combines low charging costs with minimal emissions, as a result of its electricity mix dominated by hydro and nuclear power. The country leads in EV adoption and plans to ban fossil-fuel vehicle sales by 2030, making it a frontrunner in aligning economic and environmental objectives.

Poland demonstrates notable cost advantages through optimized charging; however, its electricity system remains heavily reliant on coal. As a result, smart EV charging strategies, although economically beneficial, lead to the highest CO₂ emissions among the analyzed countries. Nevertheless, these emissions are still approximately 24% lower than those generated by a comparable fleet of internal combustion engine (ICE) vehicles, highlighting that even in coal-dependent systems, electrification offers meaningful environmental benefits. While coal remains the dominant source of emissions in Poland, ongoing national efforts to reduce coal dependence suggest that the sustainability potential of EV fleets will increase over time. This underscores the dual importance of both electrification and grid decarbonization, particularly in countries transitioning from fossil-heavy power systems.

Altogether, these examples emphasize that the sustainability of EV fleet implementation is deeply dependent on the national energy context. While optimization strategies can effectively reduce costs, their environmental impact is ultimately determined by the

cleanliness of the electricity supply. Incorporating CO₂ emissions minimization along-side cost minimization in the objective function would be a logical next step to balance both goals. Although this approach may reduce the extent of cost savings compared to a purely cost-focused strategy, it would lead to lower overall carbon emissions—without necessarily achieving the absolute minimum in either objective.

4 Conclusion

In this study, the cost and sustainability outcomes of electric vehicle (EV) charging were analyzed for a small-scale fleet of 10 EVs and 2 chargers. A two-step methodological framework was applied: first, cost optimization based on Mixed-Integer Linear Programming (MILP) was performed, followed by a post-optimization sustainability analysis. The framework was applied under consistent vehicle characteristics and statistically generated driving schedules, with variations in day-ahead electricity prices and carbon intensities across five European countries: Croatia, France, Germany, Poland, and Sweden. The objective was to evaluate and compare the potential for EV fleet optimization within and across these countries.

The results highlight that when assessing EV charging strategies, both cost efficiency and sustainability must be considered simultaneously. While centralized optimization consistently led to cost reductions (28%–62%), the environmental outcomes varied depending on each country’s electricity mix. In Germany, France, Poland, and Sweden, lower costs aligned with CO₂ emission reductions (4%–43%). However, Croatia experienced an increase in emissions, largely due to its reliance on electricity imports from neighboring coal-based systems. These findings highlight that while financial savings are achievable across different contexts, environmental benefits depend strongly on the cleanliness of the electricity supply at the time of use.

This type of cross-market analysis provides valuable insights for strategic decision-making, especially when determining which electricity markets or regions offer the most favorable conditions for fleet deployment and smart charging solutions. This work lays a solid foundation for future algorithm development. Future work will expand the current one-day case study into a multi-day analysis to better capture the temporal variability in electricity prices and carbon intensity. Additionally, a multi-objective optimization approach will be developed to incorporate both cost and sustainability objectives directly into the MILP formulation. This will enable a more comprehensive assessment of trade-offs and allow for comparison with the current cost-driven optimization results. Further enhancements will include the integration of electricity price forecasting and driving schedule prediction models, enabling greater adaptability and real-time applicability of the optimization framework across diverse energy markets.

This research has been supported and funded by Verne Croatia, a company focused on the development of technologies for autonomous vehicles. The authors gratefully acknowledge their support.

References

1. Outlook for electric mobility – Global EV Outlook 2024 – Analysis. <https://www.iea.org/reports/global-ev-outlook-2024/outlook-for-electric-mobility>. Accessed 21 Apr 2025

2. Delivering the European Green Deal - European Commission. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal/delivering-european-green-deal_en. Accessed 21 Apr 2025
3. Climate strategies & targets - European Commission. https://climate.ec.europa.eu/eu-action/climate-strategies-targets_en. Accessed 21 Apr 2025
4. Liikkanen, J., Moilanen, S., Kosonen, A., Ruuskanen, V., Ahola, J.: Cost-effective optimization for electric vehicle charging in a prosumer household. *Solar Energy*, 112122 (2023). <https://doi.org/10.1016/j.solener.2023.112122>
5. Das, R., Wang, Y., Busawon, K., Putrus, G., Neaimeh, M.: Real-time multi-objective optimisation for electric vehicle charging management. *J. Clean. Prod.* **292**, 126066 (2021). <https://doi.org/10.1016/j.jclepro.2021.126066>
6. Moradi Amani, A., Sajjadi, S.S., Somaweera, W.A., Jalili, M., Yu, X.: Data-driven model predictive control of community batteries for voltage regulation in power grids subject to EV charging. *Energy Rep.* **9**, 236–244 (2023). <https://doi.org/10.1016/j.egyr.2022.12.089>
7. Sykiotis, S., Athanasoulas, S., Temenos, N., Rallis, I., Doulamis, A., Doulamis, N.: Community-driven Smart EV charging With Multi-Agent Deep Reinforcement Learning. In: 2024 International Joint Conference on Neural Networks (IJCNN), pp. 1–8 (2024). <https://doi.org/10.1109/IJCNN60899.2024.10650988>
8. Moghaddass, R., Mohammed, O.A., Skordilis, E., Asfour, S.: Smart control of fleets of electric vehicles in smart and connected communities. *IEEE Trans. Smart Grid* **10**, 6883–6897 (2019). <https://doi.org/10.1109/TSG.2019.2913587>
9. Lee, H., Kim, J., Park, H.: Large-scale electric vehicle charging coordination for cost-effectiveness and fairness under peak power constraints. *Int. J. Electr. Power Energy Syst.* **166**, 110539 (2025). <https://doi.org/10.1016/j.ijepes.2025.110539>
10. Lee, Z.J., et al.: Adaptive charging networks: a framework for smart electric vehicle charging. *IEEE Trans. Smart Grid* **12**, 4339–4350 (2021). <https://doi.org/10.1109/TSG.2021.3074437>
11. Aygun, A.I., Kamalasadan, S.: Centralized charging approach to manage electric vehicle fleets for balanced grid. In: 2022 IEEE International Conference on Power Electronics, Smart Grid, and Renewable Energy (PESGRE), pp. 1–6 (2022). <https://doi.org/10.1109/PESGRE52268.2022.9715836>
12. Zafar, S., Blavette, A., Camilleri, G., Ben Ahmed, H., Prince Agbodjan, J.J.A.: Decentralized optimal management of a large-scale EV fleet: optimality and computational complexity comparison between an adaptive MAS and MILP. *Int. J. Electr. Power Energy Syst.* (2022). <https://doi.org/10.1016/j.ijepes.2022.108861>
13. Paudel, A., Hussain, S.A., Sadiq, R., Zareipour, H., Hewage, K.: Decentralized cooperative approach for electric vehicle charging. *J. Clean. Prod.* **364**, 132590 (2022). <https://doi.org/10.1016/j.jclepro.2022.132590>
14. Zhang, W., et al.: Decentralized electric vehicle charging strategies for reduced load variation and guaranteed charge completion in regional distribution grids. *Energies* **10**, 147 (2017). <https://doi.org/10.3390/en10020147>
15. Raju, B.A., Mohan Reddy, P.K., Vuddanti, S.: Decentralized coordination of electric vehicle charging in microgrids using game theory. In: 2024 IEEE 3rd International Conference on Electrical Power and Energy Systems (ICEPES), pp. 1–5 (2024). <https://doi.org/10.1109/ICEPES60647.2024.10653600>
16. Triviño, A., López, A., Yuste, A.J., Cuevas, J.C.: Decentralized EV charging and discharging scheduling algorithm based on Type-II fuzzy-logic controllers. *J. Energy Storage* **93**, 112054 (2024). <https://doi.org/10.1016/j.est.2024.112054>
17. Sepetanc, K., Pandzic, H.: A cluster-based model for charging a single-depot fleet of electric vehicles. *IEEE Trans. Smart Grid* **12**, 3339–3352 (2021). <https://doi.org/10.1109/TSG.2021.3064272>

18. Sevdari, K., Calero, L., Striani, S., Andersen, P.B., Marinelli, M., Ronnow, L.: Autonomous distributed control of electric vehicle chargers for grid services. In: 2021 IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe), Espoo, Finland, pp. 1–5. IEEE (2021). <https://doi.org/10.1109/ISGTEurope52324.2021.9640132>
19. Yi, Z., et al.: A highly efficient control framework for centralized residential charging coordination of large electric vehicle populations. *Int. J. Electr. Power Energy Syst.* **117**, 105661 (2020). <https://doi.org/10.1016/j.ijepes.2019.105661>
20. Engel, J., Schmitt, T., Rodemann, T., Adamy, J.: Hierarchical economic model predictive control approach for a building energy management system with scenario-driven EV charging. *IEEE Trans. Smart Grid* **13**, 3082–3093 (2022). <https://doi.org/10.1109/TSG.2022.3160390>
21. Hussain, S., et al.: Hybrid coordination scheme based on fuzzy inference mechanism for residential charging of electric vehicles. *Appl. Energy* **352**, 121939 (2023). <https://doi.org/10.1016/j.apenergy.2023.121939>
22. Fernandez, V., Pérez, V.: Optimization of electric vehicle charging control in a demand-side management context: a model predictive control approach. *Appl. Sci.* **14**, 8736 (2024). <https://doi.org/10.3390/app14198736>
23. De La Torre, S., Aguado, J.A., Sauma, E.: Optimal scheduling of ancillary services provided by an electric vehicle aggregator. *Energy* **265**, 126147 (2023). <https://doi.org/10.1016/j.energy.2022.126147>
24. Bodenschatz, N., Eider, M., Berl, A.: Mixed-integer-linear-programming model for the charging scheduling of electric vehicle fleets. In: 2020 10th International Conference on Advanced Computer Information Technologies (ACIT), pp. 741–746 (2020). <https://doi.org/10.1109/ACIT49673.2020.9208875>
25. Per capita energy from fossil fuels, nuclear and renewables. <https://ourworldindata.org/grapher/per-capita-energy-source-stacked>. Accessed 05 May 2025
26. Live 24/7 CO₂ emissions of electricity consumption. <http://electricitymap.tmrow.co>. Accessed 25 Apr 2025
27. Nord Pool | Day-ahead prices. <https://data.nordpoolgroup.com/auction/day-ahead/prices?deliveryDate=latest¤cy=EUR&aggregation=DeliveryPeriod&deliveryAreas=FR,GER,PL,SE1>. Accessed 18 Apr 2025
28. HRVATSKA BURZA ELEKTRIČNE ENERGIJE - HRVATSKA BURZA ELEKTRIČNE ENERGIJE d.o.o. - cropex.hr. <https://www.cropex.hr/hr/>. Accessed 16 Jan 2024
29. Ritchie, H.: Which form of transport has the smallest carbon footprint? Our World in Data (2023)



Electric Vehicle Based Virtual Electricity Network (EVEN) Solution for Performance Enhancement in Distribution Networks

Pei Huang^{1,2,3,4}(✉)  and Rehman Zafar^{3,4} 

¹ School of Business, Society, and Engineering, Mälardalen University, Västerås, Sweden
pei.huang@mdu.se

² Future Energy Center (FEC), Mälardalen University, Västerås, Sweden

³ Institution of Information and Technology, Dalarna University, Falun, Sweden

⁴ Dalarna University, SERC, Falun, Sweden

Abstract. Large-scale electric-vehicle (EV) uptake is challenging the power grid due to a lack of sufficient hosting capacity. Most smart-charging studies still treat EVs as stationary loads in one location and ignore their mobility. This work closes that gap by evaluating an EV-based virtual electricity network (EVEN) that lets vehicles charge at one feeder and discharge at another. We formulate a comprehensive framework that combines inter-network energy-delivery optimization, stochastic time-series hosting capacity analysis, and battery-degradation assessment. The approach is tested on two real-world distribution systems: a voltage-constrained 50-bus rural residential network and a capacity-rich 76-bus industrial network. Simulation results reveal that shifting only 10% of the evening demand from the rural to industrial network cuts the rural undervoltage index by roughly 80% and weekly violation counts by 65%, while adding no more than two minor violations upstream. At this level, average yearly battery-cycling degradation rises modestly from 0.4% to 0.8%. The study thus demonstrates EVEN as a cost-effective, scalable alternative to physical reinforcement and provides the first integrated assessment linking network-level benefits with battery-health impacts.

Keywords: Electric Vehicle · Hosting Capacity · E-mobility · Grid Integration

1 Introduction

Electric vehicles (EVs) are promising solutions for the decarbonization of transport sector. With the target of transport electrifications in many countries and regions, the past decades have witnessed dramatic increase of EV deployment. Recent reports such as the *IEA Global EV Outlook 2024* illustrates the exponential market growth and the strategic role EVs are expected to play in future energy systems [1]. The increasing number of EVs, which need to be connected and charged in the power grid, is creating significant challenges to many of distribution networks today. This is because most of the existing distribution networks were designed several decades ago, without considering the increasing EV charging loads and other new electric loads such as heat pumps [2].

As a result, there is a lack of sufficient hosting capacity in many distribution networks. Hosting capacity is the maximum amount of electricity consumption or generation that can be integrated into the power system, while still maintaining its performance within acceptable limits [3]. Due to a lack of sufficient hosting capacity, problems such as line overloads, transformer overloads, power quality deterioration, and voltage drop may occur [4].

To enhance the grid hosting capacity, existing studies have developed different measures. The most straightforward approach is to reinforce or reconfigure the grid. However, this approach is expensive and it typically takes long time. Other measures include using on-load tap charging transformers, reactive power control, energy storage or demand response [5, 6]. Among these methods, a promising solution is demand response and smart charging of EVs. For instance, in [2], a smart EV charging approach is developed to enhance the hosting capacity for EVs. In [6], different EV charging strategies are studied and compared from the hosting capacity perspective, and it was found that spreading out EV charging loads during the connection period can benefit the grid most. In [7] a coordinated charging control method is developed for EV aggregators to optimize the charging loads of an EV fleet. The developed control effectively improved the hosting capacity by 15% for a 40% EV penetration level.

Despite notable progress in smart EV charging, several limitations remain. Many studies rely on overly simplified grid models, overlooking battery-wear effects and the intricate mobility patterns that arise in real-world operation. Moreover, research typically treats EV integration as a single-location problem, ignoring the vehicles' potential as *mobile* storage assets. In practice, an EV can charge at one node and later discharge at another, thereby forming an EV-based virtual electricity network (EVEN). Although previous work has explored the benefits of EVEN for local energy balancing and cost reduction [8–10], its contribution to distribution-grid hosting capacity is still unclear.

To bridge the research gap, the present study develops a comprehensive framework that combines inter-network energy-delivery optimization, stochastic time-series hosting capacity analysis, and battery-degradation assessment. Two real-world distribution feeders serve as case studies. We quantify how EVEN affects voltage violations in both networks and evaluate the associated cycling degradation of the EV batteries. To our knowledge, this is one of the first investigations into the impact of EV-mediated energy transfer on hosting capacity. The findings advance the theoretical understanding of EVEN and provide actionable guidance for grid operators seeking to exploit EVs as dynamic, mobile energy resources.

2 Methodology

This section introduces the methodology developed for the EVEN solution. First, the EV charging/discharging profiles are optimized. Then, the EV charging/discharging profiles are assigned to specific buses in each distribution network based on the EVs' connection period. Then the grid performance is analyzed using power flow analysis in a stochastic way. After the power flow simulation, the voltage violation is estimated, and the battery cycling degradation is calculated. Details about each step are introduced below.

2.1 Control of EVs for Electricity Delivery in the EVEN Solution

The EVEN solution leverages the mobile storage of EV batteries to actively deliver electricity between different places. Within this study, the aim is to deliver electricity from areas with strong distribution network to areas with weak distribution networks. In the scope of this study, ‘strong’ and ‘weak’ are defined based on the hosting capacity of the network. This study considers EVs commuting between two distribution networks: a rural distribution network and an industrial distribution network. The optimization is to minimize the aggregated peak loads in the bus connection, considering the base loads. It is depicted by Eq. (1).

$$J = \text{minimize}(\max(e_{ex,t})), \quad (1)$$

$e_{ex,t}$ (kW) is the hourly aggregated load on a bus, as calculated by Eq. (2),

$$e_{ex,t} = e_{d,t}^b + u_t, \quad (2)$$

$e_{d,t}^b$ (kW) is the base load on the bus, and u_t (kW) is the EV charging/discharging rate. The EV charging/discharging rates should meet the following constraints:

$$u_{lm,disc} < u_t < u_{lm,char}. \quad (3)$$

$$0 \leq SOC_0 \times CAP + (u_{t_j} + u_{t_j+1} + \dots + u_{t_j+i}) \times \tau \leq CAP \quad \text{where } i = 1, 2, \dots, n^k, \quad (4)$$

$u_{lm,char}$ (kW) and $u_{lm,disc}$ (kW) are the maximum charging rate and discharging rate, respectively. SOC_0 is the initial State-of-Charge when arriving at a bus, and CAP (kWh) is the EV battery capacity. τ (hour) is 1 h. Furthermore, the total amount of electricity charged in each session is calculated by Eq. (5).

$$(u_{t_k} + u_{t_k+1} + \dots + u_{t_k+n_k}) \times \tau = \begin{cases} CAP \times (1 - SOC_0), \text{ in strong network} \\ \beta \cdot E_{load}, \text{ in weak network} \end{cases} \quad (5)$$

In the strong network, EV is always charged to full capacity. While in the weak network, it will discharge electricity. The amount to be discharged is defined by a coefficient β , which defines the load reduction percentage. A larger value means using EVs to deliver more electricity to the weak distribution network. E_{load} (kWh) is the total base electricity loads in the connection period.

The optimization is conducted at hourly interval. The optimization horizon is the connection period of EV in each charging session.

2.2 Hosting Capacity Analysis

This study adopts a stochastic time-series approach for evaluating the hosting capacity of the distribution networks. A winter week with high average loads is selected for analysis. Considering uncertainty in the EV model and location of EV connection in

each network, Monte Carlo (MC) simulation is conducted to evaluate the power grid performances in a stochastic framework.

For this study, AC power flow analysis is performed to evaluate the distribution network performance. The MATPOWER tool is utilized to run the power flow simulations. The power flow analysis is based on the well-known Newton-Raphson method, which calculates steady-state bus voltages and line flows by iteratively solving the complex power mismatch equations [11, 12].

There are many performance indicators for assessing the power grid performances. For instance, in [6], three performance indices, including voltage violation, cable overloading, and transformer overloading are considered. This study takes a similar approach and evaluates the voltage profiles for violations. Note that the current and thermal overloading are not considered as both case power grid models have large redundancy in these aspects. In this study, voltage violation is considered as the performance indicator. Considering the network work faces critical challenges with large load integration, only undervoltage problem is considered. 0.95 p.u. is set as a threshold as in [6], and once the voltage of a bus goes below 0.95 p.u., it is considered as a violation. Both the number of violations and the amount of voltage violations (i.e., how far it is from the threshold) are estimated.

2.3 Modelling of Battery Cycling Degradation

The cycling degradation of EV battery is evaluated using the Rainflow Counting algorithm, which was originally formulated for fatigue analysis. The Rainflow Counting algorithm converts the depth-of-discharge (DoD) time-series data into an equivalent set of stress cycles [13]. Then the DoD ranges and associated cycles can feed empirical life models for assessing degradation. The model from [14] is used. It is shown by Eq. (6). a and b are coefficients, which are different for different types of battery chemistry. They are taken from [15].

$$D_{CL} = \sum_{k=1}^m \frac{\text{Cycle of } R_k}{a \times (R_k)^b} \quad (6)$$

A limitation of this study is the consideration of cycling degradation only. Future work will look into the calendar degradation and make more comprehensive evaluation of the battery performance.

3 Case Studies and Results

This section first introduces the case systems. Then, the grid performance improvements via the EVEN solution are presented and discussed, followed by the battery degradation analysis results.

3.1 Description of the Systems Considered

To guarantee that the study reflects practical operating conditions, we employ two measured Norwegian medium-voltage feeders: (i) a 50-bus rural residential network [16] and

(ii) a 76-bus industrial distribution network [17]. The rural feeder is purposely chosen because it already exhibits multiple off-nominal voltage nodes in the base case, providing a stringent benchmark for additional EV penetration. A representative winter week is simulated; Fig. 1 plots the active-power demand at every load bus for both systems. In the residential feeder, several nodes carry almost constant demand while others follow the typical evening-peaking household profile. Conversely, the industrial feeder is characterized by pronounced daytime plateaus and nocturnal troughs, and its mean demand per bus is substantially higher than that of the rural system, indicating greater unused capacity for power delivery in off-peak hours.

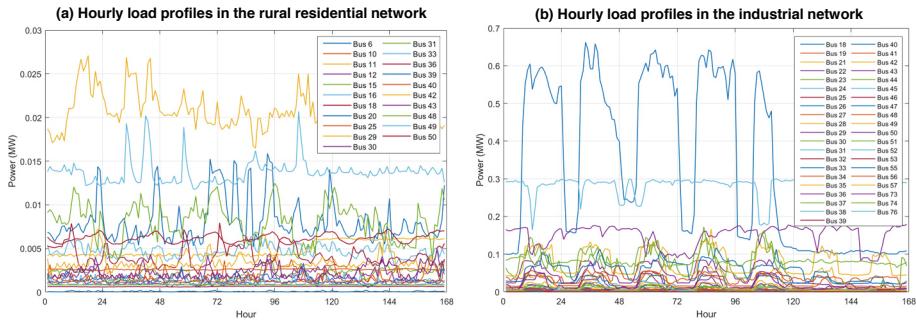


Fig. 1. Hourly load profiles on each bus in the residential network and industrial network

A set of EV models in the market are considered, with capacities ranging from 53 kWh to 77 kWh. The total number of EVs is 21, equaling the total number of buses with loads in the residential network. The mobility pattern of vehicles is taken from statistical survey. Each MC run proceeds as follows. First, the vehicles are randomly assigned to connection nodes in both test networks to reflect stochastic usage behaviour. Secondly, individual charging and discharging schedules are optimized using the method described in Sect. 2.1. The resulting power profiles are then added on the native bus demands and an AC power-flow calculation is executed. This sequence is repeated 200 times to build a statistically robust ensemble that captures uncertainty in EV mobility and location. For the residential feeder five demand-reduction objectives, expressed by the dimensionless index β in (5), are investigated: $\beta = 0.10, 0.20, 0.30, 0.40$ and 0.50 . Here $\beta = 0.10$ requires the EV fleet connected to a given bus to offset 10% of the baseline electric loads during its connection period, while higher β values representing more electricity delivered via EVs. The total simulation time for each case of β is around 1h on an Intel® Core™ i7-1355U computer.

3.2 Grid Performance Improvement via EVEN Solution

Figure 2 compares the network-wide voltage-violation index under the five scenarios. Subplots (a) and (b) display violin plots built from 200 MC runs for the residential and industrial feeders, respectively; the dashed line denotes the no-EV benchmark. In the residential system a mere 10% load reduction ($\beta = 0.10$) cuts the mean violation index

from around 2.6 to 0.5—an 80% mitigation—whereas pushing β beyond 0.10 yields only marginal additional benefit. The industrial feeder, which must supply the exported energy, shows a monotonic increase in violations as β rises, yet the absolute values remain below 3×10^{-3} , i.e., orders of magnitude lower than in the residential case and therefore operationally negligible.

Subplots (c) and (d) report the aggregate count of voltage-limit breaches observed in each MC realization for the residential and industrial feeders, respectively. In the rural residential network the introduction of EV electricity delivery for 10% load reduction ($\beta = 0.10$) cuts the mean violation count from roughly 460 to 168 times per week—a 64% improvement. Beyond $\beta = 0.20$, further load reduction yields only marginal additional gains, indicating a saturation of the hosting-capacity benefit. In contrast, the industrial feeder—initially free of violations—experiences a modest rise to an average of two times when β is increased to 0.50, reflecting its role as the upstream supplier of the exported energy. Overall, the EVEN strategy markedly strengthens the voltage profile of the weak residential system while imposing only negligible stress on the more robust industrial backbone.

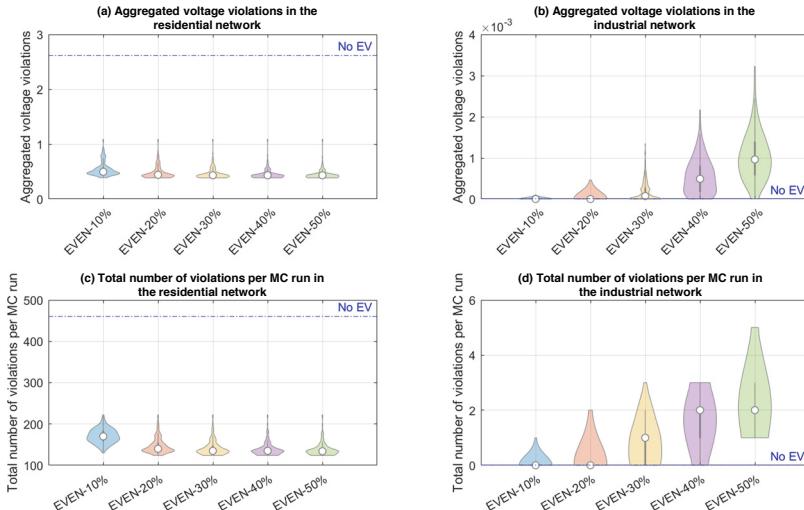


Fig. 2. Hosting capacity analysis results of the two distribution networks

Figure 3 visualizes the spatiotemporal undervoltage pattern in the residential feeder for the base case and the EVEN-10% scenario. Marker color denotes the mean deviation below the 0.95 p.u. limit, and marker area is proportional to the seven-day violation count aggregated over 200 MC runs. Without EV support, widespread and severe violations concentrate in the evening peak (18:00–24:00), driven by high household demand. Supplying just 10% of that demand from vehicle batteries attenuates both the depth and frequency of these evening breaches and slightly lowers the number of affected buses.

Daytime conditions are unchanged because the fleet is off-site; hence larger energy-offset targets (EVEN-20% to -50%) yield little extra benefit, consistent with the plateau seen in Fig. 2.

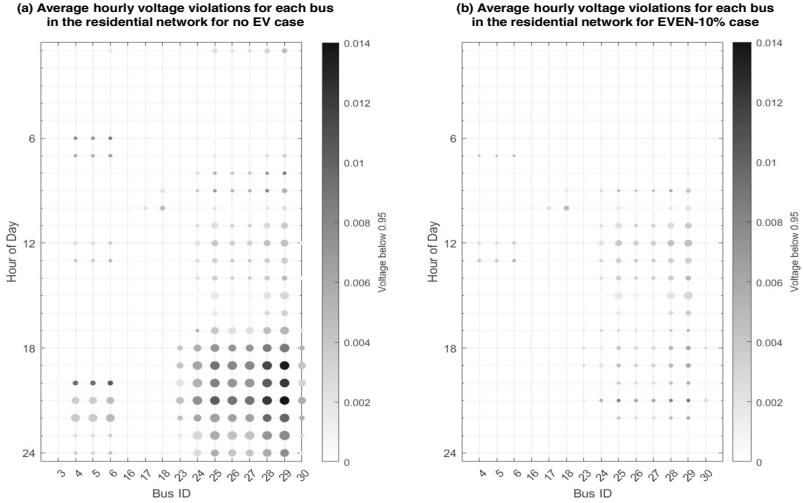


Fig. 3. Bubble plot of the average hourly voltage violations for each bus in the residential network for (a) no EV case, and (b) EVEN-10% case.

In summary, the analysis shows that the EVEN solution can significantly improve voltage regulation in the weak rural feeder while imposing only negligible additional stress on the robust industrial backbone. When the upstream network retains sufficient hosting capacity, energy withdrawn there and injected downstream cuts undervoltage depth and frequency in the residential system without breaching limits in the source feeder. These gains, however, arise solely during the hours in which EVs are connected to the weak network. Scenario sweeps further indicate the existence of an optimum demand-offset target—here around 10% of local load—beyond which incremental benefits saturate for the residential feeder and begin to raise violation counts in the industrial network.

3.3 Battery Degradation Results

Figure 4 quantifies the EV battery cycling degradation scaled to one year period in different scenarios. In the reference case, i.e., unidirectional overnight charging only, the mean capacity degradation attributable to cycling is 0.4%/year. Introducing bidirectional operation to offset 10% of local demand doubles the average to around 0.8%/year; raising the offset to 50% pushes the fleet-wide mean to about 1.5%/year. Dispersion is considerable: vehicles that supply the highest energy volumes (up to 50 kWh per day at $\beta = 0.50$) experience annual cycling losses exceeding 4%. Importantly, Sect. 3.2 showed that the EVEN-10% scenario already delivers most of the grid-support benefit, implying that a modest (i.e., about 0.4%/year) increment in battery ageing suffices to

secure substantial network gains. Because the degradation metric in (6) is chemistry-specific, absolute values will vary with cell design; the present results should therefore be interpreted as qualitative guidance on the trade-off between grid support level and battery wear.

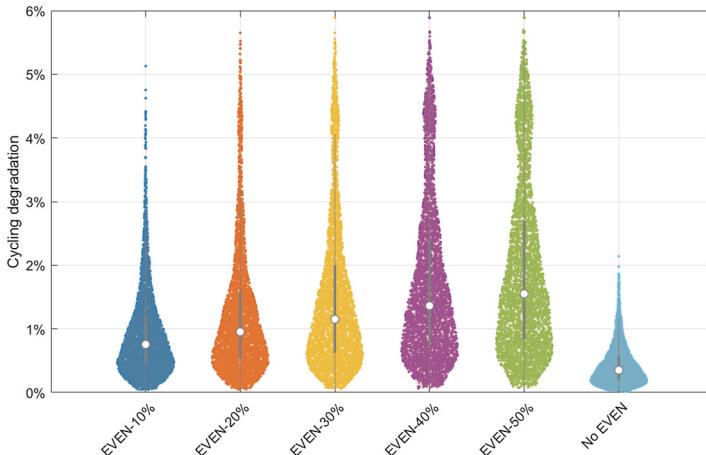


Fig. 4. Violin plot of annual EV battery degradation due to cycling in each scenario

4 Conclusions

This study analyzes how E-mobility can be used as a resource for enhancing the performances of distribution networks with limited hosting capacity. It presents a stochastic time-series framework that co-optimizes bidirectional electric-vehicle (EV) charging schedules for inter-feeder energy exchange and quantifies the resulting impacts on both network operation and battery health. Two real-world distribution systems, a voltage-constrained rural residential feeder and a capacity-rich industrial feeder, serve as case studies. The main findings are summarized below.

- EVs have large potential for delivering electricity between different places of the power grid.
- Even a small amount of electricity delivery can significantly enhance the hosting capacity of weak distribution networks. Offsetting only 10% of the evening demand at the weak feeder cuts its undervoltage index by $\sim 80\%$ and lowers violation counts by $\sim 65\%$.
- When the source feeder possesses adequate reserve capacity, the additional power export raises its violation rate only marginally (≤ 2 events per week at a 50% offset), indicating negligible operational risk.
- Delivering electricity will increase the battery cycling degradation. However, when the amount of delivered electricity is small, the cycling degradation will not increase significantly.

This study proposes a cost-effective and scalable solution for enhancing the hosting capacity in weak distribution network without the need of large-scale reinforcement

of the existing grid infrastructure. Future research will benchmark this strategy against other flexibility options and smart charging strategies and incorporate calendar ageing models to estimate total battery degradation under prolonged deployment. Future work will also investigate the economic performances of the proposed solution compared with other solutions.

Acknowledgement. This work was supported by Formas Early Career Research Fund [Grant number: 2022–00444] and Energimyndigheten [Grant number: 2023–00498].

References

1. 2024 GEO. 2024. Trends in electric cars. <<https://www.iea.org/reports/global-ev-outlook-2024/trends-in-electric-cars>>. Accessed 1 Nov 2024
2. Fachrizal, R., Ramadhani, U.H., Munkhammar, J., Widén, J.: Combined PV–EV hosting capacity assessment for a residential LV distribution grid with smart EV charging and PV curtailment. *Sustain. Energy Grids Netw.* **26**, 100445 (2021)
3. Bollen, M.H., Hassan, F.: *Integration of Distributed Generation in the Power System*. John Wiley & Sons (2011)
4. Panossian, N., Muratori, M., Palmintier, B., Meintz, A., Lipman, T., Moffat, K.: Challenges and opportunities of integrating electric vehicles in electricity distribution systems. *Current Sustain. Renewable Energy Reports* **9**, 27–40 (2022)
5. Ismael, S.M., Aleem, S.H.A., Abdelaziz, A.Y., Zobaa, A.F.: State-of-the-art of hosting capacity in modern power systems with distributed generation. *Renew. Energy* **130**, 1002–1020 (2019)
6. Sandström, M., Huang, P., Bales, C., Dotzauer, E.: Evaluation of hosting capacity of the power grid for electric vehicles – a case study in a Swedish residential area. *Energy* **284**, 129293 (2023)
7. Quijano, D.A., Melgar-Dominguez, O.D., Sabillon, C., Venkatesh, B., Padilha-Feltrin, A.: Increasing distributed generation hosting capacity in distribution systems via optimal coordination of electric vehicle aggregators. *IET Gener. Transm. Distrib.* **15**, 359–370 (2021)
8. Huang, P., Ma, Z.: Unveiling electric vehicle (EV) charging patterns and their transformative role in electricity balancing and delivery: insights from real-world data in Sweden. *Renew. Energy* **236**, 121511 (2024)
9. Board, A., Sun, Y., Huang, P., Xu, T.: Community-to-vehicle-to-community (C2V2C) for inter-community electricity delivery and sharing via electric vehicle: Performance evaluation and robustness analysis. *Appl. Energy* **363**, 123054 (2024)
10. Zhang, L., Alvarez, E.M.O., Huang, P.: Optimizing coordinated spatio-temporal control of electric vehicles for enhanced energy sharing and performance across building communities. *Energy Build.* **312**, 114167 (2024)
11. Zimmerman, R.D., Murillo-Sánchez, C.E., Thomas, R.J.: MATPOWER: steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Trans. Power Syst.* **26**, 12–19 (2010)
12. Grainger, J., Stevenson, W., Chang, G.: Power system stability. *Power Syst. Anal.* 528–550 (2016)
13. Alam, M., Saha, T.: Cycle-life degradation assessment of battery energy storage systems caused by solar PV variability. *Power Energy Soc. Gen. Meet. (PESGM)*, pp. 1–5 (2016)
14. Huang, P., Tu, R., Zhang, X., Han, M., Sun, Y., Hussain, S.A., et al.: Investigation of electric vehicle smart charging characteristics on the power regulation performance in solar powered building communities and battery degradation in Sweden. *J. Energy Stor.* **56**, 105907 (2022)

15. Electropaedia. Battery Life (and Death) (2020). <<https://www.mpoweruk.com/life.htm>>. Accessed 15 June 2021
16. Engan, L.M., Ekrheim, S., Bjarghov, S., Klemets, J.R.A., Schytte, I., Kjølle, G.: Reference dataset for semi-urban and rural Norwegian low voltage distribution grids. *Data Brief* **59**, 111453 (2025)
17. Sandell, S., Bjerkehagen, D., Birkeland, B., Sperstad, I.B.: Dataset for a Norwegian medium and low voltage power distribution system with industrial loads. *Data Brief* **48**, 109121 (2023)



Variational Quantum Eigensolver-Based CaaS Business Model for V2G

Desh Deepak Sharma¹ , Ramesh C. Bansal², and Jeremy Lin³

¹ MJP Rohilkhand University, Bareilly, India

desh.sharma@mjpnu.ac.in

² Electrical Engineering Department, University of Sharjah, Sharjah, United Arab Emirates

³ Krieger School of Art and Sciences, Johns Hopkins University, Washington, DC, USA

Abstract. In electric vehicles (EV) ecosystem, Charging-as-a-Service (CaaS) with subscription models has emerged as a promising approach to optimize EV charging infrastructure while ensuring cost-effectiveness and accessibility. However, traditional optimization methods struggle to handle the complex, dynamic nature of large-scale EV charging networks. This paper proposes a novel Charging-as-a-Service (CaaS) business model leveraging the Variational Quantum Eigensolver (VQE) algorithm to optimize energy management in Vehicle-to-Grid (V2G) systems. By integrating quantum computing with V2G infrastructures, the proposed model addresses the increasing computational complexity of large-scale energy optimization in smart grids with high EV penetration. The VQE algorithm is utilized to minimize the system Hamiltonian representing energy cost and grid stability parameters, offering superior performance over classical optimization approaches in terms of scalability and convergence. The paper explores the integration of quantum computing to enhance CaaS subscription models by leveraging quantum optimization algorithms for real-time charging scheduling, demand forecasting, and grid resilience. The CaaS model enables stakeholders including EV fleet operators, energy providers, and grid managers to access quantum computing resources on demand, reducing upfront investment and enabling adaptive, high-performance energy decision-making. Simulation results demonstrate the feasibility and benefits of the quantum-enhanced model, including improved grid resilience, reduced peak load, and enhanced economic returns for participants. By harnessing the power of quantum computing, this study aims to revolutionize EV charging by enabling faster decision-making and improved resource allocation.

Keywords: Charging-as-a-Service · EV energy scheduling · quantum hamiltonian · subscription tiers · V2G · variational quantum eigensolver

1 Introduction

Power distribution across the grid is facilitated by quantum computing. This equilibrium keeps some components from being overworked. It increases the system's resilience and lowers the danger of outages. Quantum computing can be used by engineers to anticipate and stop power grid disasters. This can be done by identifying weak points and taking early action. A model is developed to identify the optimal placement of EVCS by utilizing a novel quantum annealing (QA) algorithm and quantum computation (QC).

The knapsack-based formulation is suggested to address the EV Charging problem by utilizing the quantum approximation optimization algorithm (QAOA) [1]. A thorough summary of the PQC algorithms' present migration status in the automobile industry has been offered [2]. A quadratic unconstrained binary optimization issue of optimal power flow has been solved using quantum annealing [3]. The problems, such as late message delivery, computational complexity, and data congestion, have been analysed [4]. For EV charging systems, Quantum Reinforcement Learning (QRL) maximizes the distribution of charging resources according to current customer demand [5].

The rapid adoption of Electric Vehicles (EVs) has transformed the transportation industry, particularly in fleet-based operations such as logistics, ride-hailing, and public transport. However, scalability, cost-efficiency, and reliability remain key challenges for fleet operators managing large numbers of EVs. Traditional pay-per-use charging models often lead to unpredictable costs, operational inefficiencies, and downtime, making them less viable for fleet-based operations.

To address these challenges, Charging-as-a-Service (CaaS) in a subscription model has emerged as an innovative solution. This model allows fleet operators to subscribe to charging services at fixed or dynamic rates, providing predictable energy costs, priority access to charging stations, and optimized charging schedules. Unlike conventional charging approaches, CaaS ensures a seamless and cost-effective charging experience, reducing range anxiety and improving fleet efficiency.

The success of EVs in a market depends on figuring out the best location for EVCS. To address this issue, a new quantum-classical solution has been suggested [6]. Using atomic search optimization(ASO), this work proposes an efficient planning methodology for electric vehicle (EV) fast-charging stations (CS). The suggested technique binarizes the algorithm using quantum operations and outperforms the current binary ASO algorithm in terms of convergence rate. Quantum Neural Networks (QNNs), which are the result of quantum computing, employ the superposition states of quantum bits to encode data more efficiently [7]. QuantumCharge, a PQC extension, sets the future safe EV charging infrastructure [8]. Quantum annealing is proposed to solve combinatorial optimal power flow problems [9]. A work presents a blockchain and Quantum Reinforcement Learning based optimized Energy Trading (BQL-ET) model for E-mobility [10]. A structured and thorough overview of the current migration state of PQC algorithms in the auto industry is presented [11].

A quantum approximate optimization algorithm (QAOA) is suggested to analyze the effect of electric vehicle (EV) charging stations on the dynamic operations of microgrids [12]. A study shows the gate-based quantum optimization capability for a real-world use case. It compares benchmark results from the quantum approximate optimization algorithm (QAOA) with those from the class [13]. A combinatorial power flow (PF) algorithm and a modified version of the PF algorithm that conforms to the principles of PF analysis, the adiabatic quantum PF algorithm (AQPF). Both use model formulations and Quadratic Unconstrained Binary Optimization (QUBO) [14]. A model for optimization is offered to determine the best location for new EV charging stations [15]. The system's components are modeled using a multi-agent artificial intelligence tool, and an energy management algorithm (EMA) is employed to maximize energy utilization and aid in decision-making in hybrid electric vehicles system [16]. A study provides experimental

research and comprehensive design processes for intelligent, affordable, eco-friendly, and user-friendly electric bike charging systems [17].

1.1 Research Gap

Although Charging-as-a-Service (CaaS) and Vehicle-to-Grid (V2G) technologies are increasingly being explored to enhance electric mobility and grid integration, several critical research gaps persist when attempting to incorporate quantum computing and subscription-based service models:

- Lack of Quantum-Optimized Scheduling in Subscription-Aware V2G Systems:

Existing CaaS-V2G frameworks rely predominantly on classical optimization techniques, which are computationally inefficient for large-scale, dynamic environments involving thousands of electric vehicles (EVs), fluctuating grid loads, and differentiated service tiers. The potential of quantum computing to solve such NP-hard scheduling problems more efficiently has yet to be fully investigated.

- Underdeveloped Tiered Subscription Mechanisms in Grid-Aware Contexts

Most current research either treats EV charging as a uniform service or applies simple pricing models without accounting for differentiated user priorities and service guarantees that subscription models can offer. The integration of multi-tiered subscription strategies with quantum decision-making remains largely unexplored.

- Insufficient Real-Time Optimization under Uncertainty:

Classical solvers often fall short in adapting to rapidly changing conditions such as EV arrival/departure times, variable renewable generation, or fluctuating energy demands. Quantum algorithms offer probabilistic sampling and parallelism, yet their application to stochastic or adaptive CaaS planning in V2G systems remains limited in the literature.

1.2 Contributions

This work introduces a novel quantum computing-based charging-as-a-service (CaaS) framework with a Subscription Model tailored for Vehicle-to-Grid (V2G) systems, offering a pioneering approach to address the challenges of real-time energy optimization, user-centric flexibility, and secure transaction management in modern smart grids. The key contributions of this work are as follows:

1.2.1 Quantum-Enabled Optimization for V2G Energy Management

We propose a quantum computing paradigm leveraging quantum annealing or variational quantum algorithms (VQA) to solve the high-dimensional, nonlinear optimization problems involved in dynamic energy scheduling between electric vehicles (EVs) and the grid, outperforming classical solvers in scalability and response time.

1.2.2 Subscription-Based Charging-As-A-Service (CaaS) Architecture

A unique service model is introduced where EV users subscribe to tiered charging plans (e.g., peak/off-peak, priority access, green energy guarantees), enabling predictable revenue for service providers and flexible, contract-based engagement for consumers. The simulation results demonstrate that the proposed model significantly improves grid load balancing, energy cost efficiency, and EV owner satisfaction compared to traditional centralized charging models without quantum optimization.

2 Charging-As-A-Service with Subscription Model for V2G

As electric vehicle (EV) adoption increases, Charging-as-a-Service is emerging as a viable business model to provide reliable and cost-effective charging solutions. With the integration of Vehicle-to-Grid (V2G) technology, EVs can not only draw power from the grid but also supply energy back, creating a dynamic energy ecosystem. However, current charging models often lack flexibility in pricing, grid interaction, and fleet management. Charging-as-a-Service (CaaS) with a subscription model is formulated as an optimization problem where revenue, costs, and customer demand interact dynamically. Below is a structured mathematical framework for this business model. A subscription-based CaaS model for V2G-enabled EVs allows users (individuals and fleets) to access charging services at fixed or dynamic rates, while also benefiting from energy trading opportunities. This approach optimizes grid stability, reduces energy costs, and enhances renewable energy utilization.

2.1 Notation and Definitions

N = Number of subscribed users.

S = Subscription fee per user per month (\$)

C_f = Fixed costs (infrastructure, maintenance, rent, etc.).

C_v = Variable cost per charging session.

λ = Average number of charging sessions per user per month.

P_c = Average power consumed per charging session (kWh).

C_e = Cost per kWh of electricity(\$/kWh)

R = Total revenue per month.

C_t = Total cost per month.

2.2 Revenue Model

The total revenue is derived from the monthly subscription fees:

$$R = N \cdot S \quad (1)$$

where:

R increases with more subscribers N .

S can be static or dynamic, based on demand.

2.3 Cost Model

The total cost per month consists of fixed costs and variable costs.

$$C_t = C_f + N.\lambda.(C_v + P_c.C_e) \quad (2)$$

where fixed costs C_f include station installation, maintenance, software, and rent. Variable costs depend on energy consumption P_c per session costs, and electricity price.

$N.\lambda.$ represents the total number of charging sessions per month. This formula helps set the optimal subscription fee to balance user demand and profitability. A mathematical framework for CaaS with a subscription-based V2G model must consider energy pricing, power flow, battery degradation, grid stability, and revenue optimization. Below are key equations for modeling charging, discharging, subscription pricing, and optimization in a V2G-enabled system. The power flow between the EV battery and the grid must be balanced:

$$P_{net,i}(t) = P_{charge,i}(t) - P_{discharge,i}(t) \quad (3)$$

where, $P_{net,i}(t)$ = Net power drawn from the grid by EV i at time, $P_{charge,i}(t)$ = Power supplied to EV i from the grid at time t .

For grid stability, the total energy exchanged by all EVs should be balanced:

$$\sum_{i=1}^N P_{net,i}(t) = P_{grid}(t) - P_{renewable}(t) \quad (4)$$

$P_{grid}(t)$ = Total power available from the grid at time,

$P_{renewable}(t)$ = Renewable energy contribution at time t ,

2.4 Problem Formulation

We aim to optimize charging/discharging schedules and subscription plan allocations using Variational Quantum Eigensolver, considering: subscription tiers (e.g., Basic, Premium), charging costs & energy prices, user demands, grid constraints, vehicle availability, and battery states.

2.4.1 Decision Variables

$x_{i,t} \in \{0,1\}$: Binary variable indicating if vehicle i is charging (1) or not (0) at time t .

$y_i^s \in \{0,1\}$: Subscription selection variable for vehicle i under plan s .

p_t : Price of electricity at time t

E_i : Energy demand of vehicle i .

B_i^{max}, B_i^{min} : Battery capacity bounds for vehicle i .

Quantum variables are mapped as $\hat{x}_{i,t}$, and \hat{y}_i^s (Pauli-Z spin representation).

2.4.2 Objective Function

The cost function for Charging-as-a-Service (to be minimized):

$$\min_{\{x_{i,t}, y_i^s\}} \sum_i \sum_t p_t \cdot x_{i,t} \cdot \Delta t + \sum_i \sum_s C_s y_i^s \quad (5)$$

C_s : Fixed subscription cost for tier s

Δt : Time slot duration

2.4.3 Constraints

a. Energy Demand Satisfaction:

$$\sum_t x_{i,t} \cdot r_i \cdot \Delta t \geq E_i \forall i \quad (6)$$

r_i : Charging rate for vehicle i

b. Single Subscription Tier per Vehicle:

$$\sum_s y_i^s = 1 \forall i \quad (7)$$

c. Battery Limits:

$$B_i^{\min} \leq \sum_t x_{i,t} \cdot r_t \cdot \Delta t \leq B_i^{\max} \quad (8)$$

3 Proposed VQE for Energy Cost Optimization in V2G CaaS

The Variational Quantum Eigensolver (VQE) offers a novel approach for energy cost optimization in Charging-as-a-Service (CaaS) systems integrated with Vehicle-to-Grid (V2G) networks. In this context, the optimization challenge of minimizing electricity costs while satisfying vehicle energy demands and selecting appropriate subscription models is encoded as a cost Hamiltonian using a Quadratic Unconstrained Binary Optimization (QUBO) formulation. VQE transforms this cost function into a quantum Hamiltonian, representing the energy landscape of the system. A parameterized quantum circuit (ansatz) is prepared, and its expected energy is measured iteratively through quantum-classical feedback loops. Classical optimizers adjust the circuit parameters to minimize this expectation, effectively finding the optimal EV charging and subscription configurations. This hybrid quantum-classical technique enables efficient exploration of complex, high-dimensional solution spaces that are often intractable for classical methods, making VQE a promising tool for real-time, scalable energy cost management in CaaS and V2G infrastructures. All the above constraints and objectives are considered in a Quadratic Unconstrained Binary Optimization (QUBO) form.

Let: $X = \{x_{i,t}, y_i^s\}$ a vector of all binary variables.

The QUBO problem becomes

$$\min_X X^T Q X \quad (9)$$

where Q is a symmetric matrix encoding: Linear terms (costs), Quadratic terms (interactions), Penalties for violating constraints (using Lagrange multipliers).

3.1 Penalty-Based Reformulation:

$$Q = Q_{cost} + \lambda_1 Q_{energy} + \lambda_2 Q_{subscription} + \lambda_3 Q_{battery} \quad (10)$$

where each Q_* models a component of the problem, and λ_i are penalty weights.

In QUBO, the cost function maps to a Hamiltonian:

$$\hat{H}_c = \sum_{(i,j)} Q_{ij} \hat{Z}_i \hat{Z}_j + \sum_i Q_{ii} \hat{Z}_i \quad (11)$$

\hat{Z}_i is the Pauli-Z operator for qubit i . The minimum eigenvalue corresponds to the optimal solution. We want to minimize the total energy cost of charging/discharging EVs under different time-of-use pricing and subscription models, while satisfying vehicle energy demands and grid constraints. Form the Cost Function (Energy + Subscription Fees),

$$C(x, y) = \sum_{i,t} p_t \cdot x_{it} \cdot r_i \cdot \Delta t + \sum_{i,s} C_s \cdot y_i^s \quad (12)$$

$$\hat{H} = \sum_i Q_{ii} \frac{(I - Z_i)}{2} + \sum_{i < j} Q_{ii} \frac{(I - Z_i)(I - Z_j)}{4} \quad (13)$$

This Hamiltonian represents the total cost to be minimized as an energy expectation. We now run VQE to solve:

$$\min_{\theta} \langle \psi(\theta) | \hat{H} | \psi(\theta) \rangle \quad (14)$$

where: $\psi(\theta) = U(\theta)|0\rangle$: parameterized quantum circuit (ansatz)

\hat{H} : cost Hamiltonian encoding energy pricing and constraints

We iteratively update θ to minimize the expected energy:

$$\psi(\theta) = \text{Total Cost of Charging} + \text{Subscription} \quad (15)$$

$$\min_{\theta} \langle \psi(\theta) | \hat{H}_{caas} | \psi(\theta) \rangle \quad (16)$$

where \hat{H}_{caas} encodes the energy prices, subscription costs, and operational constraints of the V2G system.

4 Results and Discussion

A Charging-as-a-Service (CaaS) platform offers three subscription tiers Basic, Standard, and Premium, each with different monthly fees, energy limits, and per-kWh overage rates. The simulation evaluates user energy consumption, revenue generation, and service utilization.

4.1 Basic Tier

In this tier scheme, most users consumed energy close to the free quota. The overage revenue is minimal, indicating low additional usage. This tier primarily drives volume-based revenue through a large user count rather than energy overage, with the following details: Subscribers: 100, Monthly Fee: \$20, Included Energy: 50 kWh, Average Usage: – 52 kWh.

4.2 Standard Tier

Moderate energy overage results in a healthy mix of fixed and variable revenue. This tier is balanced in terms of cost and user consumption behavior. It offers a steady revenue stream with manageable energy loads with the following details: Subscribers: 70, Monthly Fee: \$40, Included Energy: 100 kWh, Average Usage: – 106 kWh.

4.3 Premium Tier

Users in this tier are heavy energy consumers, frequently exceeding their quota. Despite the smaller user base, this tier contributes significantly to total revenue through both fixed fees and overage charges with the following details. Subscribers: 30, Monthly Fee: \$60, Included Energy: 200 kWh, Average Usage: – 218 kWh.

4.4 Energy Delivery Analysis

Energy delivery aligns with expected usage patterns. Premium users are the most energy-intensive. This data is critical for charging station planning, ensuring capacity matches usage. Total Energy Delivered: –16,584 kWh, Basic: –5,213 kWh, Standard: –7,437 kWh, Premium: –3,934 kWh.

4.5 Revenue Analysis

The Standard Tier is the top revenue generator due to its balance between user count and energy overage. Premium users bring high per-user revenue despite being fewer in number. The analysis says: Total Revenue: –\$7,742.90 and Revenue Distribution: Basic Tier: –28%, Standard Tier: –41%, Premium Tier: –31%. We simulate a Variational Quantum Eigensolver (VQE) optimization process for minimizing the total cost in a Charging-as-a-Service (CaaS) system with 5 Electric Vehicles (EVs), 4 time slots, and binary subscription selection. The θ is updated to minimize the expected energy (14)–(16), Fig. 1.

The following data is chosen for a realistic CaaS system over a month.

$N = 200$ users, $S = 25$ (\$), $C_f = 3000$ \$, $C_v = 1.50$ \$, $\lambda = 10$ sessions, $P_c = 15$ kWh, $C_e = \frac{0.12}{\text{kwh}}$ The total revenue per month = $200 \times 25 = 5000$ \$, Total monthly sessions = 200×10 sessions, Total kWh = 30000 kWh, Electricity Cost = 3600 \$. In a Quantum Hamiltonian formulation using the QUBO model for a Charging as a Service (CaaS) system with 5 Electric Vehicles (EVs) and 4 time slots, the objective was

to minimize the total cost while satisfying user energy demands. In a Charging-as-a-Service (CaaS) framework, electric vehicle (EV) charging is managed dynamically using service-based models. A Variational Quantum Eigensolver (VQE)-inspired optimization process minimizes the total cost in a Charging-as-a-Service (CaaS) system.

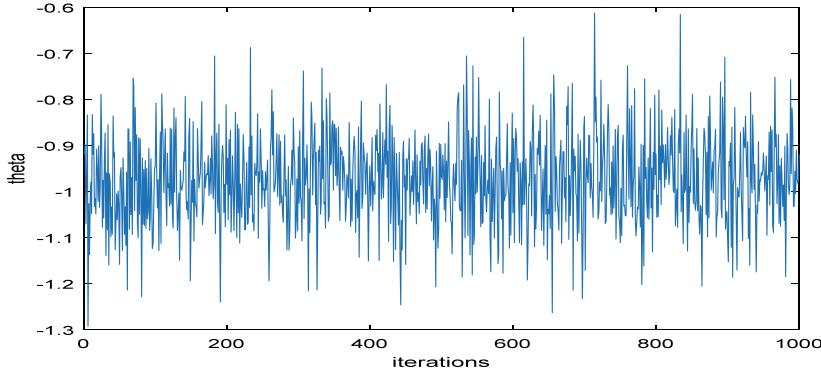


Fig. 1. A The variation θ over 1000 iterations to minimize expected energy

It handles binary decision variables to optimize energy usage, minimize cost, and enhance convenience for users and service providers. The scheduling of EV charging is a core component that ensures efficient grid utilization and customer satisfaction.4 discrete time intervals (T1, T2, T3, T4) over a typical operational window (e.g., 6 AM to 10 AM) are considered, Table 1. Each time slot has a maximum power availability constraint equal to 20 kWh per day. Charging one EV per slot delivers 10 kWh, and each charging session incurs an additional \$0.50 session cost.

Table 1. Illustrative Charging and V2G schedule (kWh)

EV\Time Slot	T1 (6–7 AM)	T2 (7–8 AM)	T3 (8–9) AM	T4 (8–9)AM	Net Energy Change
EV1	5(charge)	-2(V2G)	-	-	+3
EV2	-	-3(V2G)	5(charge)	-	+2
EV3	-	-	6(charge)	-3(V2G)	+3
EV4	3(charge)	3(charge)	-2(V2G)	-	+4
EV5	6(charge)	-	-	-4(V2G)	+2
Grid Net Flow	+14	0	+9	0	

+ve values = energy drawn from grid (charging)

-ve values = energy injected to grid (V2G)

4.6 Key Features of the Schedule

- Load Balancing: Charging loads are distributed to prevent congestion at any time slot.
- Fair Allocation: Each EV is allocated energy according to its requirement without exceeding grid constraints.
- Flexible Timing: Partial charging across multiple slots allows better resource utilization.
- Grid Efficiency: Smooth load profiles prevent transformer overloading and reduce peak shaving needs.
- Grid Support: V2G is utilized during T2, T3 and T4 when grid demand is likely high.

5 Conclusion

The integration of quantum computing into a Charging-as-a-Service (CaaS) framework with a subscription-based model for Vehicle-to-Grid (V2G) systems presents a transformative approach to managing the growing complexity of electric vehicle (EV) energy services. By leveraging the computational advantages of quantum algorithms, such as the Quantum Approximate Optimization Algorithm (QAOA) and Variational Quantum Eigensolver (VQE), the proposed model enables real-time, scalable, and efficient optimization of EV charging and discharging schedules across diverse service tiers. This not only improves the quality of service for end-users through differentiated subscription models but also enhances grid resilience by dynamically balancing energy supply and demand under uncertainty. The mathematical and algorithmic framework formulated herein addresses critical challenges in conventional V2G optimization, including high-dimensional scheduling, non-linear constraints, and stochastic energy demands. Overall, quantum-enhanced CaaS systems represent a promising frontier in smart grid innovation, with the potential to accelerate sustainable mobility, improve energy efficiency, and support the future of resilient, intelligent power systems.

Acknowledgments. This research work is supported by council of technology . U.P (Project Id 3246) and Department of Higher Education , U.P. (Project Id :78/2024/1015).

References

1. Kea, K., Huot, C., Han, Y.: Leveraging knapsack QAOA approach for optimal electric vehicle charging. *IEEE Access* **11**, (9109964–109973) (2023). <https://doi.org/10.1109/ACCESS.2023.3320800>
2. Lohmiller, N., Kaniewski, S., Menth, M., Heer, T.: A survey of post-quantum cryptography migration in vehicles. *IEEE Access*, **13**, (10160–10176) (2025). <https://doi.org/10.1109/ACCESS.2025.3528562>
3. Morstyn, T.: Annealing-based quantum computing for combinatorial optimal power flow. *IEEE Trans. Smart Grid*, **14**(2), (1093–1102) (2023). <https://doi.org/10.1109/TSG.2022.3200590>

4. Khalid, H., et al.: RAVEN: robust anonymous vehicular end-to-end encryption and efficient mutual authentication for post-quantum intelligent transportation systems. *IEEE Trans. Intell. Transp. Syst.* **25**(11), 17574–17586 (2024). <https://doi.org/10.1109/TITS.2024.3416060>
5. Xu, H., Zhang, A., Wang, Q., Hu, Y., Fang, F., Cheng, L.: Quantum Reinforcement Learning for real-time optimization in electric vehicle charging systems. *Appl. Energy*, **383**, (125279) (2025). ISSN 0306–2619
6. Rao, P.U., Sodhi, B.: Hybrid quantum-classical solution for electric vehicle charger placement problem. *Soft Comput.* **27**, 13347–13363 (2023). <https://doi.org/10.1007/s00500-022-07478-x>
7. Hairun, X., Zhang, A., Wang, Q., Yang, H., Fang, F., Cheng, L.: Quantum reinforcement learning for real-time optimization in electric vehicle charging systems. *Appl. Energy* **383**, 125279 (2025)
8. Kern, D., Krauß, C., Lauser, T., Alnahawi, N., Wiesmaier, A., Niederhagen, R.: Post-Quantum Cryptography for Electric Vehicle Charging (2023)
9. Morstyn, T.: Annealing-based quantum computing for combinatorial optimal power flow. *IEEE Trans. Smart Grid* **14**(2), 1093–1102 (2023). <https://doi.org/10.1109/TSG.2022.3200590>
10. Kumar, M., Dohare, U., Kumar, S., Kumar, N.: Blockchain based optimized energy trading for e-mobility using quantum reinforcement learning. *IEEE Trans. Veh. Technol.* **72**(4), 5167–5180 (2023). <https://doi.org/10.1109/TVT.2022.3225524>
11. Lohmiller, N., Kaniewski, S., Menth, M., Heer, T.: A survey of post-quantum cryptography migration in vehicles. *IEEE Access* **13**, 10160–10176 (2025). <https://doi.org/10.1109/ACCESS.2025.3528562>
12. Jing, H., Wang, Y., Li, Y., Du, L., Wu, Z.: Dynamics analysis of microgrids integrated with EV charging stations based on quantum approximate optimization algorithm. In: 2022 IEEE Transportation Electrification Conference & Expo (ITEC), pp. 574–578. Anaheim, CA, USA (2022). <https://doi.org/10.1109/ITEC53557.2022.9813820>
13. Federer, M., Müssig, D., Klaiber, S., Lenk, S.: Application-oriented quantum computing benchmark for an electromobility use case. In: 2022 IEEE International Conference on Quantum Computing and Engineering (QCE), pp. 749–752. Broomfield, CO, USA (2022). <https://doi.org/10.1109/QCE53715.2022.00105>
14. Kaseb, Z., Möller, M., Vergara, P.P., et al.: Power flow analysis using quantum and digital annealers: a discrete combinatorial optimization approach. *Sci. Rep.* **14**, 23216 (2024). <https://doi.org/10.1038/s41598-024-73512-7>
15. Subramanian, N., Alali, H.J., Bin Ghaith Alsuwaidi, A.R.: EV charger placement optimization using D-Wave quantum computing solvers. In: 2023 IEEE International Smart Cities Conference (ISC2), pp. 1–4. Bucharest, Romania (2023). <https://doi.org/10.1109/ISC257844.2023.10293427>
16. Mansouri, N., et al.: Control and optimization of hydrogen hybrid electric vehicles using GPS-based speed estimation. *Electronics*, **14**(1), 110 (2025). <https://doi.org/10.3390/electronics14010110>
17. Ilahi, T., et al.: Comprehensive design analysis of economical e-bike charger with IoT-empowered system for real-time parameter monitoring. *J. Adv. Transp.* 2387983. Advance Online Publication (2025). <https://doi.org/10.1155/2024/2387983>

Multi-Agent Systems and Local Market Coordination



Towards ICT-Enabled Multi-agent Based Operations in Local Energy Communities: A Proof of Concept

Haoyu Huang¹✉, Natascha Fernengel², André Xhonneux¹,
Alexander Holtwerth¹, Michael Hehemann³, Eugen Hoppe³,
Simon Waczowicz², Kevin Förderer², Veit Hagenmeyer²,
and Dirk Müller^{1,4}

¹ Institute of Energy and Climate Research, Energy Systems Engineering (ICE-1),
Forschungszentrum Jülich GmbH, Jülich, Germany
h.huang@fz-juelich.de

² Institute for Automation and Applied Informatics (IAI), KIT, Karlsruhe, Germany
natascha.fernengel@kit.edu

³ Institute of Energy and Climate Research, IET-4: Electrochemical Process
Engineering, Forschungszentrum Jülich GmbH, Juelich, Germany

⁴ Institute for Energy Efficient Buildings and Indoor Climate, E.ON Energy
Research Center, RWTH Aachen University, Aachen, Germany

Abstract. The growing decentralization of energy systems requires scalable, flexible coordination of distributed generation, energy storage, and demand-side flexibility among local energy communities. This work builds upon the agent-based scheduling framework MASSIVE, extending its capabilities to operate in real-world settings. Within the extensive framework, agents participate in the local electricity market by submitting bids based on operational constraints and preferences of local energy components or aggregates, such as a campus. Optimized setpoints derived from market clearing are sent as control signals to physical or simulated assets. To enable the transmission to be modular, interoperable, and responsive in real time, we extend the MASSIVE framework with a lightweight, MQTT-based layer. We validate the applicability of these control signals through a series of experiments involving real hardware and technical and safety constraints. Additionally, a geographically distant battery system was incorporated in real time and it effectively followed market-driven setpoints. The results confirm that a decentralized, agent-based market coordination model facilitates flexible integration of physical energy systems. Plug-and-play functionality, heterogeneous control strategies, and interconnection across regions are collectively offered by the framework, thereby providing a robust path to smart energy communities.

Keywords: Decentralized Optimization · Multi agent system · Local energy community · Market-based Coordination · Smart area · Real world hardware connection · Cloud based optimization

1 Introduction

The concept of local energy communities (LECs) in Europe is undergoing a steady increase in its prevalence. By 2030, LECs are expected to possess up to 21% of installed solar- and 17% of wind-capacity [4]. However, this growing decentralization poses significant challenges for local grid stability, especially as traditional distribution networks were not designed for high shares of bidirectional energy flow [20].

A key notion introduced to address this challenge is flexibility, which describes the ability of a system to adapt generation or demand in response to expected and unforeseen fluctuations [1]. Flexibility is particularly important in LECs, where coordinated actions such as load shifting, energy storage, or curtailment can help maintain balance and support grid resilience [6].

Coordinating these flexibility potentials requires appropriate control strategies. Centralized approaches formulate the entire community as a single optimization problem, aiming for globally optimal solutions. Although this method is effective in smaller systems, there is a limit to its scalability when it comes to more complex systems involving greater participation [35]. Consequently, it often becomes impractical for real-world applications [23].

To overcome these limitations, research has shifted towards decentralized optimization, for which local subproblems are solved independently by each system element. This not only reduces computational effort but also supports modularity and system scalability [13]. In this context, we use the term component to refer to any logical aggregation of sensing, control, and actuation capabilities that can make autonomous decisions. This may represent a building-level system, renewable generation, or a storage unit. When equipped with real-time metering, sensors, and communication capabilities, such components can autonomously monitor their operation and interact with other entities within the smart energy community [7].

These capabilities enables market-based coordination, in which components respond to price signals to optimize both local objectives and system-wide outcomes. Properly designed local energy markets have the potential to align individual incentives with social welfare, supporting the transition toward a more sustainable, decentralized energy system [23, 33].

This work builds upon a previously developed framework, Multi-Agent Scheduling Solution In a Virtual Environment (MASSIVE) [11]. Within the MASSIVE framework, each autonomous agent represents one or more energy components and handles communication and forecasting. Local optimization is utilized for components with flexibility potential. These agents collectively form a multi-agent system (MAS) that enables decentralized, market-based coordination [21].

For rapid experimentation, the framework initially integrated agents and the local energy market within a monolithic simulation script executed sequentially under perfect foresight, thereby constraining its utilization to theoretical studies. Despite the implementation of an Message Queuing Telemetry Transport (MQTT)-based communication layer to support future deployment modes, ini-

tial research remained in the tightly coupled version. In this work, the communication layer is adapted for real-time operation and validated under real-world conditions with live message exchange.

This study investigates the following research questions:

- To what extent can market-based multi-agent coordination enable reliable energy balancing across distributed resources in real-time settings?
- What role does MQTT communication infrastructure play in enabling scalable and flexible agent interaction?
- How accurately can market-cleared setpoints be applied to physical hardware, and what constraints must be considered in practice?

To explore these questions, we implement and evaluate a proof of concept by making use of the extended MASSIVE framework. Our contributions are demonstrated through a stepwise validation process, including simulated agent interaction, hardware validation with market-derived control profiles and the real-time market participation of a remotely located physical battery storage system.

This paper is structured as follows: Sect. 2 presents and distinguishes related work and the current state of technology. In Sect. 3, the infrastructural basis of our experiment is introduced. In Sect. 4, the architecture of the MASSIVE framework as well as the tools and infrastructures that were used are laid out. In Sect. 5, the results of the experiments are presented. These results are interpreted and discussed in Sect. 6. Section 7 concludes this work and motivates for future work building on the proof-of-concept presented in this paper are explored.

2 Related Work

In the early stages of research on multi-agent systems (MAS), McArthur et al. [21] proposed the utilization of MAS in power engineering in 2007. Kumar Nunna et al. [18] developed a model for a virtual energy market for microgrids, introducing a two-level agent architecture to analyze supply–demand mismatches among distributed energy resources. Focusing on increasing energy self-sufficiency in communities, Reis et al. [27] and Prasad et al. [25] applied MAS to coordinate energy communities. Prasad et al. [25] integrated deep reinforcement learning to examine cooperative behavior among buildings, while Reis et al. [27] used real-world data to explore fair cost and benefit allocation within communities. Stennikov et al. [31] extended the MAS modeling framework to encompass centralized generation, thereby facilitating analysis of the interaction between centralized and distributed energy generation. Open-source MAS frameworks, including AMES [19], AMIRIS [28], and ASSUME [12], facilitate a thorough examination of energy wholesale market coordination. The aforementioned frameworks are lacking at this time in the incorporation of physical assets. In order to facilitate real-time control, it is necessary to implement communication protocols that are well-suited to distributed and time-sensitive energy systems. In their 2017 study, Ozgur et al. [24] implemented an MQTT-based frame-

work to coordinate remote cyber-physical testbeds. This implementation demonstrated the framework's resilience under cyber-attack scenarios. Jamborsalamati et al. [17] developed a hierarchical MQTT-enabled architecture for autonomous resource allocation and grid demand reduction in smart areas. Estebsari et al. [9] proposed a real-time coordination schema for distributed energy resources and used MQTT as a communication adapter to facilitate interaction between aggregators and system operators. Despite the limitations of the studies, which rely on simulation or restricted testbed validation, they collectively underscore the aptitude of MQTT for real-time, scalable control in distributed energy systems [8].

Conversely, the FlexQGrid project conducted a field experiment in a residential neighborhood, examining household participation in a quota-based demand response program [30]. However, FlexQGrid did not prioritize component-level optimization or incorporate sector-coupling technologies, such as hydrogen storage, which are imperative for achieving seasonal flexibility.

3 Research Environment

Functioning as a LEC, the Living Lab Energy Campus (LLEC) at Forschungszentrum Jülich Germany (FZJ) serves as a flexible testbed for the exploration and evaluation of innovative monitoring and control solution for district energy systems such as LEC with a high share of renewables [2]. Beyond the shelf components, the LLEC hosts multiple advanced technologies for the conversion and storage of energy to explore cross-sectoral flexibility. This includes a low-temperature district heating network with heat pumps and associated thermal storage, large-scale battery energy storage systems and a hydrogen energy infrastructure that encompasses production, storage, and reconversion. The hydrogen infrastructure comprises e.g. a 400 kW proton exchange membrane (PEM) electrolysis test stand, which is capable of investigating different stacks, ranging from in-house developed to commercial products, to characterize their distinct electrochemical behaviors [8]. The test bench employs a programmable logic controller (PLC) from Beckhoff. This controller facilitates the integration of electrolysis into comprehensive energy system test series. In these test series, setpoints are derived from model-based controllers executed in the cloud, enabling a sophisticated and interconnected approach to energy system testing. [8]

However, to effectively utilize this potential, automated, data-driven decision-making is required to determine when and how much of each form of flexibility should be used [3]. This decision-making process starts with internet of things (IoT)-based status monitoring of components in the smart area [26], which generates a rich data set to support demand forecasting. Based on this data, energy management systems coordinate the optimal use of waste heat in the low-temperature district heating network [32] and the optimal interplay of short-term and seasonal storage [14]. A variety of control strategies are used to support this process, ranging from rule-based control to model predictive control (MPC) [15]. The interoperability of components using different field-bus protocols is enabled

by a FIWARE-based information and communication technology (ICT) platform [34], which facilitates local integration and interaction between LECs.

To support cross-regional coordination, geographically distributed co-simulations enable the assessment of LEC flexibility and their dynamic impact on system stability under realistic conditions [5, 22]. At the Karlsruhe Institute of Technology (KIT), a FIWARE Context Broker is hosted to establish a network that connects research centers by exchanging energy data in real time. In collaboration with the German Aerospace Center (DLR) and FZJ, a simulation and testing process is underway for renewable energy generation and storage in future energy systems. This process utilizes real consumer data within the EnergyLab at KIT. In this study, the FIWARE-based data collection functionality was subjected to evaluation during a live test.

4 Methodology

This section delineates the fundamental components of the MASSIVE framework for decentralized coordination in LECs. The methodology is built around three core components. First, the agent architecture models components as autonomous units capable of local optimization and decision-making. Second, the market design facilitates coordination and interaction between agents through market-based mechanisms. Third, the ICT infrastructure provides the real-time communication backbone that connects agents and the marketplace, ensuring seamless data exchange and operational responsiveness. At the end, the setup of the experiment in this paper is explained.

4.1 Agent Modeling and Roles

In order to determine the offer of agents to purchase energy from the local electricity market, it is necessary for each agent to initially generate a power profile forecast of its component. Within the MASSIVE framework, the configuration of such simulation models is facilitated by the utilization of human-readable files, thereby enabling the instantiation of agents as digital twins when provided with real-time input data. The forecasts can be model-specific, as in the case of a photovoltaic (PV) system. For a campus agent, forecasts are derived from the aggregation of multiple simulation results.

Two types of agents are introduced in our architecture. The first type functions primarily as data analyzer, translating real-time environmental data (e.g., ambient temperature, solar irradiance, wind speed) from German Weather Service (DWD) into forecasts of building heat demand or renewable generation. This facilitates the capacity of agents to dynamically adapt to external conditions. The second type of agent not only forecasts their own demand, but also determines and utilizes the flexibility potential of its components (e.g., thermal storages, battery systems, controllable loads) by optimizing their operation to minimize costs in response to local dynamic price signals. In doing so, they

generate time-resolved power profiles that reflect the physical constraints of the system.

Agents are further classified into three categories based on their market behavior. Consumers, such as office buildings, cover their electricity demand exclusively by submitting bids to the marketplace, as power supply contracts are not considered in this study. The second group comprises pure generators, including renewable sources such as PV panels and wind turbines, as well as the public grid itself, which acts as a balancing source during periods of insufficient local generation. The third group is defined as “active prosumer”, which is characterized by its dual capacity to consume and produce energy. These may operate energy storage systems (e.g. batteries or hydrogen storage systems) to engage in energy arbitrage, buying energy when it is cheap and selling it when prices are high.

4.2 Market Design and Clearing Mechanism

In our framework, agent coordination is enabled through market-driven scheduling that operates on a 15-min clearing interval, reflecting the structure of real-world intraday markets such as EPEX Spot Intraday. The market is cleared using the merit-order principle, in which bids are prioritized based on price to balance supply and demand [29]. For each clearing interval, the market compiles the bids from the agents, which represents the energy demand or supply for the subsequent 24 h. At the core of the market-driven scheduling is a linear optimization problem, which processes all collected bids to compute the optimal energy allocation for each clearing interval. Its objective is to minimize the overall operational cost of the LEC, while determining a corresponding market-clearing price for each timestep.

To focus the study on the coordination logic and flexibility aggregation, no bidding strategy is applied within the individual agents. Rather than implementing dynamic or strategic bidding behavior, each agent uses a fixed price for demand and supply. These values remain constant throughout the simulation. This simplification ensures that market outcomes solely reflect the effects of system-level optimization and agent coordination, rather than being influenced by complex or competitive bidding strategies. Despite the use of fixed prices, the market still produces dynamic clearing results, including fluctuations in market price that arise naturally from changing load and generation conditions. This demonstrates that the market design remains responsive and effective in coordinating distributed energy resources, even using simplified bidding strategies.

4.3 Agent Communication Framework

To enable structured and scalable communication within the LEC, a lightweight communication layer is used that is based on the MQTT protocol. In line with MQTT’s publish/subscribe architecture, agents submit their bids by publishing to predefined topics. In order to protect the privacy of its users, agents are not able to communicate directly with each other or accessing data from other

agents. Instead, all agent messages are routed exclusively through a centralized marketplace module that functions as the sole subscriber to agent outputs. At the end of each market-clearing interval, the marketplace publishes the resulting clearing prices and individual load profiles to agent-specific topics.

For direct communication between agents and the energy market, we rely solely on the MQTT protocol to minimize system complexity, reduce latency, and maintain high responsiveness, which are key requirements for real-time market interaction. Beyond the lightweight MQTT communication, we also enabled integration with a FIWARE-based ICT platform for data transmission from hardware to agent. FIWARE [10] offers advanced functionalities such as data contextualization, orchestration, security, and analytics. In the scope of this work, each campus operates its own instance of the ICT platform to collect sensor data from its assets.

In the initial version of the MASSIVE framework, a handshake mechanism was developed to support modularity and runtime flexibility. This mechanism enables agents to dynamically join the market [16]. Upon startup, an agent sends a registration message containing essential metadata to the marketplace that includes the agent in the upcoming market-clearing optimization. Similarly, a deregistration process allows agents to gracefully disconnect from the system. When an agent leaves, it sends a deregistration message, prompting the marketplace to remove its entry from the participant list, thereby ensuring that outdated or inactive agents do not affect future market rounds.

Furthermore, the agents were modularized and containerized using Docker, and each was deployed on an OpenStack-based virtual machine. This architectural shift enabled distributed and parallel execution, continuous runtime, and clear separation between agents and the marketplace instance, thereby supporting both scalability and realistic event-driven message exchange.

Together, these mechanisms enable plug-and-play functionality, allowing agents to be added or removed at runtime without restarting or reconfiguring the entire system. This real-time communication infrastructure provides a robust foundation for coordinating distributed components, and it also supports hardware integration, enabling direct control of physical devices using market-cleared setpoints.

4.4 Experimental Setup

The validation of the proposed decentralized, agent-based coordination framework was carried out in a hybrid experimental environment, combining purely simulative and real-world components. In the interest of simplicity, the grid constraints are not taken into consideration. Figure 1 illustrates the overall structure of the experimental setup.

At first, we established a setup of purely simulative component agents to the market environment. These components included an PV, electric vehicle (EV), building with heat pump, battery energy storage agent and a campus agent (FZJ Agent), which comprises multiple selected components of the FZJ campus. In agents such as the PV agent, bid generation is performed internally

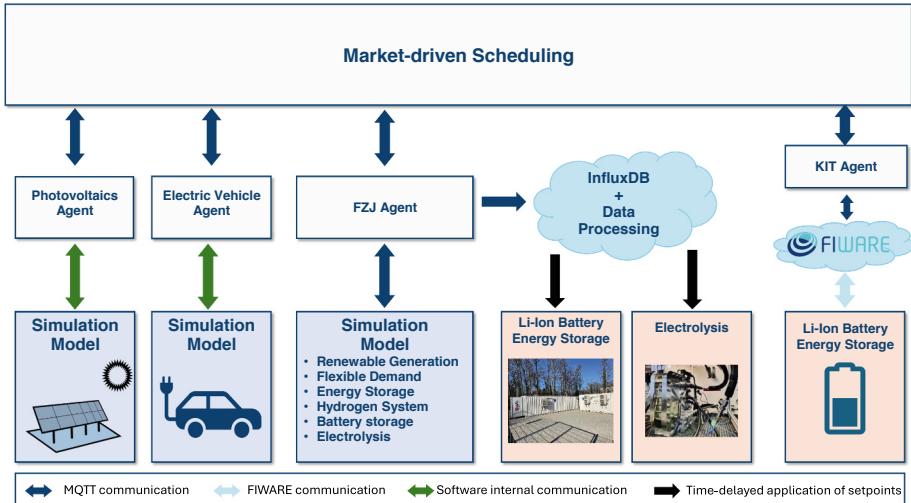


Fig. 1. Excerpt of experimental setup (not all simulative agents are shown)

and the resulting bids are communicated to the marketplace via MQTT. To reduce the solving time, the FZJ agent makes use of the controller as proposed in [15]. For bid generation, FZJ agent aggregates simulation results from multiple models using MQTT-based communication. A grid agent, representing a backup capacity (not displayed in the Fig. 1), is added to make sure that the market clearing is always solvable.

In the second step of the validation process, selected load profiles were applied to physical devices. These agents were modeled using physics-based simulations that closely reflected real hardware deployed on the campus. The application of setpoints was time-delayed, and there was no live connection between the market and the hardware. In our study, a PEM electrolysis system and a lithium-ion battery storage system were used to validate the feasibility of executing the market-cleared load profiles under real operational constraints. To ensure that the new legal requirements for cyber security (e.g. IEC 62443 series) are met as effectively as possible in the future, data transmission within the virtual local area network (VLAN) of the Jülich campus network was implemented and secured with a certificate.

In the third step, to demonstrate the validity of market derived load profile for hardware control, a live connection was established between the MASSIVE framework and a BESS located within the Energy Lab at KIT. The system under test consists of seven lithium iron phosphate battery modules, delivering a usable capacity of 15.4 kWh, and includes a battery management system, an energy management system, and an inverter.

The generation of bids for physical devices was accomplished through the utilization of a customized digital twin, which was configured using the configu-

ration file embedded in the MASSIVE framework. The optimization process was extended to incorporate real-time feedback from the physical system, allowing MPC to respond dynamically to the battery's actual performance, while also accounting for both calendric and cyclic battery aging.

Once a schedule was determined by the market-clearing process, the load profile was translated into setpoints, which were periodically transmitted to the BESS via an MQTT–Modbus bridge developed at the Energy Lab. The agent and the adapter integrating the battery's measurements into the FIWARE-based ICT platform were both containerized using Docker and run in Kubernetes to ensure stability and deployability.

At the same time, using mainly MQTT for data transport, the current state of charge and active power measurements of the battery were used to update the corresponding battery entity in the FIWARE-based ICT platform, which also stores the data in an InfluxDB database. To ensure that the model of the battery and the physical system stayed in alignment, the state of charge was also fed back into the agent to update the battery model before each optimization.

5 Results

This section presents the results of a stepwise validation workflow designed to evaluate the real-world applicability of the proposed MASSIVE framework. The validation process is carried out across three stages of increasing system realism. First, we demonstrate real-time interaction among simulated agents connected via the ICT platform, validating the communication architecture and decentralized coordination mechanism over multiple weeks. Next, we evaluate whether market-cleared setpoint trajectories can be applied to real hardware components under controlled test conditions. Finally, we establish a live connection between market and physical devices, assessing the complete system in a field test.

5.1 Step 1: Simulated Agents

In this subsection, we evaluate the behavior of a simulated agent representing a PV system within our decentralized market framework. The PV agent offers its full generation potential based on the weather forecast to the market at each time step. Based on the agent's internal configuration and real-time weather data, it generates a 24-h forecast with a resolution of 15-min. The purpose of this test is to demonstrate the ability of the framework to handle intermittent renewable generation and apply curtailment when necessary through market-based coordination.

Figure 2 presents a comparison between the offered and realized power outputs of the PV agent over two consecutive days. The offered power (yellow curve) represents the maximum generation potential as the weather forecast, calculated using irradiance data and the technical capacity of the PV system. The realized power (green curve) shows the actual dispatched output as determined by the

PV Agent - Offered vs Realized Power Generation

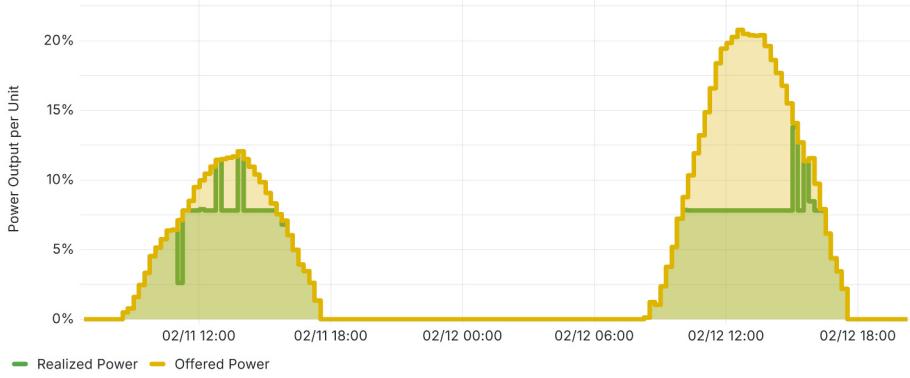


Fig. 2. PV power output - offered power (yellow) vs. actually used power (green) (Color figure online)

market-clearing process. While the offered profile follows the expected solar production curve, the realized output is frequently capped below the maximum, especially during midday peaks. This behavior is evident on both February 11 and 12, where flat-topped plateaus in the realized output indicate intentional curtailment initiated by the marketplace.

5.2 Step 2: System Response to Exemplary Market-Cleared Setpoint Trajectories

In this section, we focus on the evaluation of real-world hardware behavior. The objective of this stage is to determine whether physical components can accurately track market-cleared setpoints under real-time conditions while adhering to component-specific safety constraints.

Despite the utilization of a physics-based simulation model for signal generation, their execution was subject to stringent oversight, necessitating manual adjustment of load profiles to ensure adherence to operational constraints. This included ramp-rate smoothing and exclusion of restricted operating zones, depending on the device.

The initial live test utilized a PEM electrolysis apparatus with a commercial stack that was under development. As part of the assessment, a 9-h profile was applied, which was generated by FZJ agent following the market clearing process. In consideration of the device's operational requirements, particularly with regard to current ramp rates, the original market output was subjected to post-processing to ensure that the current gradient remained below a certain value.

Figure 3 illustrates the behavior of the electrolysis system during a live test conducted on February 6–7, 2025. The red markers represent the original setpoints issued by the market-based coordination framework, which were based on



Fig. 3. Setpoints vs. system response of PEM electrolysis (Color figure online)

an optimization not including the minimal part-load, while the blue line shows the measured current response of the electrolysis. To reflect operational safety guidelines, a restricted operation zone (highlighted in red) was defined below a certain normalized current threshold.

A second live test was performed using a large-scale battery storage system (BESS) of LLEC, which was subjected to an excerpt of the power profile determined by the initial tests with the multiple agents. Although detailed technical specifications of the hardware systems cannot be disclosed, all evaluations were supported by a dedicated data acquisition pipeline based on the ICT platform. Measurements were recorded per second and stored in a cloud-hosted InfluxDB database for subsequent analysis.

5.3 Step 3: Real-Time Connection

In this section, the behavior of a second campus agent (KIT Agent), that was deployed to represent a battery system located at KIT, is evaluated. This agent was integrated into the existing setup, establishing a live connection between the market and the real hardware.

This experiment was carried out over several days. The power measurements and the battery's state of charge over the course of 14 h can be seen in Fig. 4 alongside the schedule that the agent received from the marketplace. Here, the active power measurement as well as the scheduled setpoints are displayed in Watt, with negative values indicating charging, and positive values indicating discharging of the battery system.

At the start of the experiment the battery's SOC was 38 %. After a period of neither charging nor discharging, it was then charged with up to 1 kW starting at 7 pm, and then discharged continually during most of the night, with up to

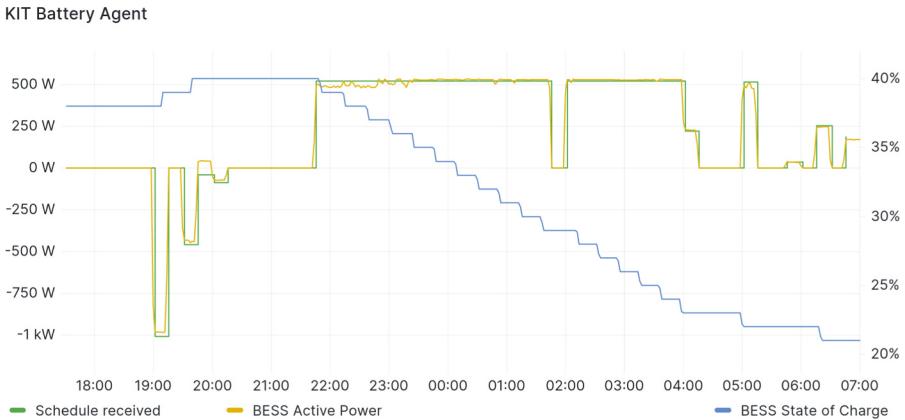


Fig. 4. Active power and state of charge of BESS, alongside the load schedule received from MASSIVE’s marketplace.

500 W. In the morning, periods of inactivity widen and it is only discharged sporadically with 250 W and less. This leaves the BESS with a SOC of 21 % at the end of the experiment.

The measured power values follow the determined schedule closely with a deviation around 35 W occurring only in one timestep in this time frame, and smaller deviations occurring shortly after the battery received a new set-point. Some deviations were expected using this BESS, and overall the observed response was reasonably precise and timely.

As can be seen in this graph, the battery is charged in the evening, when there is still some PV production, just before sundown, and discharged during the night. This results in a reduction of the time between charging and discharging, thereby minimizing losses due to self-discharge over time. In this time frame, the battery is charged up to only 40 % of its total capacity. Charging unnecessarily, which can lead to battery degradation and potential loss of capacity was avoided. This is consistent with the optimization we used, which incorporates parameters such as self-discharge, charging losses and battery aging.

6 Discussion

In this section, the results of the three validation steps will be interpreted and discussed. Additionally, we explore the practical considerations involved in hardware integration, with a particular focus on safety and operational constraints.

6.1 Interpretation of Results

In step 1 (Simulated Agents), the behavior of simulated agents within our decentralized market framework was observed, including varied agents representing

PV and wind systems, battery storage and heat pump-controlled buildings as well as an aggregate agent representing a selected section of the LLEC setup at FZJ. An experiment including an agent representing a PV system shows the offered profile following the PV forecast with most capping occurring around midday, as can be seen in Fig. 2. Thus the capability of the framework to actively curtail PV output when required by system-level constraints or market conditions was demonstrated. Furthermore, the ability of the agent-based market mechanism to integrate fluctuating renewable energy was confirmed.

In step 2 (Hardware reply to predefined load profile), real-world hardware was included and its behavior within our simulated agent framework was observed. In order to ascertain the viability of physical devices in adhering to market-cleared setpoints, a 9-h profile generated by an agent was applied to a PEM electrolysis. The results demonstrate that the electrolysis system is capable of accurately tracking discrete setpoint changes, maintaining stable operation during activation phases and returning to baseline output during inactive periods. The findings indicate that market-based control signals generated by the MASSIVE framework can be effectively followed by relatively slow-responding hardware, such as the electrolysis system, while adhering to safety constraints. However, it also highlights the importance of incorporating such safety restrictions during the design and calibration of the simulation model to ensure realistic and executable control behavior. A second live test using a battery storage system also demonstrated stable and accurate tracking, indicating that market-cleared dispatch signals can be executed reliably by storage assets.

In step 3 (Real-time connection), a physical system was included in real-time, controlling a BESS in a live connection within the MASSIVE framework. A live connection between geographically distributed agents and physical devices was successfully established, confirming the feasibility of the complete system in a field test. Through this experiment, which was running over multiple days with timely and correct responses from the BESS to the scheduled setpoints, it was demonstrated that the resulting control signals are adequate to apply on real-world hardware. Furthermore, the load schedule as an optimization result was reasonable and well-suited for the battery system, taking into account the system's specifications. While this demonstrates the feasibility of using this framework in a live connection, this validation step was only done using a single component. To further validate it, this experiment should be expanded to include more both real-world and virtual components over a longer time-span.

Overall, the outcomes of this research have demonstrated the applicability of the proposed framework. Not only were we able to demonstrate real-time interaction among simulated agents connected via the ICT platform over multiple weeks, validating the communication architecture and de-centralized coordination mechanism. By establishing a live connection between geographically distributed agents and physical devices, we confirm the feasibility of the complete system. It has been demonstrated that market-cleared setpoints derived from optimization can be followed by real hardware components.

Given the prototypical configuration of the framework, the present outcome merely substantiates the feasibility of implementing an ICT-connected MAS for LECs.

6.2 Practical Challenges of Hardware Integration

Maintaining the safety of energy systems such as electrolysis, which pose explosion and pressure risks, is a challenge when developing integrated energy systems. It is important that the developers of the software and the system technology work well together to ensure compliance with safety-relevant and operationally limits, even during operation via the ICT platform. For the electrolysis, these requirements have been jointly integrated into the communication protocols by hardware operator and software developer, and then implemented by the manufacturer in the PLC's MQTT interface program.

It is also imperative to implement additional safety limitations in future developments. For example, configurable limits for electrolysis can be specified within the test stand, customized to the specific stack installed. It is important that these limits should not be override via the MQTT interface. Furthermore, the automation of start-up and shut-down processes, in conjunction with the behavior of the electrolysis with respect to current changes, should be considered in subsequent development.

6.3 System-Level Considerations

This study substantiates the viability of coordinating LECs through market-cleared load profiles and the adaptability of these profiles to physical hardware. However, the validation remains constrained in scope. The evaluation was conducted on a standalone basis, encompassing a single BESS and electrolysis, and the outcomes were applied exclusively over brief time frames. Consequently, the system's behavior concerning larger-scale participation and seasonal dynamics remains to be explored. To address this, future work should expand the number of simulated agents. The efficiency of the central market instance in MAS-SIVE ensures the computational tractability of simulations involving thousands of agents [11]. Additionally, the incorporation of supplementary real hardware will introduce realistic uncertainty and facilitate the testing of agents with more complex optimization strategies, such as MPC.

A further crucial constraint lies in the utilization of a simplified grid model, which assumes unrestricted power flows. To capture grid constraints, the MAS-SIVE framework's multi-round clearing feature can be adapted to account for capacity prices [11]. An additional clearing round could reflect distribution limits, with the grid agent raising capacity prices during congestion to influence market outcomes accordingly.

Adopting these measures will enable the system to evolve in a more realistic direction, thereby ensuring its applicability for broader implementation in future smart energy infrastructures.

7 Conclusion

Increasing penetration of renewable energy sources necessitates the provision of flexibility, which can be offered by energy communities and smart areas. In this work, we propose a decentralized agent-based control framework using a market-driven scheduling, powered by a lightweight ICT platform. Its applicability under close-to-real conditions in smart energy system laboratories is validated, establishing a testbed for future exploration of interacting local energy communities.

Overall, the feasibility of the overall system is confirmed in a field test. We are able to show that market-cleared setpoints derived from simulated digital twins can be safely applied to real hard-ware components within Living Lab Energy Campus at Forschungszentrum Jülich and Energy Lab at Karlsruhe Institute of Technology. Furthermore, we validate the framework's suitability using geographically distributed agents and physical devices. Autonomous agents ensure the scalability of the system.

This successful validation of a decentralized, market-based MAS demonstrates the potential of this concept for the coordination of smart energy areas, addressing a possibility for the provision of flexibility in future energy systems.

Future research could integrate more components and smart areas, exploring the framework's scalability and suitability for more diverse components and participants. Furthermore, the incorporation of distribution-grid constraints through the multi-round clearing functionality of the MASSIVE framework is expected to enhance the realism of the simulation. In the subsequent phase of the study, the economic and environmental impact of agent-based flexibility coordination will be examined, with the incorporation of factors such as wholesale electricity prices and carbon taxes.

Acknowledgments. This work was supported by the Energy System Design (ESD) Program of the Helmholtz Association (HGF) within the structure 37.12.02.

References

1. Babatunde, O., Munda, J., Hamam, Y.: Power system flexibility: a review. *Energy Rep.* **6**, 101–106 (2020). <https://doi.org/10.1016/j.egyr.2019.11.048>. <https://linkinghub.elsevier.com/retrieve/pii/S2352484719309242>
2. Benigni, A., Xhonneux, A., Carta, D., Pesch, T., Muller, D.: On the development of control solutions for local energy communities: an incremental prototyping approach and related infrastructure. *at - Automatisierungstechnik* **70**(12), 1095–1115 (2022). <https://doi.org/10.1515/auto-2022-0034>. <https://www.degruyter.com/document/doi/10.1515/auto-2022-0034/html>
3. Bordons, C., Garcia-Torres, F., Ridao, M.A.: Model Predictive Control of Microgrids. *Advances in Industrial Control*. Springer, Cham (2020). <https://doi.org/10.1007/978-3-030-24570-2>
4. Caramizaru, E., Uihlein, A.: Energy communities: an overview of energy and social innovation (2020). ISBN 9789276107132 ISSN 1831-9424. <https://doi.org/10.2760/180576>. <https://publications.jrc.ec.europa.eu/repository/handle/JRC119433>

5. Carta, D., et al.: VILLASnode-based co-simulation of local energy communities. In: 2022 Open Source Modelling and Simulation of Energy Systems (OSM-SES), pp. 1–6, April 2022. <https://doi.org/10.1109/OSMSES54027.2022.9768933>. <https://ieeexplore.ieee.org/document/9768933>
6. Chicco, G., Riaz, S., Mazza, A., Mancarella, P.: Flexibility from distributed multienergy systems. Proc. IEEE **108**(9), 1496–1517 (2020). <https://doi.org/10.1109/JPROC.2020.2986378>. <https://ieeexplore.ieee.org/abstract/document/9082595>
7. De São José, D., Faria, P., Vale, Z.: Smart energy community: a systematic review with metanalysis. Energy Strategy Rev. **36**, 100678 (2021). <https://doi.org/10.1016/j.esr.2021.100678>. <https://linkinghub.elsevier.com/retrieve/pii/S2211467X2100064X>
8. Emonts, B., et al.: A holistic consideration of megawatt electrolysis as a key component of sector coupling. Energies **15**(10), 3656 (2022). <https://doi.org/10.3390/en15103656>. <https://www.mdpi.com/1996-1073/15/10/3656>
9. Estebsari, A., Mazzarino, P.R., Bottaccioli, L., Patti, E.: IoT-enabled real-time management of smart grids with demand response aggregators. IEEE Trans. Ind. Appl. **58**(1), 102–112 (2022). <https://doi.org/10.1109/TIA.2021.3121651>. <https://ieeexplore.ieee.org/document/9582830/>
10. FIWARE: Open APIs for Open Minds, May 2025. <https://www.fiware.org/>
11. Fritz, J., Riebesel, L., Xhonneux, A.: MASSIVE: multi agent scheduling solution in a virtual environment, June 2025. <https://doi.org/10.5281/zenodo.15768504>. <https://zenodo.org/records/15768504>
12. Fritz, J.M., Riebesel, L., Xhonneux, A., Müller, D.: MASSIVE: a scalable framework for agent-based scheduling of micro-grids using market mechanisms. Energy Inf. **8**(1), 101 (2025). <https://doi.org/10.1186/s42162-025-00558-w>
13. Harder, N., Miskiw, K., Maurer, F., Khanra, M., Parag, P.: ASSUME: agent-based electricity markets simulation toolbox — ASSUME: agent-based electricity markets simulation toolbox. <https://assume.readthedocs.io/en/latest/>
14. He, Z., et al.: A reliability assessment approach for integrated community energy system based on hierarchical decoupling optimization framework. In: 2018 IEEE Power & Energy Society General Meeting (PESGM), pp. 1–5, August 2018. ISSN 1944-9933. <https://doi.org/10.1109/PESGM.2018.8586025>. <https://ieeexplore.ieee.org/document/8586025/>
15. Holtwerth, A., Xhonneux, A., Müller, D.: Modelling of energy systems with seasonal storage and system state dependent boundary conditions using time series aggregation and segmentation. In: 34th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems, ECOS 2021. ECOS 2021 Program Organizers, Taormina, Italy, pp. 344–354 (2022). <https://doi.org/10.5220/062738-0031>. <http://www.proceedings.com/062738-0031.html>
16. Holtwerth, A., Xhonneux, A., Müller, D.: Closed loop model predictive control of a hybrid battery-hydrogen energy storage system using mixed-integer linear programming. Energy Convers. Manage. X **22**, 100561 (2024). <https://doi.org/10.1016/j.ecmx.2024.100561>. <https://linkinghub.elsevier.com/retrieve/pii/S2590174524000394>
17. Jamborsalamati, P., Fernandez, E., Moghimi, M., Hossain, M.J., Heidari, A., Lu, J.: MQTT-based resource allocation of smart buildings for grid demand reduction considering unreliable communication links. IEEE Syst. J. **13**(3), 3304–3315 (2019). <https://doi.org/10.1109/JSYST.2018.2875537>. <https://ieeexplore.ieee.org/document/8509108/>

18. Kumar Nunna, H.S.V.S., Doolla, S.: Multiagent-based distributed-energy-resource management for intelligent microgrids. *IEEE Trans. Industr. Electron.* **60**(4), 1678–1687 (2013). <https://doi.org/10.1109/TIE.2012.2193857>. <https://ieeexplore.ieee.org/document/6179527/>
19. Li, H., Tesfatsion, L.: Development of open source software for power market research: the AMES test bed. *J. Energy Markets* **2**(2), 111–128 (2009). <https://doi.org/10.21314/JEM.2009.020>. <http://www.risk.net/journal-of-energy-markets/technical-paper/2160804/development-source-software-power-market-research-ames-test-bed>
20. Manditereza, P.T., Bansal, R.: Renewable distributed generation: the hidden challenges - a review from the protection perspective. *Renew. Sustain. Energy Rev.* **58**, 1457–1465 (2016). <https://doi.org/10.1016/j.rser.2015.12.276>. <https://linkinghub.elsevier.com/retrieve/pii/S1364032115016597>
21. McArthur, S.D.J., et al.: Multi-agent systems for power engineering applications-part I: concepts, approaches, and technical challenges. *IEEE Trans. Power Syst.* **22**(4), 1743–1752 (2007). <https://doi.org/10.1109/TPWRS.2007.908471>. <https://ieeexplore.ieee.org/document/4349106>
22. Monti, A., et al.: A global real-time superlab: enabling high penetration of power electronics in the electric grid. *IEEE Power Electron. Mag.* **5**(3), 35–44 (2018). <https://doi.org/10.1109/MPEL.2018.2850698>. <https://ieeexplore.ieee.org/document/8458285/>
23. Nagpal, H., Avramidis, I.I., Capitanescu, F., Madureira, A.G.: Local energy communities in service of sustainability and grid flexibility provision: hierarchical management of shared energy storage. *IEEE Trans. Sustain. Energy* **13**(3), 1523–1535 (2022). <https://doi.org/10.1109/TSTE.2022.3157193>. <https://ieeexplore.ieee.org/document/9729638/>
24. Ozgur, U., Nair, H.T., Sundararajan, A., Akkaya, K., Sarwat, A.I.: An efficient MQTT framework for control and protection of networked cyber-physical systems. In: 2017 IEEE Conference on Communications and Network Security (CNS), October 2017, pp. 421–426 (2017). <https://doi.org/10.1109/CNS.2017.8228674>. <https://ieeexplore.ieee.org/document/8228674/>
25. Prasad, A., Dusparic, I.: Multi-agent deep reinforcement learning for zero energy communities. In: 2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe), September 2019, pp. 1–5 (2019). <https://doi.org/10.1109/ISGTEurope.2019.8905628>. <https://ieeexplore.ieee.org/document/8905628/>
26. Redder, F., et al.: Information and communication technologies (ICT) for holistic building energy system operation in living labs: conceptualization, implementation, evaluation, February 2024. <https://doi.org/10.2139/ssrn.4743282>. <https://papers.ssrn.com/abstract=4743282>
27. Reis, I.F., Gonçalves, I., Lopes, M.A., Antunes, C.H.: A multi-agent system approach to exploit demand-side flexibility in an energy community. *Utilities Policy* **67**, 101114 (2020). <https://doi.org/10.1016/j.jup.2020.101114>. <https://linkinghub.elsevier.com/retrieve/pii/S0957178720301089>
28. Schimeczek, C., et al.: AMIRIS: agent-based Market model for the Investigation of Renewable and Integrated energy Systems. *J. Open Source Softw.* **8**(84), 5041 (2023). <https://doi.org/10.21105/joss.05041>. <https://joss.theoj.org/papers/10.21105/joss.05041>
29. Sensfuß, F., Ragwitz, M., Genoese, M.: The merit-order effect: a detailed analysis of the price effect of renewable electricity generation on spot market prices in Germany. *Energy Policy* **36**(8), 3086–3094 (2008). <https://doi.org/10.1016/j.enpol.2008.03.035>. <https://linkinghub.elsevier.com/retrieve/pii/S0301421508001717>

30. Sloot, D., Lehmann, N., Ardone, A.: Explaining and promoting participation in demand response programs: the role of rational and moral motivations among German energy consumers. *Energy Res. Soc. Sci.* **84**, 102431 (2022). <https://doi.org/10.1016/j.erss.2021.102431>. <https://linkinghub.elsevier.com/retrieve/pii/S2214629621005181>
31. Stennikov, V., Barakhtenko, E., Mayorov, G., Sokolov, D., Zhou, B.: Coordinated management of centralized and distributed generation in an integrated energy system using a multi-agent approach. *Applied Energy* **309**, 118487 (2022). <https://doi.org/10.1016/j.apenergy.2021.118487>. <https://linkinghub.elsevier.com/retrieve/pii/S0306261921017086>
32. Stock, J., Althaus, P., Johnen, S., Xhonneux, A., Müller, D.: Method development for lowering supply temperatures in existing buildings using minimal building information and demand measurement data, November 2023. arXiv [arXiv:2311.01800](https://arxiv.org/abs/2311.01800) [eess]. <https://doi.org/10.48550/arXiv.2311.01800>
33. Wang, T., Yamashita, D., Takamori, H., Yokoyama, R., Niimura, T.: A dynamic pricing model for price responsive electricity consumers in a smart community. In: 2013 IEEE Power & Energy Society General Meeting, July 2013, pp. 1–5 (2013). ISSN 1932-5517 <https://doi.org/10.1109/PESMG.2013.6672417>. <https://ieeexplore.ieee.org/document/6672417/>
34. Westphal, L., Schröder, M., Carta, D., Xhonneux, A., Benigni, A., Müller, D.: Development and application of a FIWARE-based ICT-platform for multi-energy systems on building and district level. In: 2024 Open Source Modelling and Simulation of Energy Systems (OSMSES), September 2024, pp. 1–6 (2024). <https://doi.org/10.1109/OSMSES62085.2024.10668993>. <https://ieeexplore.ieee.org/document/10668993>
35. Zhou, Y., Wei, Z., Sun, G., Cheung, K.W., Zang, H., Chen, S.: A robust optimization approach for integrated community energy system in energy and ancillary service markets. *Energy* **148**, 1–15 (2018). <https://doi.org/10.1016/j.energy.2018.01.078>. <https://linkinghub.elsevier.com/retrieve/pii/S0360544218300963>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.





Peer-to-Peer Energy Management Model for Residential Homes

Najmeh Khajoei, Runar Unnþorsson, and Steinn Gudmundsson^(✉)

Faculty of Industrial Engineering, Mechanical Engineering and Computer Science,
University of Iceland, Reykjavík, Iceland
{nak18,runson,steinng}@hi.is

Abstract. A linear optimization model for peer-to-peer energy exchange for residential homes is presented. The model takes into account heterogeneous consumption patterns, shared battery storage and supports intra-group trading among households within the same building, as well as inter-group trading between multiple buildings. Using energy usage data from Icelandic households and meteorological data for the same period, a scenario with two groups of users was simulated, and the effects of direct trading within and between the groups were studied. Results indicate that most of the cost savings arise from intra-group trading instead of inter-group and exchanges with the utility grid. Sensitivity analysis of battery maintenance costs shows that this parameter needs to be estimated carefully.

Keywords: Peer-to-peer energy trading · Optimization · Battery storage · Solar photovoltaic production

1 Introduction

This paper presents an energy trading model for residential buildings, where energy is produced locally using solar photovoltaic panels. Trading occurs between individual households within a single building, across buildings, and between the buildings and the utility grid. This is an example of a peer-to-peer energy (P2P) trading system. In this system, households are categorized as either *prosumers*, who both produce and consume electricity, or *consumers* who only consume electricity. Households within a building can share a single battery for electricity storage, and there is also the possibility of a building having no battery. Assuming that the cost of electricity from the utility grid is higher than the cost of purchasing from neighbours, in the same building, or other buildings, it is possible to lower the total cost of the community by decreasing reliance on the grid. Surplus energy, that is remaining when local demand has been satisfied, can be stored in a local battery, if it is present, sold to neighbours, or exported to the grid.

Recent studies have shown that P2P energy sharing offers several benefits, including reduced electricity costs, lower carbon emissions, increased consumer

empowerment through direct energy trading participation, enhanced demand-side flexibility, and grid stability in systems with significant renewable energy penetration [13]. P2P energy trading models have emerged during the transition toward decentralized and sustainable energy systems, for a recent overview of the field, see e.g., [11] and the references therein.

A frequently used approach to such P2P systems is mathematical optimization, including linear programming, mixed-integer linear programming, and nonlinear programming. Nguyen et al. [10] apply a mixed-integer model for rooftop solar production and energy storage, while Long et al. [9] propose an alternative model for managing community energy storage using nonlinear programming. A peer-to-peer energy-sharing framework for managing both day-ahead and real-time energy across communities of prosumers is proposed in [3]. The model includes inter- and intra-community strategies aimed at minimizing energy costs and maximizing renewable energy usage. Gbadega and Yanxia [6] present a comprehensive P2P model that includes rules for energy transactions, aiming to maximize social welfare while minimizing inequity among trading participants. Alam et al. [1] investigate the impact of P2P energy trading within a microgrid of interconnected smart homes, focusing on cost optimization through shared energy resources. In [7] A two-stage P2P energy trading system between prosumers and consumers based on real-time pricing was implemented to reduce electricity costs through demand response and optimized appliance scheduling. Note that the above list of previous P2P studies is by no means exhaustive, but represents some cases that are related to this study. For more comprehensive reviews, see [12, 13].

This paper proposes a P2P energy model specifically designed for residential buildings in Iceland. While progress towards P2P trading in Iceland has been hampered by lack of incentives and unclear regulatory framework, energy communities are just beginning to form, currently focusing on internal sharing, rather than trading. This work aims to explore mechanisms that could support the future development of energy communities in the Icelandic context. Since residential heating is mainly supplied by geothermal energy, electricity is generally not used for heating purposes. Furthermore, air conditioning is rarely needed, as average summer temperatures range between 10 and 13 °C. As a result, electricity is used primarily for powering household appliances and for charging electric vehicles. Seasonal variation in electricity demand is therefore relatively small, and there is little or no variation in demand due to short-term fluctuations in outside temperature. Moreover, utility electricity prices are currently fixed. These characteristics collectively support the use of a linear optimization model.

To account for household diversity and usage patterns, residential buildings are categorized as shown in Fig. 1. The model differentiates between *prosumers*, households that both produce and consume energy and *consumers*, who only consume energy. Each category is further subdivided into *single* and *multi*-household dwellings (e.g., apartment buildings). Households are also classified based on whether they possess individual or shared battery storage, or none at all. The simulation results are based on electricity demand data from Icelandic

households and matching historical weather conditions. Additionally, a sensitivity analysis was conducted to evaluate the impact of battery depreciation costs.

2 Methods

This section presents a detailed description of the linear programming model for peer-to-peer energy exchange among residential households in Iceland.

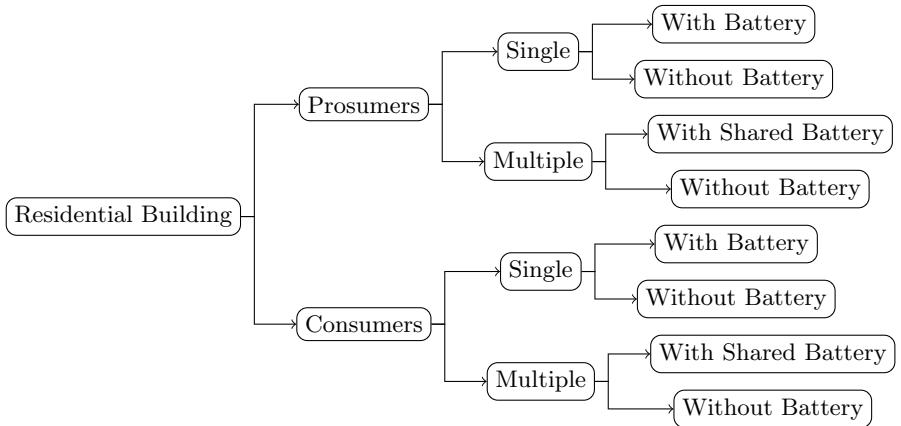


Fig. 1. Classification of residential buildings based on prosumers, consumers and battery storage. A residential building can have only consumers, only prosumers or both.

To account for the scenarios shown in Fig. 1 we consider an apartment building with $N_1 \geq 0$ prosumers and $N_2 \geq 0$ consumers. These users may share a battery storage system, or there may be no battery at all. We allow energy exchange both within a single building and between two or more buildings. We consider two scenarios. The first involves a single building where prosumers can sell surplus energy to each other and share excess energy with consumers within the same building. This forms a local P2P energy market. The second scenario considers multiple interconnected buildings, where users can engage in electricity trading across building boundaries.

We assume there are $\mathcal{R} = 1, \dots, r$ apartment buildings. In each building, some residents may have solar panels, and the building may either have a shared battery or none at all. We assume there are N_1^r residents with solar panels and N_2^r residents without them, where the index r denotes the r -th apartment building. The objective is to fulfill the local demand within each building and to share surplus energy with other buildings, or the utility grid.

2.1 Mathematical Model

In this section we provide a description of the optimization model. We start by defining sets of users and helper functions that identify the location and type of users. Let $P = \{p_1, \dots, p_n\}$ be the set of all prosumers, $C = \{c_1, \dots, c_m\}$ is the set of all consumers and $U = P \cup C$ is the set of all users. Define function $b : u \rightarrow \mathbb{N}$ that maps users to buildings. Two users i and j are said to be *neighbours* if they live in the same building. Let $N_{pp}(i)$ denote the set of prosumers that are neighbours to prosumer i , let $N_{cp}(i)$ denote the set of consumers that are neighbours to prosumer i . Define set $M_{pc}(i)$ as the set of prosumers that are neighbours to consumer i . Analogously define sets $\overline{N}_{pp}(i)$, $\overline{N}_{cp}(i)$ and $\overline{M}_{pc}(i)$ of prosumers that are not neighbours to prosumer/consumer i . The time horizon is indexed by $t = 1, \dots, T$ where T is the length of the control horizon, here $T = 24$ h. The following decision variables are needed: $G^b[t, i]$ is the amount of energy i buys from the grid, $G^s[t, i]$ is the amount of energy i sells to the grid, $E^b[t, i, j]$ is the amount of energy that i buys from j , $E^s[t, i, j]$ is the amount of energy that i sells to j , $c[t, i]$ is battery charging of user i , $d[t, i]$ is the battery discharging of user i , and $B[t, r]$ is the battery level for building r at time t . The following parameters are assumed to be known (fixed): $D[t, i]$ is the demand of user i , $S[t, i]$ is the amount of solar photovoltaic energy produced by user i .

Energy balance for prosumer $i \in P$ gives

$$\begin{aligned} D[t, i] = & S[t, i] + d[t, i] + G^b[t, i] + \sum_{j \in N_{pp}(i)} E^b[t, i, j] + \sum_{j \in \overline{N}_{pp}(i)} E^b[t, i, j] - c[t, i] \\ & - G^s[t, i] - \sum_{j \in N_{pp}(i)} E^s[t, i, j] - \sum_{j \in N_{cp}(i)} E^s[t, i, j] - \sum_{j \in \overline{N}_{cp}(i)} E^s[t, i, j] \\ & - \sum_{j \in \overline{N}_{pp}(i)} E^s[t, i, j] \end{aligned} \quad (1)$$

and energy balance for consumer $i \in C$ gives

$$D[t, i] = d[t, i] + G^b[t, i] + \sum_{j \in M_{pc}(i)} E^b[t, i, j] + \sum_{j \in \overline{M}_{pc}(i)} E^b[t, i, j] \quad (2)$$

Symmetry between buying and selling gives, for $i \in U$ and $j \in P$ gives

$$E^b[t, i, j] = E^s[t, j, i] \quad (3)$$

The current demand provides an upper bound on energy purchases of prosumers, for all $i, j \in P$,

$$G^b[t, i] \leq D[t, i] \quad (4)$$

$$E^b[t, i, j] \leq D[t, i] \quad (5)$$

Solar production provides an upper for prosumer selling, for all $i \in P$ and $j \in U$,

$$G^s[t, i] \leq S[t, i] \quad (6)$$

$$E^s[t, i, j] \leq S[t, i] \quad (7)$$

If two prosumers have surplus energy, they should not trade amongst themselves, i.e., for all $i, j \in P$, if $S[t, i] - D[t, i] > 0$ and $S[t, j] - D[t, j] > 0$ then we add constraints

$$E^b[t, i, j] = E^b[t, j, i] = 0. \quad (8)$$

A prosumer should not buy and sell at the same time, for all $i, j \in P$,

$$G^b[t, i] \cdot G^s[t, i] = 0 \quad (9)$$

$$E^b[t, i, j] \cdot E^s[t, j, i] = 0. \quad (10)$$

These two nonlinear constraints can be linearized by introducing binary variables resulting in a mixed-integer optimization problem. However, if pricing assumptions described below are satisfied, the constraints become redundant and may be omitted from the model. Furthermore, charging and discharging of batteries can not take place at the same time, for all $i \in P$,

$$d[t, i] \cdot c[t, i] = 0. \quad (11)$$

Since both charging and discharging incur some costs, this constraint is also automatically satisfied and can be omitted.

The batteries are modeled as follows, for all buildings r with a shared battery,

$$B^{\min}[r] \leq B[t, r] \leq B^{\max}[r] \quad (12)$$

$$B[t, r] = B[t - 1, r] + \mu_r^{ch} \sum_{i \in P: b(i)=r} c[t, i] - \frac{1}{\mu_r^{dis}} \sum_{i \in U} d[t, i] \quad (13)$$

$$B[T, r] = B[0, r] \quad (14)$$

$$0 \leq \sum_{i \in P} c[t, i] \leq C_r^{ch} \quad (15)$$

$$0 \leq \sum_{i \in U} d[t, i] \leq D_r^{dis} \quad (16)$$

where $B^{\min}[r]$ and $B^{\max}[r]$ are the minimum and maximum battery levels for building r , respectively, μ_r^{ch} and μ_r^{dis} are the battery charge and discharge efficiencies (Table 1).

The objective is to minimize the total cost, namely

$$\begin{aligned}
\text{Cost} = & \sum_{t=1}^T \sum_{i \in P} \left(p_{\text{buy}}^{\text{grid}} \cdot G^b[t, i] - p_{\text{sell}}^{\text{grid}} \cdot G^s[t, i] \right. \\
& + p_{\text{buy-internal}}^{\text{neighbor}} \cdot \sum_{j \in N_{pp}(i)} E^b[t, i, j] - p_{\text{sell-internal}}^{\text{neighbor}} \cdot \sum_{j \in N_{pp}(i)} E^s[t, i, j] \\
& - p_{\text{sell-internal}}^{\text{neighbor}} \cdot \sum_{j \in N_{cp}(i)} E^s[t, i, j] + \text{batt_cost} \cdot (d[t, i] + c[t, i]) \\
& + p_{\text{buy-external}}^{\text{neighbor}} \cdot \sum_{j \in \bar{N}_{pp}(i)} E^b[t, i, j] - p_{\text{sell-external}}^{\text{neighbor}} \cdot \sum_{j \in \bar{N}_{pp}(i)} E^s[t, i, j] \\
& \left. - p_{\text{sell-external}}^{\text{neighbor}} \cdot \sum_{j \in \bar{N}_{cp}(i)} E^s[t, i, j] \right) \\
& + \sum_{t=1}^T \sum_{i \in C} \left(p_{\text{buy}}^{\text{grid}} \cdot G^b[t, i] + p_{\text{buy-internal}}^{\text{neighbor}} \cdot \sum_{j \in M_{pc}(i)} E^b[t, j, i] \right. \\
& \left. + \text{batt_cost} \cdot d_{[t, k, r]} + p_{\text{buy-external}}^{\text{neighbor}} \cdot \sum_{j \in \bar{M}_{pc}(i)} E^b[t, j, i] \right) \tag{17}
\end{aligned}$$

where $p_{\text{sell}}^{\text{grid}}$ is the selling price to the grid, $p_{\text{sell-internal}}^{\text{neighbor}}$ is the selling price to neighbours in the same building, $p_{\text{sell-external}}^{\text{neighbor}}$ is the selling price to users outside the building, and $p_{\text{buy-internal}}^{\text{neighbor}}$, $p_{\text{buy-external}}^{\text{neighbor}}$ and $p_{\text{buy}}^{\text{grid}}$ are the corresponding purchase costs. Since both the objective function and constraints are linear, the resulting optimization problem is linear. The following assumptions are made:

1. $p_{\text{sell}}^{\text{grid}} < p_{\text{sell-external}}^{\text{neighbor}} < p_{\text{sell-internal}}^{\text{neighbor}} < p_{\text{buy-internal}}^{\text{neighbor}} < p_{\text{buy-external}}^{\text{neighbor}} < p_{\text{buy}}^{\text{grid}}$ to ensure that binary constraints in Eqs. 9 and 10 are always satisfied. The values used in the study are listed in Table 1.
2. Simultaneous charging and discharging incur additional costs; thus, even without explicitly enforcing the constraint $d[t, i] \cdot c[t, i] = 0$, the optimization inherently avoids such behavior.
3. A linear approximation to battery costs is used that includes both charging and discharging [3, 8],

Prosumers: $\text{batt_cost} \cdot (d[t, i] + c[t, i])$

Consumers: $\text{batt_cost} \cdot d[t, i]$

The participation willingness index [9] was used as a measure of P2P participation. It is defined as the proportion of prosumers whose energy costs under P2P trading are lower than those under traditional peer-to-grid (P2G) trading.

Table 1. Battery properties and trading price used in the model.

Battery assumptions, r denotes the building		
$B_{[r]}^{\min}$	Minimum battery level	1.05 kWh
$B_{[r]}^{\max}$	Maximum battery level	4.2 kWh
μ_r^{ch}	Battery charge efficiency	0.98
μ_r^{dis}	Battery discharge efficiency	0.96
C_r^{ch}	Maximum charging rate	1.05 kW
D_r^{dis}	Maximum discharging rate	1.05 kW
Trading price assumptions		
$p_{\text{buy}}^{\text{grid}}$	Buying price from grid (price + transmission)	20.93 kr/kWh
$p_{\text{sell}}^{\text{grid}}$	Selling price to grid	4.186 kr/kWh
$p_{\text{buy-neighbor}}$	Internal buying price (neighbor)	14.65 kr/kWh
$p_{\text{sell-neighbor}}$	Internal selling price (neighbor)	10.465 kr/kWh
$p_{\text{buy-external}}$	External buying price (neighbor)	16.744 kr/kWh
$p_{\text{sell-external}}$	External selling price (neighbor)	6.279 kr/kWh
batt_cost	Battery depreciation cost	2.0 kr/kWh

2.2 Properties of Batteries and Solar Panels

The battery parameters are derived from the NE-48D100-NP LiFePO₄ battery, which operates at a nominal voltage of 52.5 V and has a capacity of 100 Ah.

The number of solar panels installed per household varies depending on energy requirements and available space. The total active area increases proportionally with each additional panel. The solar energy output of a single solar panel at time step t is calculated using the model from [4]:

$$S_{[t,i]} = \eta_{\text{PV}} \cdot \eta_{\text{inv}} \cdot A_{\text{pv}[i]} \cdot \psi \cdot \frac{G_{[t]}}{G_{\text{ref}}} [1 + \psi_T (T_{[t]} - T_{\text{STC}})] \quad (18)$$

Here, η_{PV} and η_{inv} represent the conversion efficiencies of the photovoltaic (PV) module and the inverter, respectively, ψ is a factor accounting for environmental and operational losses. ψ_T denotes the temperature coefficient of maximum power, indicating the drop in power output per degree Celsius increase in temperature. Here, $G_{[t]}$ denotes the total radiation incident at time t , G_{ref} is the incident radiation at reference conditions, $T_{[t]}$ is the operating temperature of the PV cells at time t and T_{STC} is the standard test condition temperature. A_{pv} is the panel's active area, which is scaled by the number of panels (n_{panels}) when applicable. The parameter values are given in Table 2.

2.3 Data and Software

Hourly temperature and solar radiation data were obtained from the Icelandic Meteorological Office, recorded in Reykjavik from 1.5.2024 to 30.8.2024.

Table 2. Key Parameters of the Hyundai HG 430W PERC Shingled Full Black - HiE-S430HG(FB) Solar Panel.

Parameter	Symbol	Value	Source
PV Module Efficiency	η_{PV}	0.207	Datasheet (430W model)
Inverter Efficiency	η_{inv}	0.95	Assumed (typical value)
PV Active Area	A_{PV}	2.08 m ²	Datasheet
Derating Factor	ψ	0.9	[4]
Total radiation incident	G	input	Input
Standard incident radiation (STC)	G_{ref}	1000 W/m ²	[4] / STC standard
Temperature Coefficient of Power	ψ_T	-0.0034 °C	Datasheet (-0.34%)
Cell Temperature	T	input	Input
Reference Cell Temp (STC)	T_{STC}	25 °C	[4]/STC standard

Anonymized, proprietary residential electricity consumption data from five households in the same region, collected via smart meters by a local utility provider (sampled at one hour intervals), were used in the simulations.

The optimization problem was formulated using the CVXPY modeling language [5] and solved with the GUROBI solver.

3 Results

To evaluate the model, we consider two groups of users, referred to as Group 1 and Group 2. The model is solved on a 24-h basis over a 122 day period during the summer of 2024, when solar production in Iceland is at maximum. Group 1 includes two prosumers and one consumer, all connected to a shared battery. Group 2 includes one prosumer and one consumer, who also share a battery. Both groups are connected to the utility grid, which sells energy to users and purchases surplus energy from prosumers, both without restrictions. We evaluate three operational settings. (1) *Peer-to-Grid*, where all users are individually connected to the grid (prosumers and consumers) and may share a battery if available. (2) *Intra-trading*, where prosumers can exchange surplus energy within the same group. (3) *Inter-trading*, which extends intra-group to include exchanges with users in other groups. Figure 2 illustrates the simulation setup.

Table 3 presents the simulation results for all three scenarios. As expected, the P2G scenario has the highest cost, followed by intra-trading setting, while inter-trading results in the lowest overall cost. Although the absolute differences in cost are modest, these results are illustrative and reflect the limited scale of the current test case. Additionally, average solar production relative to household energy demand in this scenario is relatively low. The results also indicate that most cost savings are obtained through intra-group resource sharing, with only marginal gains from inter-group trading under the current configuration.

Inter-trading activity remains infrequent in this example, primarily due to high intra-group demand compared to local solar energy production. A representative trading profile illustrating trading between the two groups is shown in Fig. 3. The figure shows clearly how demand for house 3 in group 1 is practically zero in the period. The participation willingness among prosumers reached 100% in both the intra- and inter-trading scenarios (Table 3).

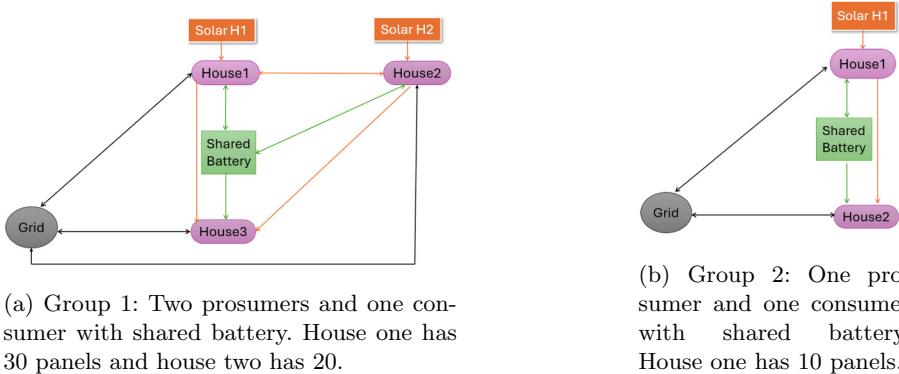


Fig. 2. Simulation setup with group 1 on the left and group 2 on the right. Black arrows denote connections to the grid, green arrows denote connections to the shared battery. Orange arrows denote direct connections that become available in the intra-trading setting. (Color figure online)

Table 3. Comparison of Costs (IKR) for P2G, Inter-trading and Intra-trading. Negative values indicate net profit.

Group	Participant	P2G	Intra-trading	Inter-trading
1	Prosumer 1	2172	-5554	-6460
1	Prosumer 2	34231	31436	34116
1	Consumer 3	46355	53703	51497
1	Subtotal	84181	79584	79153
2	Prosumer 1	2579	-2952	-2283
2	Consumer 2	13280	18422	17188
2	Subtotal	15859	15470	14904
1+2	Total	100040	95054	94057

We then analyzed battery discharge patterns across the 122-day period by aggregating 24-h profiles over time (see Fig. 4). As expected, charging patterns

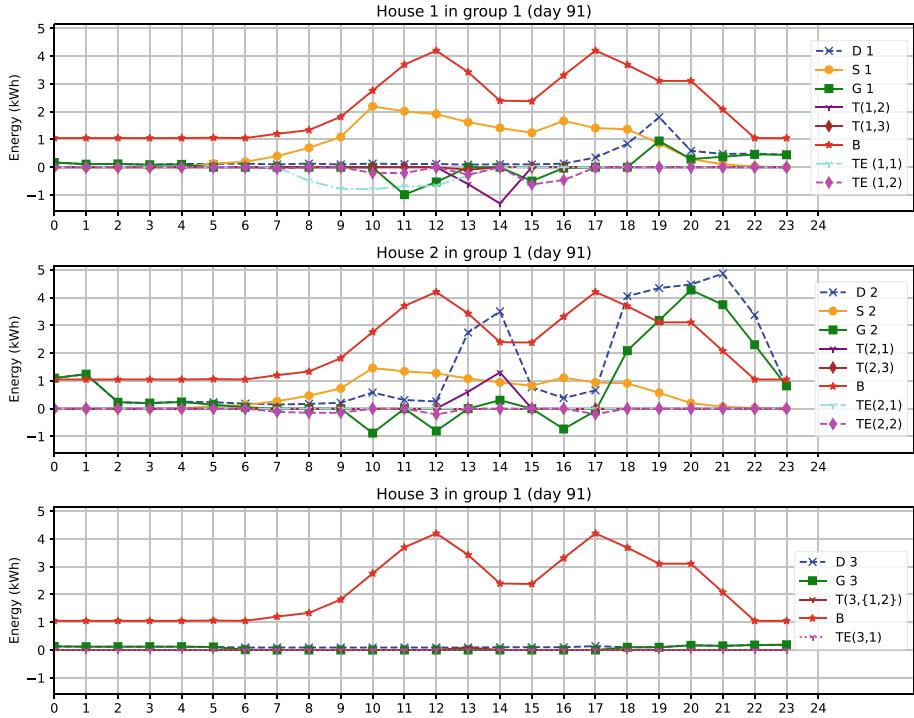


Fig. 3. Example of inter-trading for group 1 in a single 24-h period. Figure legends: Demand (D), Solar production (S), Grid (G), $T(i, j)$ intra-trading from user i to user j , $TE(i, j)$ from user i in group 1 to user j in group 2 and B is the shared battery level. Positive values of G , T and TE denote purchasing, negative values denote selling.

generally align with solar availability during daylight, while discharge activity increases into the evening, peaking near midnight. The discharge behavior reflects the relatively high demand compared to solar production, and is also influenced by the modeling assumption that battery levels must return to their initial state at the end of each 24-h period.

To evaluate the impact of battery depreciation costs, a sensitivity analysis was conducted by repeating the simulations using a range of cost values, including the baseline value of 2 kr/kWh (Table 4).

As expected, the total energy costs for both groups increase as battery depreciation costs rise. Interestingly, as battery costs rise, some prosumers experience net benefit (i.e., negative net costs indicating revenue generation), while consumers disproportionately absorb the additional cost burden. These findings suggest that battery depreciation costs must be carefully estimated and allocated fairly if the system is to be implemented in practice.

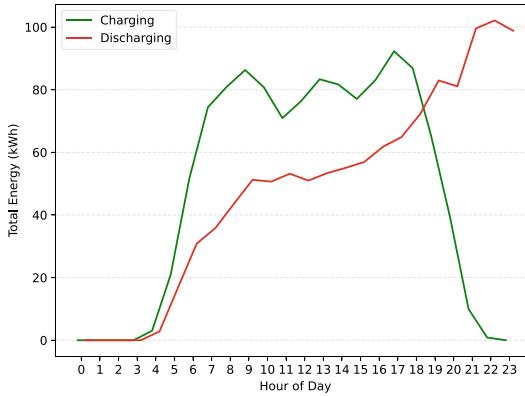


Fig. 4. Charging and discharging patterns over the 122-day period on a 24-h basis (summation over days) for all households in Group 1 under the inter-trading scenario.

Table 4. The effect of battery depreciation cost (IKR) on the total cost in the inter-trading scenario. Negative values indicate net profit.

Group	Participant	0.5	1	2	4
1	Prosumer 1	-160	235	-6460	-4466
1	Prosumer 2	34322	34650	34116	33236
1	Consumer 3	41799	42246	51497	51981
1	Subtotal	75961	77131	79153	80751
2	Prosumer 1	1208	1589	-2283	-1652
2	Consumer 2	11716	12045	17181	17801
2	Subtotal	12924	13634	14904	16149
1+2	Total	88885	90765	94057	96900

4 Conclusions

This study presented a linear optimization model for peer-to-peer energy trading. Using data from Iceland, we simulated three scenarios: (1) peer-to-grid, (2) P2P trading within pre-defined user groups (intra-trading), and (3) P2P trading between different groups (inter-trading). Among the three scenarios, the P2G configuration resulted in the highest total cost, followed by intra-trading. Inter-trading produced the lowest costs, although the absolute differences between scenarios were relatively small. The simulations showed that most cost savings were achieved through intra-group trading, while inter-group trading yielded only marginal additional benefits. We also conducted a sensitivity analysis on battery depreciation costs. The results indicate that when these costs exceed a certain threshold, consumers bear a disproportionately large share of the total expense. This underscores the importance of accurately estimating and allocating battery-related costs in any practical deployment.

Historical data for user demand and solar production were treated as known inputs. Future work will incorporate solar production forecasts from meteorological models, develop predictive models for user demand, for example using deep neural networks, incorporate energy transmission losses or network charges for inter-building energy exchanges. In the present study, the optimization objective is the minimization of the total energy cost of all users. While this approach ensures overall system efficiency, it does not explicitly account for how costs and benefits are distributed among individual participants. This can be addressed in several ways, e.g., by using max-min fairness of [2].

Acknowledgements. The work of the first author was supported by the Northern Periphery and Arctic Programme through the project “COPower based virtual POWER plant” (COPower, contract no. NPA0300090). Additional support was provided by The Reykjavik Energy Research and Innovation Fund (VOR).

Disclosure of Interests. The authors declare that they have no conflicting interests.

References

1. Alam, M.R., St-Hilaire, M., Kunz, T.: Peer-to-peer energy trading among smart homes. *Appl. Energy* **238**, 1434–1443 (2019)
2. Bertsimas, D., Farias, V.F., Trichakis, N.: The price of fairness. *Oper. Res.* **59**(1), 17–31 (2011)
3. Cui, S., Wang, Y.W., Shi, Y., Xiao, J.W.: An efficient peer-to-peer energy-sharing framework for numerous community prosumers. *IEEE Trans. Industr. Inf.* **16**(12), 7402–7412 (2019)
4. Cui, Y., Zhu, J., Meng, F., Zoras, S., McKechnie, J., Chu, J.: Energy assessment and economic sensitivity analysis of a grid-connected photovoltaic system. *Renew. Energy* **150**, 101–115 (2020)
5. Diamond, S., Boyd, S.: CVXPY: a Python-embedded modeling language for convex optimization. *J. Mach. Learn. Res.* **17**(83), 1–5 (2016)
6. Gbadega, P.A., Sun, Y.: Centralized peer-to-peer transactive energy market approach in a prosumer-centric residential smart grid environment. *Energy Rep.* **8**, 105–116 (2022)
7. Kanakdhurga, D., Prabaharan, N.: Demand response-based peer-to-peer energy trading among the prosumers and consumers. *Energy Rep.* **7**, 7825–7834 (2021)
8. Li, T., Dong, M.: Residential energy storage management with bidirectional energy control. *IEEE Trans. Smart Grid* **10**(4), 3596–3611 (2018)
9. Long, C., Wu, J., Zhou, Y., Jenkins, N.: Peer-to-peer energy sharing through a two-stage aggregated battery control in a community microgrid. *Appl. Energy* **226**, 261–276 (2018)
10. Nguyen, S., Peng, W., Sokolowski, P., Alahakoon, D., Yu, X.: Optimizing rooftop photovoltaic distributed generation with battery storage for peer-to-peer energy trading. *Appl. Energy* **228**, 2567–2580 (2018)

11. Shah, M.I.A., Wahid, A., Barrett, E., Mason, K.: Multi-agent systems in peer-to-peer energy trading: a comprehensive survey. *Eng. Appl. Artif. Intell.* **132**, 107847 (2024)
12. Soto, E.A., Bosman, L.B., Wollega, E., Leon-Salas, W.D.: Peer-to-peer energy trading: a review of the literature. *Appl. Energy* **283**, 116268 (2021)
13. Tushar, W., et al.: Peer-to-peer energy systems for connected communities: a review of recent advances and emerging challenges. *Appl. Energy* **282**, 116131 (2021)



Agent-Based Flexibility Aggregation for a Distributed Redispatch

Malin Radtke^{1,2}(✉) , Sanja Stark¹, and Stefanie Holly¹ 

¹ OFFIS - Institute for Information Technology, Escherweg 2, 26121 Oldenburg, Germany

{malin.radtke,sanja.stark,stefanie.holly}@offis.de

² Carl von Ossietzky Universität, Ammerländer Heerstraße, 114-118, 26129 Oldenburg, Germany

Abstract. Integrating the flexibility of small-scale assets into the redispatch process is a necessary measure for the reliable operation of future energy systems with high shares of renewable energy, flexible loads and battery storages, especially at low voltage levels. Market-based procurement of redispatch flexibility from such assets is a useful complement to the existing process, enabling a scalable solution that reduces complexity for grid operators by decentralizing operations and pooling assets through aggregators. We present the novel flexibility model BASE for household-level assets that is tailored to this use case. The model provides fundamental features and principles, but the calculation method may differ for different asset types. This is demonstrated by detailing the calculation methods for battery storages and heat pumps. The model enables efficient local computation on edge hardware, and can be used to generate optimized aggregated redispatch offers in a flexibility market. Intelligent aggregation facilitates offers that leverage the heterogeneity of the asset pool, thereby conferring a substantially greater degree of flexibility compared to that achievable through individual flexibility offers. The conceptual extension to other types of flexible plants is planned for future research.

Keywords: decentralized redispatch · small-scale flexibility · multi-agent system

1 Introduction

The increasing deployment of small-scale assets in the energy system introduces new opportunities and challenges for grid management [8]. These assets, including distributed generation sources such as distributed generation sources such as systems, Heat Pumps (HPs), Battery Energy Storage Systems (BESSs) and Electric Vehicles (EVs), contribute to enhanced system flexibility by adjusting

M. Radtke and S. Holly—These authors contributed equally.

their power generation or consumption dynamically in response to grid conditions, market signals, or operational constraints. However, their integration poses challenges for network operators, who are responsible for maintaining grid stability. Traditionally, redispatch measures have been used to mitigate congestion and ensure stable system operation [5]. Yet, with the growing number of small-scale assets, redispatch processes become increasingly complex, requiring scalable solutions [9,10].

To address these challenges, a partially market-based redispatch is currently discussed for the practical implementation of the so-called Redispatch 3.0 in Germany [1]. An agent-based system design to facilitate the flexibility calculation of small assets on-the-edge and to aggregate individual flexibilities of a pool of assets is proposed by Hess et al. [9]. However, the challenge of effective flexibility calculation and aggregation methods that fully leverage pooling effects and meet market restrictions remains.

In this paper, we present an effective method to model and aggregate power flexibilities of small-scale assets, tailored specifically for application in a market-based redispatch. Our approach allows for effective flexibility aggregation while ensuring compliance with market requirements. Furthermore, we provide a comparative analysis of our method against an alternative aggregation strategy. This comparison illustrates the advantages of our approach, as it demonstrates the ability to leverage the full range of pooled flexibility while aligning with market restrictions.

The paper is organized as follows: First, we introduce the use case under consideration, including the market design and the resulting requirements for flexibility modeling. We then give a brief overview of existing flexibility models and explain why they are not suitable for the specific use case. Next, we present the generic baseline optimization at the household level, the flexibility model called BASE - illustrated by two types of assets - and aggregation and disaggregation. This is followed by an illustrative example of flexibility at the household level and a simulative evaluation at the aggregation level.

2 Setting

Our approach to integrating small-scale power assets into the existing redispatch process in Germany is based on a complementary market-based redispatch mechanism [1]. A simplified illustration of the setting, using an agent-based system design as proposed by Hess et al. [9], is illustrated in Fig. 1.

Transmission System Operators (TSOs) and Distribution System Operators (DSOs) are responsible for maintaining grid stability. To achieve this, they continuously monitor planned power consumption and feed-in schedules (referred to as baselines). If congestion is anticipated based on these schedules, they act as buyers of power flexibility on a market platform. Power flexibility is provided by *Flexible Device Owners*, who voluntarily offer the ability to adjust power consumption or feed-in. The assets under consideration include BESSs, HPs and EV pre-charging storages.

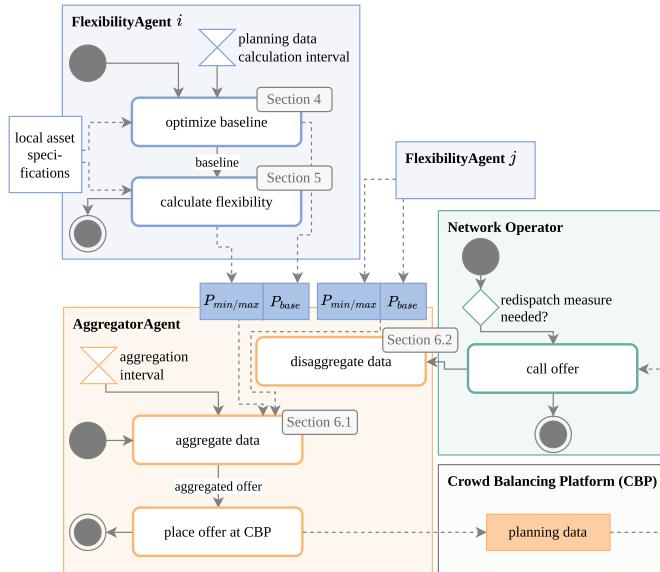


Fig. 1. Simplified activity diagram to illustrate the setting

Each flexible device is represented by a software agent called a *FlexibilityAgent*. This agent is responsible for monitoring and controlling its associated asset(s), computing a cost-optimal baseline P_{base} for those flexible assets, and determining the asset's maximum potential deviations from this baseline, which constitute its flexibility potential $P_{min/max}$ (see Sect. 5).

Planning data – including the baseline schedule of the flexible assets and flexibility offers – are submitted to a Crowd Balancing Platform (CBP), which acts as a marketplace for flexibility trading. The CBP facilitates the registration of flexible assets, submission and processing of planning data, activation of accepted flexibility offers, and the validation of flexibility provision.

Individual *FlexibilityAgents* interact with an *AggregatorAgent*, an intermediary entity that aggregates flexibility offers and baselines from multiple assets behind a single grid connection point before submitting them to the CBP (for details see Sect. 6).

For valid participation in the CBP, the *AggregatorAgent* must comply with specific market rules, which (among others) include Time Requirements (TRs) and Data Format Requirements (DRs):

- **TR1** Flexibility offers must be submitted before the Day-Ahead market. Planning data submission starts on Day D-7 at 11:00 a.m. and closes on Day D-2 at 08:00 a.m.
- **TR2** Flexibility offers start from each 15-minute interval but must cover one-hour intervals.
- **TR3** Calculations for baseline and flexibility must be efficiently performed on edge devices with limited computational resources.

- **DR1** Planning data must include the aggregated baseline schedule and flexibilities for its flexible assets behind a grid connection point.
- **DR2** The planning data must be provided at a 15-minute resolution.
- **DR3** The aggregated baseline and flexibility must be provided as power values in watts (W).
- **DR4** The flexibility values represent relative deviations from the baseline schedule (min/max power values for each interval).

This setting presents significant challenges, particularly in ensuring accurate calculation and aggregation of flexibility in accordance with market requirements.

3 Related Work

The concept of flexibility in energy systems has been explored from various perspectives in scientific literature. Degefa et al. [7] provide a comprehensive overview of the different definitions and characteristics of flexibility across the field, highlighting the diverse interpretations depending on stakeholder perspectives.

Brandt et al. [3] conducted a systematic evaluation of flexibility modeling approaches for energy systems, focusing specifically on models capable of representing flexibility in a unified form across different types of power devices. The authors selected five approaches for detailed analysis: *Flexibility Trinity* [17], *Multienergy Node* [6], *Support Vector Data Description* [4], *OpenTUMFlex* [13,20], and *FlexOffer* [18].

The *OpenTUMFlex* model [13] shares similarities with our approach, defining flexibility as possible deviations from a cost-optimal baseline schedule. However, it does not fully address our requirements, as it does not meet the temporal requirement TR2 (one-hour intervals for flexibility offers). Furthermore, *OpenTUMFlex* only contains a simple aggregation approach.

The *FlexOffer* model [18] serves as a generalized representation of flexibility that can be exchanged between different entities in a cellular energy system [2]. The primary goal is to provide a scalable, device-independent approximation of flexibility that retains most of the flexibility while being computationally efficient for aggregation and optimization. The application of *FlexOffers* to BESSs was explored in [12] and to HPs in [11]. Essentially, *FlexOffers* represent the energy available for consumption at each time slice modeled by a set of constraints, optionally including interdependencies between time slices. Since market offers must be power values in watts (DR3), the lack of explicit power bounds may result in infeasible offers without additional knowledge of the asset's power limits. Due to this need for information on power restrictions (DR3) and the use of the flexibility model to calculate a deviation from the baseline (DR4), *Flexibility Trinity*, *Multienergy Node* and the *Support Vector Data Description* are also not suitable for the use case.

Another related model is *Amplify*, a modeling approach by Tiemann et al. [16]. It was designed specifically for multi-purpose usage of BESSs in a swarm context. It prioritizes primary applications (like peak shaving) while offering

remaining flexibility for secondary applications and focuses on battery-centric flexibility calculation with emphasis on State of Charge (SoC) management. It has both an energy-bound and a power-bound representation. However, the approach lacks an integrated aggregation methodology, which is essential for our market-oriented use case.

Many approaches have also been taken to explicitly model the flexibility of HPs, mostly in combination with thermal storage. Two examples are the approaches of Steinle et al. [15] and You et al. [19]. Both model the flexibility of HPs using operational constraints of the HP considering power and energy aspects, and comfort constraints related to the thermal energy demand of the household. We model HPs in the same way, but use this model as a basis to retrieve flexibility in a generalized asset type independent form.

4 Baseline Optimization

The optimization of the baseline on the household level is the first step in our approach. It is a linear or non-linear optimization problem, depending on the involved assets. Considered assets are non-flexible PV-plants and flexible BESSs and HPs as well as a combination of these. We always aim to find a cost-optimal schedule that fulfills the demands of the household. For power consumption we either assume a fixed price or a dynamic tariff that depends on the day ahead spot market prices ($j_c(i)$). The assumed feed-in remuneration (j_f) is based on the current German tariff for PV systems with partial feed-in. The demand of the household may be electrical and/or thermal energy. The thermal demand $P_{th,d}(i)$ is only considered if a HP is involved, while electricity demand $P_{el,d}(i)$ is always taken into account. We define our signs from the perspective of the grid connection point, i.e. feeding into the grid is positive and drawing energy from the grid is negative. Since we do not want to regulate the PV feed-in, we consider the electrical household demand and the PV feed-in as the fixed power values of the household, i.e. $P_{fix}(i) = P_{el,d}(i) + P_{pv}(i)$. The cost optimization is achieved by minimizing the objective function:

$$C = \sum_{i=1}^{i_{max}} \left[\left(P_{fix}(i) + \sum_{u=1}^{|U|} P_u(i) \right) \cdot \Delta t \cdot \begin{cases} j_c(i) & \text{if } P_{fix}(i) + \sum_{u=1}^{|U|} P_u(i) \leq 0 \\ j_f & \text{else} \end{cases} \right] \quad (1)$$

where $|U|$ is the number of flexible units that contribute to the power generation or consumption of the household, and i_{max} is the maximum number of time intervals considered.

Flexible assets can include any asset that has some type of energy storage, such as BESSs, HPs, or EVs. If a HP is considered, we assume that the HP is operated in combination with a thermal storage tank to allow flexible operation. To further increase the flexibility, a continuous controllability of the HP is required.

The constraints of the optimization problem arise from the technical restrictions of the flexible assets. The electrical power $P_u(i)$ consumed or provided by a flexible asset is restricted by maximal and minimal power values.

The SoC must always be within its allowable limits, i.e. in $[0, 1]$. For simplicity, we assume that the thermal storage energy is equal to the sink temperature and that the storage is always perfectly mixed. The SoC of the thermal storage is therefore 0 at the minimum temperature and 1 at the maximum permissible temperature (as in [15, 19]).

The power value $P_{base}(i)$ in one time step affects the SoC value F_{base} in the next time step:

$$F_{base}(i+1) = F_{base}(i) - \frac{(P_{base}(i) \cdot \eta(i) - d(i)) \cdot \Delta t}{E_{cap}} \quad (2)$$

where $\eta(i)$ is an efficiency parameter. For a BESS it differs for charging and discharging. For HPs, it corresponds to the Coefficient of Performance (COP), which depends on the ambient and sink temperatures and thus varies over time. The COP is calculated using linear modeling as in [15]. The demand $d(i)$ denotes a demand to be satisfied by the considered asset. For BESSs, it is usually 0, since the household's electrical demand can be met by other sources, and the choice is part of the optimization problem. For HPs, however, the heat demand of the household must be met by the HP in any case. E_{cap} denotes the capacity of the storage that is used to normalize the Energy to SoC. For the HP it corresponds to the energy of the thermal energy storage, i.e. the water tank, when it has the maximum allowed temperature.

To prevent premature aging of the system, for BESSs, the number of cycles per day is limited, and for HPs, frequent cycling (i.e., turning on and off) is penalized, while balanced operation is rewarded. In addition, there may already be redispatch obligations due to previous trading in the redispatch market. This results in fixed power values for certain time intervals that must be considered in the baseline optimization.

5 Flexibility Calculation

Optimizing the household baseline results in a schedule for each of the flexible assets that contributes to cost-optimal coverage of household needs. Since it is difficult to predict consumption at the individual household level, we offer the aggregated baseline of the controllable units and their flexibility to deviate from that baseline to the CBP (see Sect. 2).

Flexibility calculations for individual assets must be performed on edge devices (TR3), taking into account technical constraints and local household demand. However, they are not yet required to comply with market restrictions. Instead, they must enable the efficient calculation of aggregated market offers while leveraging the potential of pooling heterogeneous assets. Therefore, we will first define several general principles that allow for the aggregation of flexibility. Following this, we will delve into the specifics of calculating the flexibility of two different asset types.

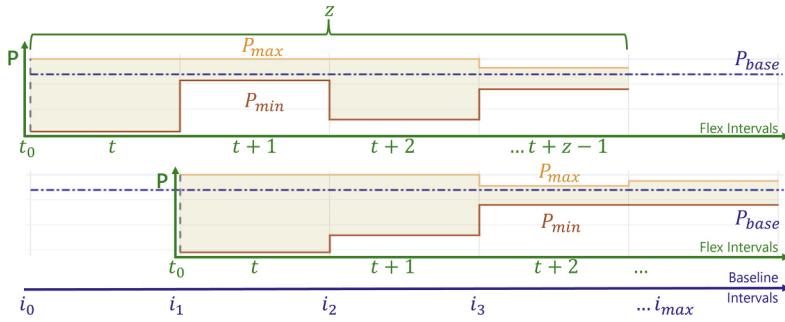


Fig. 2. Overview on time window and variables in flexibility calculation

Basic Principles. Figure 2 illustrates the temporal relationships and basic structure of the flexibility bands.

1. Flexibility spans over z intervals and is calculated as a flexibility band consisting of power values $(P_{min}(t), P_{max}(t))$, with $t \in \{1, \dots, z\}$.
2. The individual flexibility bands contain absolute power values - not deviations from the baseline.
3. Flexibility bands consider technical and comfort constraints.
4. An asset can perform any power value $P(t)$ between the boundaries, i.e. $P_{min}(t) \leq P(t) \leq P_{max}(t)$
5. The calculation of successive time steps is always based on the maximum use of flexibility in the previous time steps.
6. Unused flexibility potential can be transferred to successive time steps
7. The baseline is considered
 - by using the SoC of the baseline (F_{base}) at the beginning of each interval i , i.e. $F(t_0) = F_{base}(i)$
 - by constraining the flexibility bands such that the baseline values are feasible at all time steps in z , and thus the baseline falls within the flexibility bands, i.e. $P_{min}(t) \leq P_{base}(t) \leq P_{max}(t)$

The resulting flexibility model is called **BASE** (**B**ounded **A**ggregatable **S**tateful **E**dge **F**lexibility).

5.1 Battery Storage Flexibility

For BESSs technical and baseline constraints can be represented and combined using SoC values. For the sake of clarity, only discharge flexibility is discussed below. Charge flexibility is calculated in the same way, taking into account the corresponding efficiencies and SoC limit (maximum SoC = 1). Starting from the SoC at an interval start i determined by the baseline $F_{base}(i) = F_{reach}^{min}(0)$,

the technically reachable SoC values are calculated. They can be achieved with maximum discharge power (P_{bat}^{dis}) in each time step:

$$F_{reach}^{min}(t+1) = \max(0, F_{reach}^{min}(t) - P_{bat}^{dis} \cdot \frac{1}{\eta_{dis}} \cdot \frac{\Delta t}{E_{cap}}) \quad (3)$$

where η_{dis} corresponds to the discharging efficiency.

To ensure baseline power discharge requirements P_{base}^{dis} are met in the future, the required SoC bound F_{req}^{min} is calculated backwards:

$$F_{req}^{min}(t-1) = \max(0, F_{req}^{min}(t) + \frac{\Delta t}{E_{cap}} \cdot P_{base}^{dis}(t)) \quad (4)$$

The most restrictive combination of reachable and required SoC results in the allowed SoC range for each interval t :

$$F^{min}(t) = \max(F_{reach}^{min}(t), F_{req}^{min}(t)) \quad (5)$$

Using the adjusted SoC bound $F^{min}(t)$, the power flexibility for each interval t is calculated:

$$P_{max}(t) = \frac{(F^{min}(t-1) - F^{min}(t)) \cdot E_{cap}}{\Delta t} \cdot \eta_{dis} \quad (6)$$

5.2 Heat Pump Flexibility

HP flexibility is largely modeled in the same way as BESS flexibility. One main difference is that the heat demand of the household must already be included in the calculation of the technically feasible values. Thus, P_{max} represents the minimum power required to meet the heat demand and not drop below the minimum temperature of the tank. P_{min} indicates the maximum power that can be consumed without exceeding the maximum temperature of the tank (also taking into account the heat demand). Both values refer to power consumption and are therefore negative. Another difference is that the efficiency, i.e. the COP, depends on both the current sink temperature and the ambient temperature. Since the current state of the storage tank and thus the temperature is not known during aggregation, we calculate the HP flexibility by assuming the most unfavorable COPs for each direction. $COP_{min}(t)$ is the smallest COP and $COP_{max}(t)$ is the largest COP, both for time t and thus the associated ambient temperature. The heat demand of a household $P_{th,d}(t)$ consists of space heating and domestic hot water [19]. We assume that this heating demand, or a forecast of it, is given.

The basic calculation of the flexibility bands is performed sequentially for each time interval. First the COP is calculated, using a linear model as in [15]. With this input the upper and lower bound of the flexibility band can be calculated for the next time step:

$$P_{req}(t+1) = \left(P_{th,d}(t) - \frac{F_{tes}(t) \cdot E_{cap}}{\Delta t} \right) \cdot \frac{1}{COP_{min}(t)} \quad (7)$$

where F_{tes} is the SoC of the thermal energy storage and thus $F_{tes}(t) \cdot E_{cap} \cdot \frac{1}{\Delta t}$ is the thermal power that can be generated by the discharge of that remaining SoC. Using the worst efficiency ($COP_{min}(t)$) results in the maximum electrical power required to achieve the required minimum thermal power.

$$P_{allo}(t+1) = \left(P_{th,d}(t) + \frac{(1 - F_{tes}(t)) \cdot E_{cap}}{\Delta t} \right) \cdot \frac{1}{COP_{max}(t)} \quad (8)$$

where $(1 - F_{tes}(t)) \cdot E_{cap} \frac{1}{\Delta t}$ is the thermal power that can still be absorbed by the storage tank until the maximum temperature ($F_{tes} = 1$) is reached. Applying the highest efficiency ($COP_{max}(t)$), the minimum electrical power that can result in the maximum allowable thermal power is obtained.

After calculating the value for P_{req} or P_{allo} , the SoC for the next time step can be determined (see Eq. 2). In this way, P_{req} and P_{allo} can be calculated successively. In every time interval it must be ensured that the required values are technically feasible, i.e. that they lie within $[0, P_{hp}^{max}]$, and encompass the power values in the baseline:

$$P_{max}(t) = \min(P_{base}(t), P_{req}(t), P_{hp}^{max}) \quad (9)$$

$$P_{min}(t) = \max(P_{base}(t), \min(P_{allo}(t), P_{hp}^{max})) \quad (10)$$

where P_{hp}^{max} is the maximal charging power of the HP. If P_{max} or P_{min} differs from P_{allo} respectively P_{req} , the heat demand in the interval is either over- or under-fulfilled. Therefore, either more or less heat has to be generated in previous intervals. In this case, a backtracking algorithm first calculates the unsatisfied heat demand $P_{th,d}^{us}(t)$ or the excess heat supply $P_{th,s}^{ex}(t)$ (negative) that would result from the required power values $P_{max}(t)$ or $P_{min}(t)$. Again, the most unfavorable efficiency is assumed, i.e. the unsatisfied heat demand is maximized by the worst COP $COP_{min}(t)$ and the excess heat supply is maximized by the best COP $COP_{max}(t)$:

$$P_{th,d}^{us}(t) = P_{th,d}(t) - \left(P_{max}(t) \cdot COP_{min}(t) - \frac{F_{tes}(t) \cdot E_{cap}}{\Delta t} \right) \quad (11)$$

$$P_{th,s}^{ex}(t) = P_{th,d}(t) - \left(P_{min}(t) \cdot COP_{max}(t) - \frac{(1 - F_{tes}(t)) \cdot E_{cap}}{\Delta t} \right) \quad (12)$$

where E_{cap} is the energy capacity of the thermal storage. At time step $t - 1$, all values are recalculated, using $P_{th,d}(t-1) = P_{th,d}(t-1) + P_{th,d}^{us}(t)$ or respectively $P_{th,d}(t-1) = P_{th,d}(t-1) + P_{th,s}^{ex}(t)$ for Eqs. 7 and 8, but the original $P_{th,d}(t-1)$ for the SoC adaptation. If a valid solution is found in $t - 1$, the process continues with forward propagation. Then, the SoC in interval t ensures that a valid solution exists. If problems arise due to the changed heat demand in step $t - 1$, the heat demand is shifted to $t - 2$, and so on. When the backtracking algorithm reaches time step 0, the flexibility power value is set equal to the power value of the baseline, and the recursion is terminated. Thus, the baseline always remains as a fallback for P_{max} and P_{min} , while ensuring that $P_{min}(t) \leq P_{base}(t) \leq P_{max}(t)$ for all t . There is always a feasible solution, presuming a feasible baseline.

6 Optimization of Flexibility Pools

The aggregator (see Sect. 2) optimizes the marketing of the pool of small-scale plants. This includes the aggregation and disaggregation of flexibility offers. Aggregation is an optimization problem, given the individual flexibility of each agent and the goal of finding the best combined offer for CBP. If a flexibility offer is later called by a grid operator, the remaining offers can be recalculated, allowing to make mutually exclusive offers.

Our notion of flexibility comprises two different aspects: the technically feasible flexibility, represented by the flexibility bands, and the virtual flexibility, which results from the non-execution of the baseline. Furthermore, Aggregation is done separately in both directions. Due to space limitations, only the positive direction is described below. Negative aggregation is done analogously. The objective is to maximize the offered power P_{offer} over the product duration of z intervals and all n agents:

$$P_{offer} = \sum_{a=1}^n \sum_{t=1}^z P_{\Delta}^a(t) \quad (13)$$

where $P_{\Delta}^a(t)$ is the power deviation from the baseline of the agent a in the interval t . It is a deviation from the agent's scheduled baseline power $P_{base}^a(t)$. This is subject to a number of constraints. First, the values of the combined offer must be equal over the full product duration of z intervals:

$$\forall t \in \{1, \dots, z-1\} : \sum_{a=1}^n P_{\Delta}^a(t) = \sum_{a=1}^n P_{\Delta}^a(t+1) \quad (14)$$

Second, the power values for each agent must lie within the bounds that are defined by the technical and virtual flexibility of this agent. The lower limit for the positive technical flexibility is zero which would correspond to offer no flexibility and perform the baseline. The upper limit is determined by the technical constraint that the combination of $P_{\Delta}^a(t)$ and the baseline must remain below the maximum power limit of the asset. In addition, this limit may be extended by unused technical flexibility from previous time steps $P_{tra}^a(t)$. The general formula is:

$$P_{\Delta}^a(t) \leq v^a(t) \cdot (P_{max}^a(t) - P_{base}^a(t)) + P_{tra}^a(t) \quad (15)$$

In addition, virtual flexibility usage must be constrained to ensure future baseline feasibility. Thus, if the baseline has an alternate direction after t ($P_{base}^a(k) < 0$), $v^a(t)$ is always zero; otherwise, it is 1.

Thus, the first value chosen for the agent a , i.e. $P_{\Delta}^a(0)$, must be below the maximum power flexibility in this time step $P_{max}^a(0) - P_{base}^a(0)$. The transferable flexibility $P_{tra}^a(t)$ is defined as follows:

$$P_{tra}^a(t) = \sum_{k=1}^{t-1} \begin{cases} P_{max}^a(k) - \min\{0, P_{base}^a(k) + P_{\Delta}^a(k)\}, & P_{max}^a(k) < 0, \\ P_{max}^a(k) - \max\{0, P_{base}^a(k) + P_{\Delta}^a(k)\}, & P_{max}^a(k) \geq 0. \end{cases} \quad (16)$$

where $(P_{base}^a(k) + P_{\Delta}^a(k))$ corresponds to the resulting schedule in interval k . If this value has the same sign as the flexibility limit P_{max}^a , the technical flexibility is reduced by this value, but the remainder can be transferred. If it has the opposite sign, then the virtual flexibility has been used in the interval. Therefore, $P_{max}^a(k)$ can be transferred entirely and added to P_{tra}^a . For BESSs it is always true that $P_{max}^a \geq 0$ and for HPs it is true that $P_{max}^a \leq 0$.

After a flexibility offer has been requested by a grid operator, the aggregator has to allocate the redispatch obligations to the individual assets (see Sect. 2). For this, it uses essentially the same optimization as for aggregation. Additional constraints ensure that the targeted redispatch flexibility is achieved by the asset pool, while the objective aims to minimize costs. However, if the disaggregation problem is infeasible, e.g., due to short-term changes in flexibility or rescheduling after a disturbance, cost minimization is neglected. Instead, the objective changes to minimizing the cumulative deviation from the target instead of using hard constraints.

7 Illustrative Example

To illustrate the flexibility concept, we show the calculated baselines and flexibilities for an exemplary household in Fig. 3. The household has a 6.5 kWh BESS (as in [9]) and a HP modeled according to the specifications in [19]. The two columns show the flexibility for two consecutive time steps. The electrical and thermal household demand is shown in the top row. For each of the two flexible assets, the power flexibility (P_{max} , P_{min}) is plotted along with the SoC curve that would result from executing these power bands (SoC P_{max} , SoC P_{min}). The baseline values are marked by the dashed lines. The figure illustrates that the flexibility power bands quickly lead to the respective SOC limits, while the baseline values lie between them. When the baseline causes the SOC to be close to one of the limits, there is little flexibility left in that direction, and the flexibility power values are close to or identical to the baseline. In addition, flexibility can change significantly from one time step to the next.

8 Comparative Case Study

To evaluate the effectiveness of our approach, we conduct a comparative analysis against an alternative approach that directly aligns with market requirements at the household level. While our proposed method delays market restrictions until the aggregation phase to maximize pooling effects, the alternative approach implements market constraints directly during the initial flexibility calculation at each individual agent at household level.

Our comparative case study is structured as follows: First, we describe the market-restricted approach that serves as our benchmark. Then, we detail our experimental setup, including the assets used, input configurations, and evaluation metrics. Finally, we present and discuss our results, focusing on both effectiveness and scalability aspects.

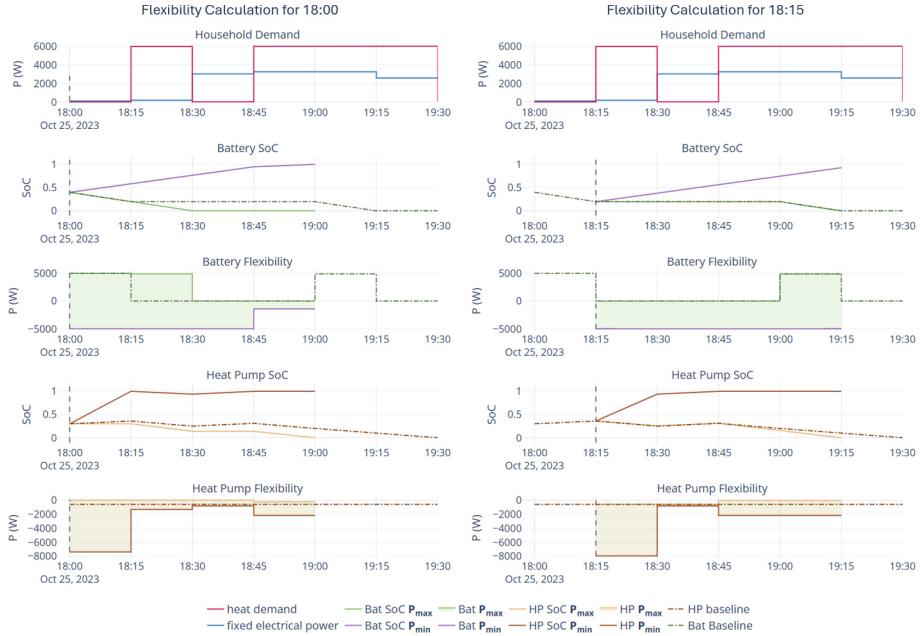


Fig. 3. Example for the baseline and BASE flexibility of a household with BESS and HP in two consecutive time steps

8.1 Market-Restricted Power Flexibility Calculation

In the market-restricted approach, flexibility is calculated to ensure compliance with market requirement TR2, which demands that power values remain constant over the entire one-hour interval. This simplifies the aggregation process for the market platform, as individual flexibility offers already meet format requirements. However, this approach may not fully capture the true flexibility potential of the assets.

For simplification, we again only consider discharge (maximum) power flexibility \hat{P}_{max} . For this, we first determine the energy requirement $E_{base}^{dis}(i)$ of the baseline power schedule within the interval i :

$$E_{base}^{dis}(i) = \sum_{j=1}^{z-1} \max(P_{base}(i+j), 0) \cdot \Delta t \cdot \frac{1}{\eta_{dis}} \quad (17)$$

With this, we can then calculate the amount of energy still available in the storage:

$$E_{avail}^{dis} = \max(0, F_{base}(i) \cdot E_{cap} - E_{base}^{dis}(i)) \quad (18)$$

The available power is then given as:

$$P_{avail}^{dis} = \frac{E_{avail}^{dis}}{\Delta t \cdot z} \cdot \eta_{dis} \quad (19)$$

The discharge flexibility is then calculated as:

$$\hat{P}_{max}(i) = \min(P_{bat}^{dis} - P_{base}^{max}, P_{avail}^{dis}) \quad (20)$$

with $P_{base}^{max} = \max_{t \in \{1, \dots, z\}} P_{base}(i + t)$ being the maximum discharge power in the baseline values during interval i . This approach ensures that the offered flexibility is technically feasible and maintains a constant power value over the required market interval, simplifying integration with the market platform but potentially limiting the pooling effects that could be achieved through more sophisticated aggregation.

8.2 Scenario Set-Up

The objective of our comparative analysis is to evaluate our aggregated flexibility offers with BASE according to two key requirements [9]: scalability (evaluating whether aggregation remains functional with increasing system sizes) and effectiveness (determining how effectively individual flexibilities are used in the pool by comparing total offered flexibility between approaches). Both feasibility and market-compatibility are met by design in both approaches.

For our study, we focus exclusively on BESSs to simplify the interpretation of the results. We use three different sizes of BESSs similar to the case study in [9]. In the first phase of our analysis, we comprehensively tested all scenario configurations with smaller asset pools (10 and 100 assets) to assess effectiveness across diverse real-world conditions:

1. Pricing scenarios
 - *Constant*: Constant prices throughout the planning horizon
 - *Incentive*: Non-constant prices with incentives for charging or discharging at specific times
2. Fixed power patterns (representing household load and PV generation)
 - *Equal cons/feed-in*: Consumption equal to feed-in (no influence of the household)
 - *One Peak*: Higher consumption in one specific interval
 - *Higher cons/feed-in*: Higher consumption or feed-in (alternating)
 - *Random Peak*: Higher consumption in a random interval
 - *Random*: Random power values
3. Battery state configurations
 - *Empty*: All BESSs are initially empty (SoC = 0)
 - *Full*: All BESSs are initially full (SoC = 1)
 - *Random*: Initial SoC values are equally distributed between 0 and 1
 - *Normal*: Initial SoC values are normally distributed with a mean of 0.5 and a standard deviation of 0.25

In the second phase, we focused on scalability by increasing the asset pool to 1,000 and 10,000 assets using a representative subset of configurations.



Fig. 4. Comparison of total flexibility range between the (BASE flexibility) optimized and market-restricted approaches across different pricing scenarios (left), battery SoC configurations (center), and power patterns (right) with 10 and 100 assets.

The simulation covered a three-day period with 15-minute intervals (288 time steps total). We conducted our experiments on a standard server with 4 CPU cores and 8 GB of memory. For solving the optimization problems, we utilized the open-source CBC solver. Our agent system was built using the *mango* framework [14]¹.

8.3 Results

In order to analyze the effectiveness across various scenarios, Fig. 4 illustrates the total flexibility range (difference between maximum and minimum flexibility $P_{max} - P_{min} / \hat{P}_{max} - \hat{P}_{min}$) achieved by both approaches. The figure presents a three-panel comparison highlighting the flexibility range across different conditions. In the first panel, we observe that the optimized approach consistently outperforms the market-restricted approach, with particularly significant improvements under incentive pricing conditions compared to constant pricing scenarios. The center panel demonstrates how initial battery state configurations impact flexibility ranges. The optimized approach shows better performance across all SoC configurations, with the most pronounced advantage occurring in the “Normal” distribution scenario, where batteries start with normally distributed SoC values. This suggests that our approach effectively takes advantage of heterogeneous battery states to maximize overall flexibility. The right panel analyzes performance across different power consumption and generation patterns. Here again, the optimized approach consistently delivers greater flexibility ranges, with particularly notable improvements in scenarios with higher variability (“Random Peak” and “Random” patterns). This indicates that our aggregation method is especially effective when managing diverse and dynamic power profiles, precisely the conditions expected in real-world implementations.

The flexibility gains are most pronounced in scenarios with diverse initial SoC distributions and during periods with variable price signals, where the pooling effect of coordinated aggregation can effectively leverage complementary battery

¹ Source code: <https://github.com/OFFIS-DAI/deer/releases/tag/v1.0.0>.

Table 1. Average Calculation Times for different steps (in seconds). For the baseline optimization and flexibility calculation, calculation durations are given for all 288 intervals. The calculation duration of the aggregation is given for each interval.

Number of Assets	Baseline	Flexibility	Aggregation
10	0.203	0.004	0.075
100	0.214	0.005	0.563
1000	0.322	0.005	5.654
10000	0.320	0.006	78.009

states. This demonstrates that our approach is particularly valuable for real-world implementations where asset heterogeneity is the norm rather than the exception.

For computational performance, Table 1 presents the execution times for different components of the system across varying pool sizes. The results demonstrate a clear architectural advantage: distributed components (baseline and flexibility calculations) show remarkable scalability with nearly constant computation times regardless of pool size. This is achieved because each asset performs these calculations independently using its own computational resources. In contrast, the centralized aggregation component exhibits approximately linear scaling with the number of assets, increasing roughly 10x with each order of magnitude in pool size. However, an aggregation of 10,000 assets should be considered a theoretical observation rather than a practical recommendation. In real-world implementations, an aggregator would likely organize assets into hierarchical sub-pools of manageable size, which would significantly reduce computation times while preserving most of the flexibility benefits. Additionally, commercial aggregation platforms would utilize more powerful computational resources than our experimental setup, potentially enabling direct aggregation of larger asset pools when necessary.

9 Conclusion and Outlook

Integrating small-scale plants into redispatch is an important step toward creating resilient and flexible future energy systems. We have developed the efficient, edge-device-based flexibility modeling approach BASE that can account for and encapsulate local constraints and preferences. At the same time, it is suitable for effectively creating offers for a redispatch market at an aggregated level, thereby leveraging the advantages of a heterogeneous plant pool. We illustrated the calculation method for two types of assets: BESS and HP. Transferring the method to other types of assets with energy storage is part of our future work. However, we will soon demonstrate the methodology in a field test with real assets, including pre-charging storage systems for charging stations, which have their own unique characteristics and specifications. We will also investigate the suitability of the model for other flexibility applications, such as other energy

markets or ancillary services. This will include modified aggregation mechanisms and targeted improvements in aggregation efficiency.

Acknowledgments. This manuscript has been created as part of a project funded by the Federal Ministry for Economic Affairs and Climate Action as part of the “Edge Data Economy” technology program. We gratefully acknowledge our “DEER” project partners’ support in this research.

References

1. Blumberg, G., Schneller, C., Schuster, H., Ocker, F., Ried, S., Stenglein, J.: Redispatch 3.0: Regulatorischer Rahmen, Markt- und Produktdesign. Tech. rep., TenneT TSO GmbH, TransnetBW GmbH (2022). https://www.transnetbw.de/_Resources/Persistent/c/4/6/9/c469f1b0ef6bae7e7bf7260b0b22bdcb29d83db0/221013_Bericht-Redispatch3.0_final1.pdf
2. Boehm, M., et al.: Data management in the mirabel smart grid system. In: Proceedings of the 2012 Joint EDBT/ICDT Workshops, Berlin, Germany, March 30, 2012, pp. 95–102. Association for Computing Machinery (ACM), United States (2012). <https://doi.org/10.1145/2320765.2320797>, eDBT/ICDT 2012 Joint Workshops ; Conference date: 30-03-2012 Through 30-03-2012
3. Brandt, J., et al.: Choosing the right model for unified flexibility modeling. Energy Inform. **5**(1), 10 (2022). <https://doi.org/10.1186/s42162-022-00192-w>
4. Bremer, J., Sonnenschein, M.: Constraint-handling with support vector decoders. In: Filipe, J., Fred, A. (eds.) Agents and Artificial Intelligence, pp. 228–244. Springer, Berlin, Heidelberg (2014), isbn 978-3-662-44440-5
5. Bundesnetzagentur: Redispatch (2024). www.bundesnetzagentur.de/DE/Fachthemen/ElektrizitaetundGas/Versorgungssicherheit/Netzengpassmanagement/Engpassmanagement/Redispatch/start.html. Accessed 12 Dec 2024
6. Chicco, G., Riaz, S., Mazza, A., Mancarella, P.: Flexibility from distributed multienergy systems. Proc. IEEE **108**(9), 1496–1517 (2020). <https://doi.org/10.1109/JPROC.2020.2986378>
7. Degfa, M.Z., Sperstad, I.B., Sæle, H.: Comprehensive classifications and characterizations of power system flexibility resources. Electric Power Syst. Res. **194**, 107022 (2021). <https://doi.org/10.1016/j.epsr.2021.107022>, <https://www.sciencedirect.com/science/article/pii/S037877962100002X>
8. Farhangi, H.: The path of the smart grid. IEEE Power Energ. Mag. **8**(1), 18–28 (2010). <https://doi.org/10.1109/MPE.2009.934876>
9. Heess, P., et al.: A multi-agent approach with verifiable and datasovereign information flows for decentralizing redispatch in distributed energy systems. Energy Inform. **8**(1), 24 (2025). <https://doi.org/10.1186/s42162-024-00464-7>
10. Krueger, C., Otte, M., Holly, S., Rathjen, S., Wellssow, A., Lehnhoff, S.: Redispatch 3.0 – congestion management for German power grids – considering controllable resources in low-voltage grids, pp. 1–7 (2023)
11. Lilliu, F., Pedersen, T.B., Siksny, L.: Heat flexoffers: a device-independent and scalable representation of electricity-heat flexibility. In: Proceedings of the 14th ACM International Conference on Future Energy Systems, pp. 374–385. e-Energy '23, Association for Computing Machinery, New York, NY, USA (2023). <https://doi.org/10.1145/3575813.3597347>

12. Lilliu, F., Pedersen, T.B., Šikšnys, L.: Capturing battery flexibility in a general and scalable way using the flexoffer model. In: 2021 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm), pp. 64–70 (2021). <https://doi.org/10.1109/SmartGridComm51999.2021.9631999>
13. Nalini, B.K., et al.: OpenTUMFlex: a flexibility quantification and pricing mechanism for prosumer participation in local flexibility markets. *Int. J. Electrical Power Energy Syst.* **143**, 108382–108382 (2022). <https://doi.org/10.1016/j.ijepes.2022.108382>, mAG ID: 4290078217 S2ID: d0bc9ec3016c1674125a76b951946e5258f8de38
14. Schrage, R., Sager, J., Hörding, J.P., Holly, S.: mango: a modular python-based agent simulation framework. *SoftwareX*, 101791 (2024). <https://doi.org/10.1016/j.softx.2024.101791>
15. Steinle, S., Zimmerlin, M., Mueller, F., Held, L., Suriyah, M.R., Leibfried, T.: Time-dependent flexibility potential of heat pump systems for smart energy system operation. *Energies* **13**(4), 903 (2020)
16. Tiemann, P.H., et al.: Operational flexibility for multi-purpose usage of pooled battery storage systems. *Energy Inform.* (2022). <https://doi.org/10.1186/s42162-022-00209-4>
17. Ulbig, A., Ulbig, A., Andersson, G., Andersson, G.: On operational flexibility in power systems. IEEE Power Energy Soc. General Meeting (2012). <https://doi.org/10.1109/pesgm.2012.6344676>
18. Šikšnys, L., Pedersen, T.B., Aftab, M., Neupane, B.: Flexibility modeling, management, and trading in bottom-up cellular energy systems. In: Proceedings of the Tenth ACM International Conference on Future Energy Systems, pp. 170–180. e-Energy '19, Association for Computing Machinery, New York, NY, USA (2019). <https://doi.org/10.1145/3307772.3328296>
19. You, Z., Zade, M., Kumaran Nalini, B., Tzscheutschler, P.: Flexibility estimation of residential heat pumps under heat demand uncertainty. *Energies* **14**(18) (2021). <https://doi.org/10.3390/en14185709>, <https://www.mdpi.com/1996-1073/14/18/5709>
20. Zade, M., You, Z., Kumaran Nalini, B., Tzscheutschler, P., Wagner, U.: Quantifying the flexibility of electric vehicles in Germany and California—a case study. *Energies* **13**(21) (2020). <https://doi.org/10.3390/en13215617>, <https://www.mdpi.com/1996-1073/13/21/5617>



A Visualization Framework for Exploring Multi-agent-Based Simulations: Case Study of an Electric Vehicle Home Charging Ecosystem

Kristoffer Christensen^(✉) , Bo Nørregaard Jørgensen , and Zheng Grace Ma^(✉) 

SDU Center for Energy Informatics, Maersk Mc-Kinney Moeller Institute, The Faculty of Engineering, University of Southern Denmark, Odense, Denmark

`{kric, bnj, zma}@mimi.sdu.dk`

Abstract. Multi-agent-based simulations (MABS) of electric vehicle (EV) home charging ecosystems generate large, complex, and stochastic time-series datasets that capture interactions between households, grid infrastructure, and energy markets. These interactions can lead to unexpected system-level events, such as transformer overloads or consumer dissatisfaction, that are difficult to detect and explain through static post-processing. This paper presents a modular, Python-based dashboard framework—built using Dash by Plotly—that enables efficient, multi-level exploration and root-cause analysis of emergent behavior in MABS outputs. The system features three coordinated views (System Overview, System Analysis, and Consumer Analysis), each offering high-resolution visualizations such as time-series plots, spatial heatmaps, and agent-specific drill-down tools. A case study simulating full EV adoption with smart charging in a Danish residential network demonstrates how the dashboard supports rapid identification and contextual explanation of anomalies, including clustered transformer overloads and time-dependent charging failures. The framework facilitates actionable insight generation for researchers and distribution system operators, and its architecture is adaptable to other distributed energy resources and complex energy systems.

Keywords: Multi-agent simulation · data visualization · electric vehicles · dashboard · smart grid · emergent behavior

1 Introduction

Agent-Based Modeling (ABM) is a computational approach that simulates the actions and interactions of autonomous agents to analyze complex systems and emergent phenomena [1]. By representing individual entities—such as households, vehicles, or market participants—with distinct behaviors and decision rules, ABM enables the study of emergent phenomena that arise from local interactions. In practice, these models are often implemented as Multi-Agent-Based Simulations (MABS), where large numbers of agents operate within dynamic environments over time. MABS has become a widely used methodology for modeling complex socio-technical systems, generating rich, high-resolution datasets that capture detailed temporal and spatial patterns across thousands

of agents. While this granularity is essential for understanding emergent phenomena, it also introduces significant challenges: raw simulation outputs are typically stored in flat-file formats (e.g., CSV), resulting in data that is both voluminous and difficult to analyze without dedicated post-processing tools [2, 3]. Effectively interpreting these outputs requires advanced data processing pipelines and interactive visualization frameworks that allow users to explore patterns, trace anomalies, and compare scenarios across multiple analytical levels.

Within the energy domain, MABS are increasingly applied to simulate residential electricity networks—particularly in scenarios involving the widespread adoption of distributed energy resources such as electric vehicles (EVs), heat pumps, and solar photovoltaics. In these contexts, agent behaviors (e.g., charging, heating, or consumption) are governed by heterogeneous rules and stochastic inputs, leading to complex and often unexpected system-level outcomes [4, 5]. A particularly pressing challenge arises in understanding the impacts of EV home charging on low-voltage distribution grids, where synchronized charging behavior can overload critical infrastructure components such as 10/0.4 kV transformers [6]. Identifying the causes of such overload events requires not only aggregate system analysis but also the ability to drill down to individual agent behavior—examining charging schedules, departure times, battery states, and pricing signals in a temporally coherent manner.

In response, this paper introduces a modular, Python-based dashboard framework for the interactive exploration of MABS outputs, with a specific focus on EV charging simulations in residential distribution networks. The system addresses the analytical gap between large, high-resolution simulation datasets and the need for interpretable, decision-support tools by providing:

1. Automated transformation of raw outputs into optimized formats;
2. a multi-page dashboard with interactive visualizations (e.g., time-series, heatmaps, spatial charts); and
3. Coordinated drill-down capabilities for agent-level analysis and anomaly tracing.

The contributions of this work are threefold: (i) a generalizable dashboard architecture tailored to post-simulation analysis of MABS; (ii) a demonstration of its utility in diagnosing emergent behaviors in a full-year, high-frequency EV charging simulation; and (iii) methodological guidance for researchers applying visualization to complex simulation output data. As the use of MABS continues to expand in smart grid research, this work addresses a critical bottleneck: the lack of accessible, scalable, and user-friendly tools for extracting actionable insights from rich agent-level simulations.

2 Related Work

MABS are increasingly used to study complex socio-technical systems, including energy transitions, EV charging ecosystems, and smart grid interactions. These simulations generate high-frequency, multivariate datasets that are often difficult to analyze using static post-processing alone. Recent literature has emphasized the importance of advanced visualization and interactive tools to support analysis, communication, and decision-making in such contexts [2].

2.1 Visualization in Energy and Mobility ABMs

Visualization has long played a central role in helping researchers and stakeholders interpret ABM outputs. In energy and mobility domains, geospatial, temporal, and network-based displays are commonly used to illustrate emergent patterns at multiple scales. For example, in energy adoption studies, spatial heatmaps are used to show geographic clustering of distributed energy resources such as solar panels or EV chargers [3]. In mobility models, vehicle flows or traveler trajectories are often mapped in 2D or 3D to highlight congestion zones or behavioral trends [2, 7]. These approaches help link individual agent behaviors to system-level phenomena but often rely on preconfigured, static visualizations.

Several tools—commercial (e.g., AnyLogic) and open-source (e.g., GAMA, Repast)—support embedded GIS integration and agent mapping [8], enabling visual overlays of model outcomes on real-world geography. However, these tools are often tied to their simulation runtime environments, limiting post-hoc, drill-down exploration flexibility. More recent efforts have addressed this by developing specialized analysis and visualization frameworks that separate simulation and evaluation layers. For instance, AgentLens introduces a visual analytics system for exploring Large Language Model-based agent behaviors through temporal decomposition and traceable cause–effect mapping [9]; Similarly, AgentPy [10] offers interactive exploration of agent behaviors in Python environments using dynamic charts and live controls. These systems reflect a growing trend toward flexible, post-simulation visualization tools, as also pursued in this work.

2.2 Temporal and Interactive Visualization Approaches

Given the time-evolving nature of MABS outputs, time-series plots and animations are widely used to trace system dynamics such as energy demand, charging load, or emissions profiles [11]. Animated replays can highlight spatio-temporal behavior—e.g., EV charging peaks during nighttime hours or dissatisfaction events triggered by scheduling conflicts. However, raw animations often lack analytical precision, and static plots can obscure agent-level variability.

Recent work advocates for more interactive visualization frameworks capable of linking system-level metrics to individual agent actions. For instance, Grignard et al. [12] propose “Agent-Based Visualization”, where visual elements are treated as agents in their own right, enabling immersive inspection in 3D environments. Similarly, the GAMA and Mesa platforms offer browser-based interfaces that support real-time charts, parameter manipulation, and filtering of agents [11].

Despite these advances, the literature reveals persistent gaps in post-simulation analysis tools for MABS, especially in energy domains with large agent populations and high temporal resolution. Many studies report only aggregate system Key Performance Indicators (KPIs) (e.g., total load, average emissions) with limited ability to diagnose specific emergent events, such as transformer overloads or charging failures. Moreover, few tools support drill-down analysis that traces a system-level anomaly to its underlying causes—such as the behavior of specific EVs, their state-of-charge, trip history, or charging strategy.

2.3 Need for Scalable, Multi-level Exploration Tools

This paper addresses this gap by introducing a dashboard framework tailored for agent-level post-simulation analysis of EV-grid interactions. Unlike embedded visualization modules or static reporting tools, our approach supports synchronized exploration across system, transformer, and consumer layers. It enables researchers and distribution system operators (DSOs) to explain emergent, high-impact events by interrogating time-aligned data streams at multiple levels of granularity. This type of analysis—currently under-represented in the literature—is essential for interpreting complex charging ecosystems and guiding future energy policy and infrastructure design.

Despite these advances, a gap remains in tools that support post-simulation exploration of high-frequency, stochastic energy data at multiple analytical levels. Many studies emphasize aggregate KPIs but lack support for tracing emergent phenomena to specific agent behaviors—e.g., transformer overloads caused by clustered EV charging. This paper addresses this gap through a modular dashboard tailored for synchronized, agent-level investigation.

3 Case Study: EV Home Charging

The dashboard framework is applied to a MABS of an EV home charging ecosystem in the Danish residential network of Strib, comprising 126 consumer nodes. The underlying model and simulation environment are based on the laxity-based aggregation strategy proposed in [6]. For this demonstration, the scenario assumes 100% EV adoption with smart charging behavior, where vehicles attempt to charge during periods of lowest electricity price. The simulation spans a full calendar year (2025) with a temporal resolution of one minute.

In such high-resolution, agent-driven simulations, emergent behaviors—such as transformer overloads, suboptimal charging, or consumer dissatisfaction—are common and often non-trivial to diagnose through static or aggregate analyses alone. Understanding the root causes of these phenomena requires layered insight across multiple dimensions of agent behavior. For example, explaining a specific overload event involves identifying which EVs were actively charging at the time, and further drilling down into each consumer's behavior, including their driving patterns (i.e., distance driven before and after the event), charging schedules, EV model specifications (battery capacity and charging power), state-of-charge (SoC) upon arrival, baseload consumption, and respective arrival and departure times. The dashboard enables precisely this kind of structured, drill-down investigation—turning what would otherwise be an opaque anomaly into an interpretable, data-driven narrative. By making such analysis intuitive and reproducible, the tool bridges the gap between raw simulation output and actionable insight.

The MABS model employed here is adopted from [6], which presents a laxity-based EV aggregation strategy aimed at minimizing charging costs while avoiding transformer overload. The model simulates individual households with varying baseload consumption, EV models, and trip profiles, and it incorporates a smart charging mechanism that shifts demand to low-price periods. The outputs from this simulation—covering variables such as charging load, SoC, arrival and departure times, and grid stress indicators—serve

as the primary input to the dashboard framework described in this work. Readers are referred to [6] for detailed modeling assumptions and algorithmic descriptions.

4 Dashboard System Architecture

The presented dashboard framework is implemented using Dash by Plotly, a Python-based web application framework specifically designed for creating interactive, analytical dashboards. This choice was driven by several key considerations:

Tight Integration with Python and Plotly. Dash is natively built for Python, making it highly compatible with the broader Python ecosystem used for data processing and scientific computing (e.g., Pandas, NumPy, and PyArrow). Most critically, it leverages Plotly.js under the hood—a robust, open-source JavaScript graphing library that supports high-quality, publication-ready interactive visualizations. This ensures that the dashboard can seamlessly generate time-series plots, heatmaps, and geospatial charts with full interactivity, pan/zoom, and export support directly from Python without needing separate front-end code.

Web-Native Architecture for Flexible Deployment. Dash applications are rendered as HTML/CSS/JavaScript in the browser, meaning they are web-native by design. This architecture facilitates platform-independent deployment on local machines, intranets, or cloud services (e.g., Heroku, AWS, or internal servers). Because the entire interface is defined in Python, it simplifies development for teams that may not have front-end developers, while also enabling easy adaptation for broader access via browser-based tools.

Modularity and Customization. Dash supports a modular, callback-driven design, which is essential for building coordinated multi-view layouts (e.g., linking maps with time-series plots or enabling drill-down analysis from KPIs to agent-level charts). Furthermore, it supports plug-ins like Dash Bootstrap Components for more advanced layouts and styling, aligning with the needs of a professional-grade visualization tool.

Selecting the appropriate framework for building an interactive analytical dashboard is a critical design decision, particularly in simulation-heavy domains such as energy systems and smart grids. The dashboard must support not only visually compelling plots, but also the ability to handle large volumes of high-resolution, agent-level data and provide flexible, responsive interactivity for exploration and insight generation. To contextualize the selection of Dash by Plotly, this section compares it against several other commonly used dashboard development tools—Power BI, Tableau, Streamlit, and Shiny—across a range of technical and practical dimensions. These include language environment, visualization capabilities, interactivity, scalability, deployment options, and fit for scientific simulation use cases.

A comparison overview can be found in Table 1 in the Appendix. The comparison illustrates that while tools like Power BI and Tableau excel in business intelligence applications, they are generally less suited to the demands of scientific simulations and domain-specific analytics. These platforms offer limited customization, are tightly coupled to proprietary ecosystems, and are optimized for tabular data rather than large-scale, time-resolved agent-based simulations.

Streamlit and Shiny offer more flexibility for technical users and are well-suited for prototyping and lightweight applications. However, their simplicity can be a limitation when building multi-page, highly coordinated dashboards with complex user interactions and visual hierarchies.

In contrast, Dash by Plotly provides the best alignment with the project requirements: it is Python-native, integrates directly with the scientific computing stack (e.g., Pandas, Plotly, NumPy), supports scalable deployment through web-native rendering (HTML/CSS/JavaScript), and offers extensive customization through callback-driven logic and modular components. Its interoperability with Plotly ensures high-quality, publication-ready visualizations that can be exported in SVG format for reporting or academic use. These features collectively make Dash an ideal choice for simulation practitioners who need fine-grained control over data exploration, agent-level insight, and visual storytelling within the context of energy systems modeling.

The dashboard's architecture follows visualization design principles proposed by Munzner [13], emphasizing domain-driven task abstraction and coordinated visual idioms aligned with multi-agent simulation data. To support this structure and aid understanding of the dashboard's structural organization, Fig. 1 presents a high-level architecture diagram of the system. It illustrates the transformation flow from raw simulation outputs through data processing components to the coordinated visualization views, highlighting the modular design and interaction between back-end scripts and front-end dashboard elements.

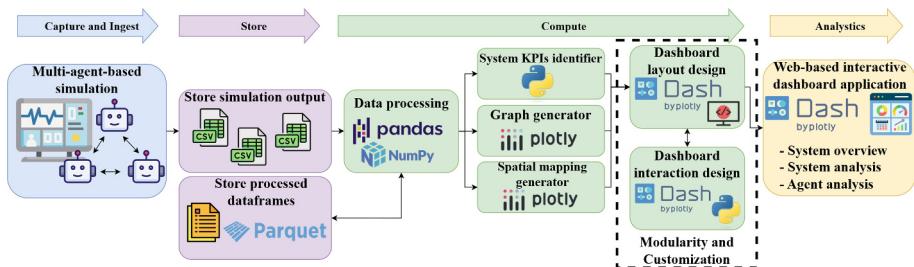


Fig. 1. System architecture of the dashboard framework.

4.1 Data Processing Pipeline

Simulation outputs are generated as CSV files, containing, among others, agent-level time series data for baseload consumption, EV consumption, EV battery State-of-Chart (SoC), CO₂ emissions, electricity prices, and network tariffs. The data pipeline consists of multiple Python modules imported into a central app.py file:

- **data_processing.py** reading data files, formatting time columns, calculating system metrics used as KPIs in the dashboard, and includes a suite of utility functions for preparing specific visualizations.

- **dashboard_layout.py** contains all Dash layout definitions, organizing the interface into modular, responsive sections. The layout utilizes the Dash Bootstrap Components to more easily build consistently styled apps with complex, responsive layouts.
- **graph_generator.py** includes all functions for generating line charts, heatmaps, and bar plots.
- **map_creator.py** defines and updates spatial visualizations using `plotly.express.scatter_map` library, enabling location-based filtering and drill-down functionality.

Dataframes that are created from scratch from the simulation data, such as departure and arrival times, are preprocessed into Parquet format using Pandas to reduce memory footprint and accelerate load times, especially for repeated visual queries.

CSV was chosen as the simulation output format due to its simplicity and compatibility with existing tools. To mitigate memory and performance limitations, large datasets are converted to columnar Parquet files using Pandas, enabling faster loading and selective querying. While no in-memory caching is used in the current version, future enhancements will explore query engines like DuckDB for on-disk filtering.

4.2 Dashboard Design

The dashboard interaction model reflects Shneiderman’s information-seeking mantra [14], supporting ‘overview first, zoom and filter, then details-on-demand’ through the three-layer interface. Therefore, the front end, built using Dash by Plotly, comprises three presentation views, as described in the following.

1) System Overview

KPI Summary Cards: The top of the dashboard page in Fig. 2, displays global metrics such as transformer overload duration, first overload date, load factor, coincidence factor, total dissatisfaction events, average charging cost (DKK/kWh), average CO₂ emissions (kg/kWh), and DSO tariff revenue. Furthermore, the KPIs show the percentage difference for a reference scenario for easy comparison. The KPIs include a percentage difference relative to a baseline scenario when provided. Users can load a reference scenario via a file input, enabling dynamic comparison across key performance indicators. Differences are displayed as percentage changes while relevant graphs include an overlay of reference scenario, streamlining comparative scenario analysis.

Spatial Map: Shows agent locations color-coded by selected parameter (e.g., total electricity expenses, EV charging load, dissatisfaction count) as shown in Fig. 3. Clicking a point triggers a table with detailed information about the user. The map shows the statistics for the current selection with the sum, maximum, mean, and minimum for the system.

Transformer Load Chart: Aggregated load time series with capacity threshold indicated. The investigation of the system concerns the loading of the transformer and is therefore chosen to be located on the System Overview (bottom right in Fig. 2).

Load and Charging Information Chart: A chart (Fig. 4) showing in one graph the correlation between aggregated baseload, charging load, total load, electricity prices,

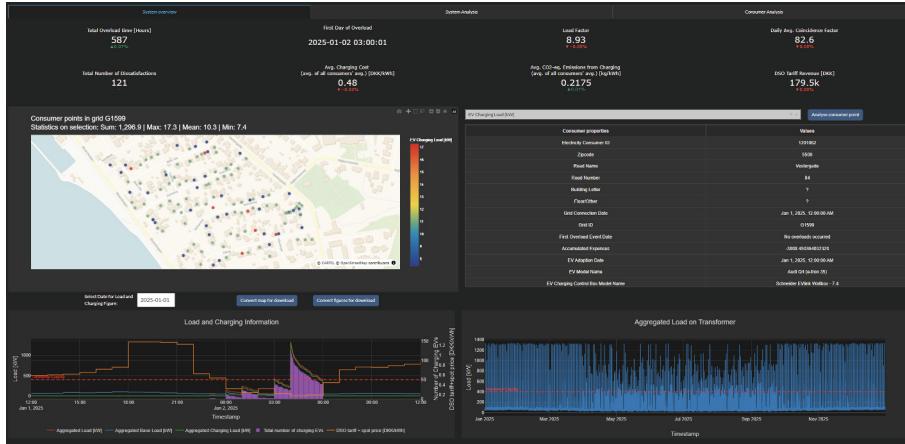


Fig. 2. The dashboard's system overview.

Consumer points in grid G1599

Statistics on selection: Sum: 1,296.9 | Max: 17.3 | Mean: 10.3 | Min: 7.4

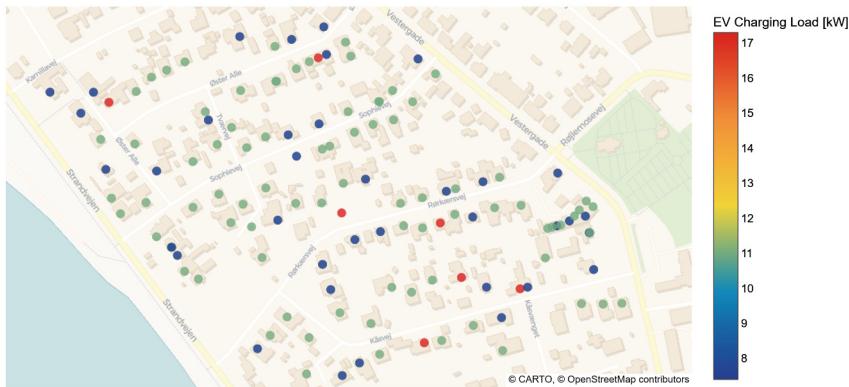


Fig. 3. Spatial map of the simulated system.

the number of charging EVs, and the transformer capacity. This chart has a feature implemented enabling clicking on the bars for the number of charging EVs, which will show on the map the consumers charging at that specific point in time. This allows for easy access to consumer information at specific times in case some behavior requires deeper insights.

2) System Analysis

Time-Series Views: Number of charging EVs; baseload; electricity price breakdown (spot, DSO tariff, total); CO₂ emissions as seen at the top of Fig. 5.

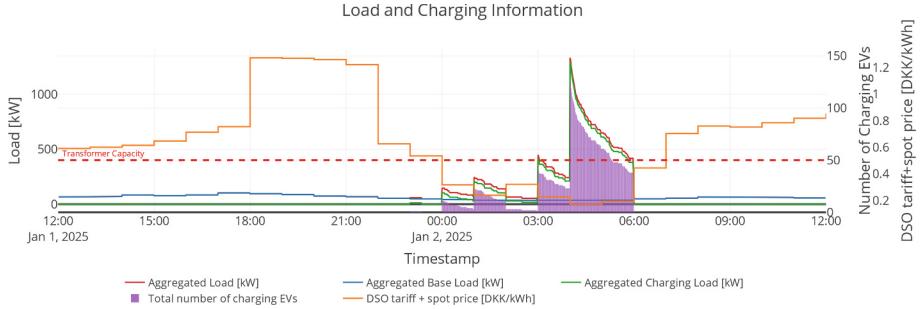


Fig. 4. Aggregated system load, EV charging load, baseload, electricity price, and number of charging vehicles over time. Clicking a bar displays consumer-level data on the spatial map.

Daily Heatmap: EV charging load per user for a selected date, as shown in Fig. 6. The heatmap shows the charging load (color code) at time (x-axis) for all 126 users in the system (y-axis).

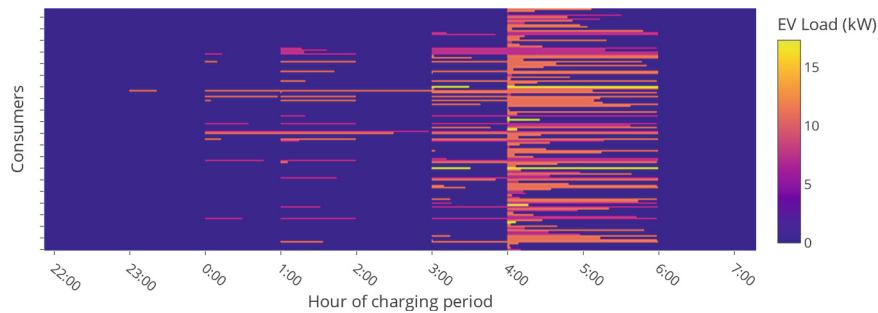
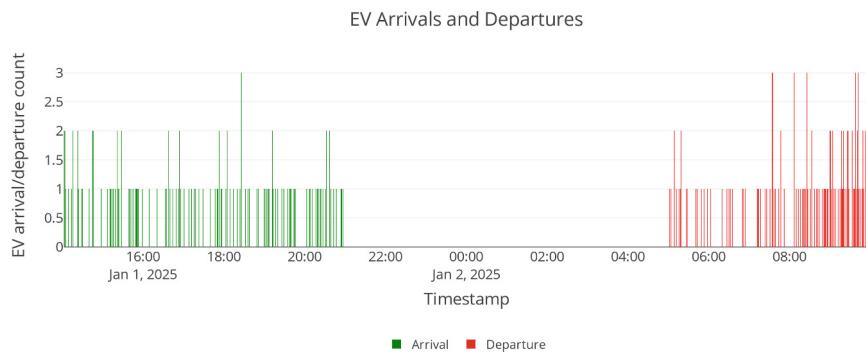
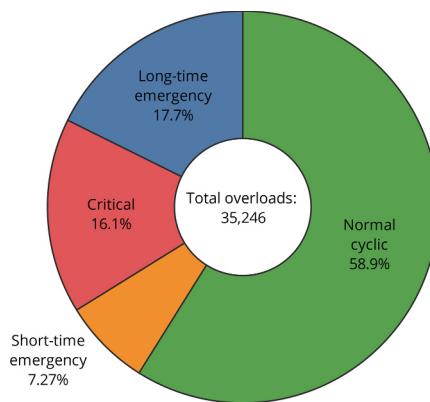
Arrival/Departure Bar Chart: Displays counts per time bin with hoverable agent IDs for anomaly tracing and is shown in Fig. 7.

Overload Distribution: Shows frequency of loading categories based on IEC 60076-7 [15] (Normal cyclic: 100–150%, Long-time emergency: 150–180%, Short-time emergency: 180–200%, Critical: >200%) and is shown in Fig. 8.



Fig. 5. The dashboard's System Analysis view page.

EV Charging Load Heatmap from 2025-01-01 to 2025-01-02

**Fig. 6.** Zoomed version of heatmap.**Fig. 7.** EV arrivals and departure graph.**Fig. 8.** Distribution of transformer overload types.

3) Consumer Analysis

Agent Selection Input: Text box for specific user ID, as shown in Fig. 9. From the System Overview page there is a button that automatically inserts the selected consumer (in the map) to the selection in Consumer Analysis page for fast access to information.

Charging Load Step Chart: User charging schedule with departure/arrival markers as shown in Fig. 10. This graph includes an optional reference overlay if selected from program initialization.

Baseload, Driving Distance & SoC Charts: Step charts showing the baseload consumption, daily driving distance of the consumer, and the SoC of the EV battery throughout the simulation. The SoC chart includes the feature highlighting days with insufficient charging level before departure, resulting in dissatisfaction. This is shown in Fig. 11, with the red cross added to each datapoint for the day with insufficient charge. The explanation for this unexpected behavior is found in a mismatch between the smart charging algorithm's handling of dates and the simulation dates when experiencing a change between summer and winter time.

All visualizations support pan, zoom, and have an export button for high-resolution SVG, with a “publication mode” to enlarge fonts and line styles for use in reports and articles. The publication mode has been used for all individual figures in this paper.

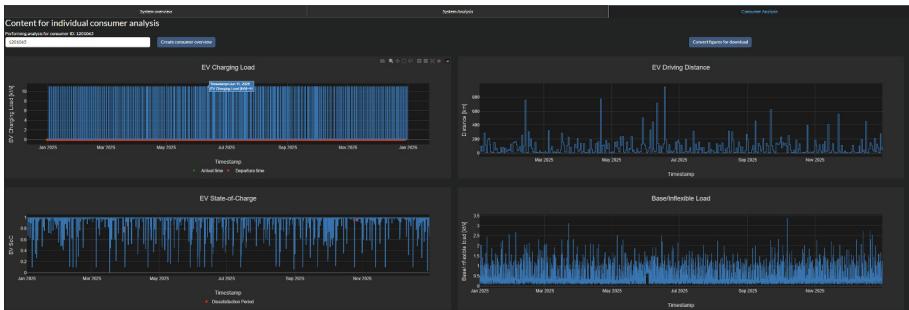


Fig. 9. The dashboards’ Consumer Analysis view page.

4.3 Case Study Highlights

- **Frequent overload events:** The correlation between the number of EVs at home, cheap electricity prices, and smart charging behavior results in daily overloads, with 16% of 587 h of overload being critical ($>200\%$).
- **Unwanted user dissatisfaction:** An analysis of user dissatisfaction events revealed that mismatches in daylight saving time transitions between the smart charging algorithm and simulation timestamps led to insufficient EV charging. This highlights the importance of consistent temporal alignment in simulation pipelines
- **Low utilization of existing capacity:** The low load factor indicates a very uneven distribution of the load, and the high coincidence factor confirms that many consumers locate their maximum load in the same period.

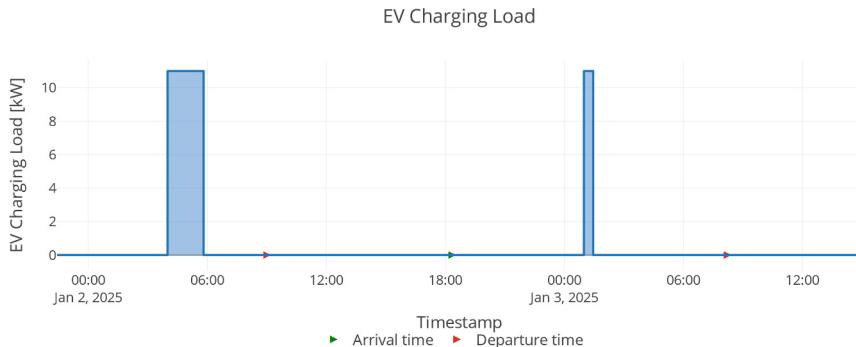


Fig. 10. Charging schedule for individual consumer analysis.

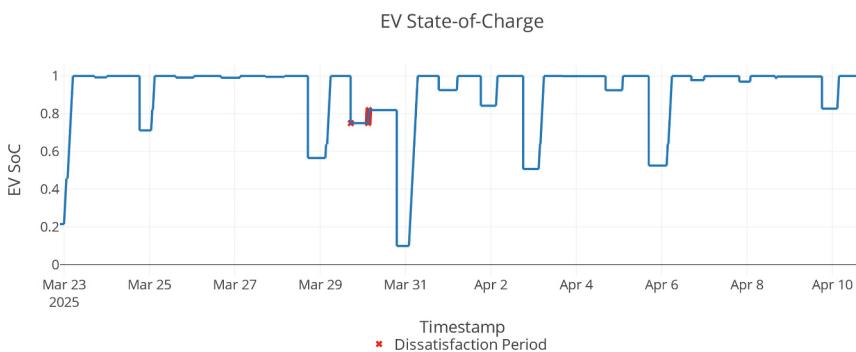


Fig. 11. SoC of individual EV batteries.

A deep understanding of the highlighted results above was rapidly identified and contextualized through synchronized dashboard views, which would have been difficult to detect with static scripts.

The dashboard supports both researchers and DSOs in analyzing scenario performance and investigating unexpected system behaviors. While the tool itself offers a shared set of indicators and interactive views, its application depends on the simulated scenario and user objective. For example, a DSO might explore grid impact under different tariff schemes, while a researcher could examine behavioral responses or test control strategies. In both cases, the dashboard facilitates decision support by making complex results more accessible.

5 Discussion

The proposed dashboard framework facilitates comprehensive analysis of complex outputs from MABS, enabling both system-level pattern recognition and detailed investigation of agent-specific anomalies. By prioritizing the identification and interpretation of

emergent behaviors—such as localized transformer overloads or user-level dissatisfaction events—the tool provides researchers and distribution system operators with a practical means to derive actionable insights. This is particularly valuable when evaluating the effectiveness of strategies like time-of-use tariffs or controlled charging schemes.

A key strength of the framework lies in its integration of high-resolution interactive visualizations with a domain-specific data processing backend. Users can explore large volumes of stochastic simulation data through synchronized views that support anomaly detection and root-cause analysis without extensive manual preprocessing.

The dashboard supports both researchers and DSOs in analyzing scenario performance and investigating unexpected system behaviors. While it offers a shared interface and consistent indicators, its application varies by user objective. For instance, DSOs may focus on infrastructure stress under different tariff schemes, while researchers might examine behavioral adaptations or validate control strategies. This flexibility enhances the tool’s value across both operational planning and academic exploration.

However, the framework does have limitations regarding data volume. In its current implementation, the dashboard has been tested with datasets comprising 1-min resolution measurements over a one-year simulation period, producing CSV files of approximately 8 GB per scenario. Despite the substantial size, the dashboard performs reliably after the initial loading phase, including when comparing two full scenarios (e.g., baseline and test case).

As simulation scales increase—whether in temporal resolution, spatial granularity, or model complexity—future work should consider integrating efficient, queryable storage solutions such as DuckDB or similar in-process analytical databases. These would enable on-demand filtering and aggregation directly on disk, thereby reducing memory overhead and improving responsiveness for larger datasets.

Overall, the framework offers a scalable foundation for exploratory analysis of complex MABS outputs, and its modular design supports future enhancements both in terms of data volume handling and extended visualization capabilities. For example, in the case study, the modular structure allows for simple extension to other distributed energy resources (e.g., home batteries, heat pumps), further increasing the dashboard’s relevance.

While formal usability evaluation is pending, early informal feedback from internal stakeholders highlighted the value of the dashboard’s layered navigation and drill-down capabilities. A structured validation process is planned, including direct involvement of external stakeholders, such as DSOs, to assess usability, interpretability, and operational impact. Future evaluations will follow established protocols such as heuristic analysis and the System Usability Scale, supporting iterative improvements and more robust validation of the tool’s decision-support effectiveness.

6 Conclusion

This paper introduced an interactive dashboard framework tailored for analyzing outputs from multi-agent-based simulations (MABS) of electric vehicle (EV) home charging ecosystems. The tool enables researchers, distribution system operators (DSOs),

and energy planners to explore complex, high-frequency simulation data across multiple analytical levels—from system-wide metrics to individual user behavior—using coordinated visualizations and agent-level drill-downs.

By focusing on emergent phenomena rather than raw computational performance, the framework supports the identification and contextual explanation of unexpected system behaviors, such as clustered transformer overloads or spikes in consumer dissatisfaction. This capability is critical in domains where stochastic agent interactions and temporal flexibility lead to highly dynamic outcomes that are difficult to trace using traditional static plots or aggregated indicators.

The architecture leverages a modular Python backend and Dash by Plotly for front-end interactivity, with efficient data preprocessing strategies to manage year-long, minute-resolution datasets. Through a detailed case study, the dashboard was shown to rapidly uncover system inefficiencies, behavioral anomalies, and potential mismatches in temporal logic—insights that would be challenging to obtain through manual script-based analysis alone.

While current limitations include data scalability and visualization support for non-EV distributed energy resources, the framework is designed for extensibility. Future work will incorporate additional distributed energy resource types such as photovoltaic systems, household batteries, and heat pumps, already present in the simulation engine but not yet visualized. Furthermore, integration with query-optimized data handling (e.g., DuckDB or Apache Arrow) is planned to support the exploration of even larger and more complex simulation scenarios.

Overall, the dashboard lowers the barrier to high-resolution exploratory analysis of MABS outputs and serves as a practical decision-support tool for evaluating the effectiveness of smart grid interventions in complex, agent-driven systems.

Acknowledgement. This paper is part of the project titled “Automated Data and Machine Learning Pipeline for Cost-Effective Energy Demand Forecasting in Sector Coupling” (jr. Nr. RF-23-0039; Erhvervsfyrтn Syd Fase 2), The European Regional Development Fund.

Appendix

Table 1. Dashboard tool feature comparison between popular tools.

Feature	Dash by Plotly [16]	Power BI [17]	Tableau [18]	Streamlit [19]	Shiny [20]
Languages	Python, R, Julia	Power Query (M), DAX, supports Python/R scripts	Drag & drop UI, VizQL; R/Python via extensions	Python	R (original), Python (new)
Visualization Engine	Plotly.js (HTML/SVG/WebGL); deeply integrated	Microsoft proprietary engine (TypeScript/HTML)	VizQL + Hyper engine	Vega-Lite, Plotly, Matplotlib, etc.	ggplot2, Plotly, htmlwidgets, Vega, Leaflet
Interactivity	High: callback-driven, multi-page, fully programmable	High: dashboard filters, drill-throughs	High: filters, parameters, tooltips	Medium: script reruns on input	High: reactive model, dynamic UI

(continued)

Table 1. (continued)

Feature	Dash by Plotly [16]	Power BI [17]	Tableau [18]	Streamlit [19]	Shiny [20]
Scientific Use Fit	Excellent: MABS, time series, system modeling	Limited: Business KPIs; less flexible for research	Moderate: supports R/Python, but GUI-bound	Good: fast for prototyping, but limited structure	Excellent: highly used in academic/statistical apps
Data Volume Handling	Good: needs pre-processing/parquet for large data	Very high (with Premium/DirectQuery)	High (Hyper extracts or live DB connections)	Moderate: in-memory; backend scaling needed	Moderate: in-memory limits; scale with DBs
Deployment	Self-hosted (Flask), Dash Enterprise (proprietary)	Microsoft ecosystem: Desktop, Cloud, Report Server	Tableau Server, Cloud, Public	Local, Docker, Streamlit Cloud	Shiny Server (open-source), RStudio Connect (pro)
Licensing	Open-source SDK (MIT); proprietary hosting optional	Proprietary (no open-source components)	Proprietary (owned by Salesforce)	Fully open-source (Apache 2.0); cloud service optional	Open-source (GPL/MIT); commercial hosting optional
Learning Curve	Moderate: Python + callback logic	Low-moderate: GUI-based, Excel-like	Low-moderate: visual modeling, Tableau calc lang	Very low: basic Python scripting	Moderate: easy to start, advanced logic requires understanding reactivity
Customization & Extensibility	Very high: HTML/CSS, JS, React components	Limited to prebuilt/custom visuals	Moderate: themes, extensions API	Moderate: layout primitives, custom components	Very high: full HTML/CSS/JS control

References

1. Epstein, J.M., Axtell, R.: Growing artificial societies: Social science from the bottom up (Complex adaptive systems.). Brookings Institution Press, Washington, D.C. (1996)
2. Lee, J.-S., et al.: The complexities of agent-based modeling output analysis. *J. Artif. Soc. Soc. Simul.* **18**(4) (2015). <https://doi.org/10.18564/jasss.2897>
3. Akhatova, A., Kranzl, L., Schipfer, F., Heendeniya, C.B.: Agent-based modelling of urban district energy system decarbonisation—a systematic literature review. *Energies* **15**(2), 554 (2022). <https://doi.org/10.3390/en15020554>
4. Mehdizadeh, M., Nordfjærn, T., Klöckner, C.A.: A systematic review of the agent-based modelling/simulation paradigm in mobility transition. *Technol. Forecast. Soc. Change* **184** (2022). <https://doi.org/10.1016/j.techfore.2022.122011>
5. Weinberger, G., et al.: Spatial agent-based modelling and simulation to evaluate on public policies for energy transition. In: Proceedings of the 3rd International Sustainable Energy Conference (ISEC 2024) (2024)
6. Christensen, K., Jørgensen, B.N., Ma, Z.G.: A multi-agent, laxity-based aggregation strategy for cost-effective electric vehicle charging and local transformer overload prevention. *Sustainability* **17**(9), 3847 (2025). <https://www.mdpi.com/2071-1050/17/9/3847>

7. Malleson, N., Heppenstall, A., See, L., Evans, A.: Using an agent-based crime simulation to predict the effects of urban regeneration on individual household burglary risk. *Environ. Plan. B Plan. Des.* **40**(3), 405–426 (2013). <https://doi.org/10.1068/b38057>
8. Bastariano, F.F., Hancock, T.O., Choudhury, C.F., Manley, E.: Agent-based models in urban transportation: review, challenges, and opportunities. *Eur. Transp. Res. Rev.* **15**(19) (2023). <https://doi.org/10.1186/s12544-023-00590-5>
9. Lu, J., et al.: AgentLens: visual analysis for agent behaviors in LLM-based autonomous systems. *TVCN* **31**, 4182–4197 (2024). <https://doi.org/10.1109/TVCN.2024.3394053>
10. Foramitti, J.: AgentPy: a package for agent-based modeling in Python. *J. Open Source Softw.* **6**(62), 3065 (2021). <https://doi.org/10.21105/joss.03065>
11. ter Hoeven, E., et al.: Mesa 3: agent-based modeling with Python in 2025. *J. Open Source Softw.* **10**(107), 7668 (2025). <https://doi.org/10.21105/joss.07668>
12. Grignard, A., Drogoul, A.: Agent-based visualization: a real-time visualization tool applied both to data and simulation outputs. In: Proceedings of the AAAI-17 Workshop on Human-Machine Collaborative Learning (AAAI Spring Symposium) (2017)
13. Munzner, T.: *Visualization Analysis and Design*. A K Peters Visualization Series, 1st edn. A. K. Peters, Natick, MA, United States (2014)
14. Schneiderman, B.: The eyes have it: a task by data type taxonomy for information visualizations. In: Proceedings 1996 IEEE Symposium on Visual Languages, 3–6 September 1996, pp. 336–343 (1996). <https://doi.org/10.1109/VL.1996.545307>
15. Power Transformers – Part 7: Loading Guide for Mineral-Oil-Immersed Power Transformers. International Electrotechnical Commission, Geneva, Switzerland (2018). <https://studylib.net/doc/25758704/iec60076-7%7Bed2.0%7Den>
16. Plotly: Dash by plotly documentation page. <https://dash.plotly.com/>. Accessed 17 May 2025
17. Microsoft Power BI: Microsoft Power BI product page. <https://www.microsoft.com/da-dk/power-platform/products/power-bi?market=dk>. Accessed 17 May 2025
18. Tableau: Tableau web page. <https://www.tableau.com/>. Accessed 17 May 2025
19. Streamlit: Streamlit web page. <https://streamlit.io/>. Accessed 17 May 2025
20. Shiny: Shiny web page. <https://shiny.posit.co/>. Accessed 17 May 2025

Author Index

A

Abouebeid, Sara II-173
Aguayo, Cristian II-187
Ahammer, Florian II-95
Aiswarya, T. S. II-78
Ajwani, Deepak II-271
Akashi, Yasunori II-114, II-144, II-154
Alaliyat, Saleh Abdel-Afou I-3
Alex, A. II-78
Al-Habaibeh, Amin I-198
Amideo, Annunziata Esposito I-255
Amini Toosi, Hashem II-221
Amos, Matt II-321, II-338, II-371
Ayyathurai, Vignesh Pechiappan I-95

B

Bansal, Ramesh C. I-309
Becker, Robert A. II-285
Beigaité, Rita II-35
Bengtsson, Gustaf II-321, II-338, II-371
Billanes, Joy Dalmacio I-79, I-215, II-254
Biswas, Debajyoti II-187
Blaabjerg, Frede I-60
Bleistein, Thomas II-285
Bolaños-Zuñiga, Johanna II-67
Bordin, Chiara I-60
Botea, Adi II-271

C

Caetano, Luis II-130
Campodonico-Avendano, Italo Aldo II-300
Carroll, Paula I-255, I-267, II-187, II-271
Çeliker, Cihat Emre II-338
Chen, Zhonghe II-271
Christensen, Kristoffer I-371
Clauß, John II-130
Cong, Lu I-237

D

Daniel, Lius II-35
Davidsson, Henrik I-131

Derbas, Abd Alelah I-60

E

Engström, Jesper I-131

F

Faisal, Mohammed Farhan I-15
Fellerer, Jonathan II-3
Fernengel, Natascha I-323
Ford, Ryan II-35
Förderer, Kevin I-323

G

Garcia, Rafeal Gomez II-321
German, Reinhard II-3
Gölles, Markus II-95
Guðjónsdóttir, Maria Sigríður II-23
Gudmundsson, Steinn I-341

H

Hagenmeyer, Veit I-323
Hajsok, Lucija I-283
Hameed, Akram II-321, II-338, II-371
Harini, N. II-78
Hasan, Agus I-3
He, Yue II-144
Hehemann, Michael I-323
Helbrecht, Jana II-3
Holly, Stefanie I-354
Holtwerth, Alexander I-323
Hoppe, Eugen I-323
Huang, Haoyu I-323
Huang, Pei I-299

J

Johansson, Dennis I-131
Johra, Hicham II-300
Jørgensen, Bo Nørregaard I-215, I-237, I-371, II-254, II-321, II-338, II-355, II-371, II-388

K

Kamenev, Nikolai [II-285](#)
 Kandemir, Ege [I-3](#)
 Kayo, Genku [I-110](#)
 Khajoei, Najmeh [I-341](#)
 Kind, Reidar [II-130](#)
 Koelsch, Celina [II-285](#)
 Kowli, Anupama [II-204](#)
 Kulkarni, Aashay [II-285](#)

L

L., Alberto J. Lamadrid [II-67](#)
 Lassen, Thomas Elvrum [II-130](#)
 Liberado, Eduardo [I-185](#)
 Lichtenegger, Klaus [II-95](#)
 Lien, Synne Krekling [II-321, II-338](#)
 Lin, Jeremy [I-309](#)
 Lygnerud, Kristina [I-79](#)

M

Ma, Zheng Grace [I-215, I-237, I-371, II-321, II-338, II-388](#)
 Ma, Zheng [I-79, II-254, II-355, II-371](#)
 Madsen, Frederik Wagner [II-254](#)
 Malakhatka, Elena [II-173](#)
 Marumoto, Shoya [II-114](#)
 Mazidi, Mohammadreza [II-173](#)
 Medjdoub, Benachir [I-198](#)
 Mehmood, Khawaja Khalid [I-32](#)
 Miller, Clayton [II-338, II-371](#)
 Mishra, Sambeet [I-60](#)
 Miyata, Shohei [II-114, II-144, II-154](#)
 Moazami, Amin [II-300](#)
 Mori, Taro [I-149, I-164, I-173](#)
 Moura, Ranier Alexander Arruda [I-32](#)
 Müller, Dirk [I-323](#)
 Muschick, Daniel [II-95](#)
 Mutule, Anna [II-187](#)

N

Najmadin, Abdelmomen [I-95](#)
 Nguyen, Phuong Hong [I-32](#)
 Niklasson, Felix [II-173](#)
 Nixon, Nimisha [I-15](#)
 Nordanger, Knut [II-130](#)

O

Osawa, Hisato [I-149, I-164, I-173](#)
 Otani, Kazuma [I-164, I-173](#)

P

Pandiyan, Surya Venkatesh [II-321](#)
 Parvaz, Md [I-131](#)
 Plompen, Hendrik [I-32](#)
 Pradeep, Jayarama [II-78](#)
 Þorvaldsson, Einar Örn [I-131](#)

R

Radtke, Malin [I-354](#)
 Rajasekharan, Jayaprakash [II-321](#)
 Ravindra, Nikhil [I-198](#)
 Richter, Christiaan [I-230](#)

S

Safavi, Aysan [I-230](#)
 Santana, Emerson [I-185](#)
 Sartori, Igor [II-321, II-338, II-371](#)
 Schopper, Fabian [II-95](#)
 Seema, [II-51](#)
 Serrao, Nikith Jude [II-78](#)
 Sharma, Desh Deepak [I-309](#)
 Shé, Clíodhna Ní [I-267](#)
 Shi, Shanrui [II-144, II-154](#)
 Sirjani, Reza [II-51](#)
 Slimani, Mohamed El-Amine [II-23](#)
 Somawanshi, Aditya [II-204](#)
 Sridhar, Araavind [II-173](#)
 Stark, Sanja [I-354](#)
 Steen, David [II-173](#)
 Suzuki, Konatsu [I-149, I-164, I-173](#)

T

Taniguchi, Keiichiro [II-114](#)
 Theocharis, Andreas [II-51](#)
 Titz, Maurizio [II-237](#)
 Tolnai, Balázs András [II-321, II-338, II-355, II-371, II-388](#)
 Tran, Ngoc [II-338](#)
 Tuan, Le Anh [II-173](#)

U

Unnþorsson, Runar [I-230, I-341](#)

V

van der Molen, Anne [I-32](#)
van der Wielen, Peter [I-32](#)
Varma, Pamba Raja [I-15](#)

W

Waczowicz, Simon [I-323](#)
Walker, Sara [I-119](#)
Wallbaum, Holger [II-173](#)
Walnum, Harald Taxt [II-321](#)
White, Stephen [II-371](#)
Witthaut, Dirk [II-237](#)

X

Xhonneux, André [I-323](#)
Xiang, Jingyu [I-255](#)
Xie, Yangxinyu [II-338](#)

Y

Yan, Hui [I-119](#)
Yang, Aileen [II-300](#)

Z

Zafar, Rehman [I-299](#)
Žakula, Tea [I-283](#)
Zimmermann, Rafael Sudbrack [II-338](#)