





Showcasing research from the Institute of Energy and Climate Research, Theory and Computation of Energy Materials (IEK-13) and Electrochemical Process Engineering (IEK-14), Forschungszentrum Jülich GmbH, Germany.

Deep learning-enhanced characterization of bubble dynamics in proton exchange membrane water electrolyzers

This study introduces a deep learning tool for analyzing bubble dynamics in water electrolyzers using optical imaging. Our software rapidly processes thousands of frames to deliver detailed visualizations of bubble ratios, distributions and shape analyses. These insights advance understanding of cell behavior across different voltage levels, enhancing diagnostic techniques. This tool offers a step forward in the mechanistic interpretation of bubbles, contributing to the characterization of energy materials.

As featured in:



See Mohammad J. Eslamibidgoli *et al.*,
Phys. Chem. Chem. Phys.,
2024, **26**, 14529.



Cite this: *Phys. Chem. Chem. Phys.*,
2024, 26, 14529

Deep learning-enhanced characterization of bubble dynamics in proton exchange membrane water electrolyzers

André Colliard-Granero,^{†ab} Keusra A. Gompou,^{†a} Christian Rodenbücher,^{id c}
Kourosh Malek,^{ab} Michael H. Eikerling,^{id abd} and Mohammad J. Eslamibidgoli,^{id *ab}

The ever-increasing utility of imaging technology in proton exchange membrane water electrolyzer research raises the demand for rapid and precise image analysis. In particular, for optical video recordings, the challenge primarily lies in the large number of frames that impede the delineation of bubble dynamics with standard methods. In order to address this problem, the present study supports the automation of data analysis to facilitate swift, comprehensive, and measurable insights from captured imagery. We present a deep learning-based framework to perform high-throughput analyses of bubble dynamics using optical images of proton exchange membrane water electrolyzers. Leveraging a relatively small annotated imaging dataset of just 35 images, various configurations of the U-Net architecture were trained to perform bubble segmentation tasks. The best model achieved a precision of 95%, a recall of 78%, and an F1-score of 86% on the validation set. Subsequent to segmentation, the methodology enabled the rapid extraction of parameters such as time-resolved bubble area, size distributions, bubble position probability density, and individual bubble shape analytics. The findings underscore the potential of deep learning to enhance the analysis of polymer electrolyte membrane water electrolyzer imaging, offering a path toward more efficient and informative evaluations in electrochemical research.

Received 1st December 2023,
Accepted 4th March 2024

DOI: 10.1039/d3cp05869g

rsc.li/pccp

1 Introduction

Among the technologies for hydrogen production, polymer electrolyte membrane water electrolyzers (PEMWE) stand out for their high current densities, reduced propensity for gas crossover phenomena, and more compact cell design. However, challenges persist, including the high cost of construction materials, reduced membrane resistance at elevated pressures, and susceptibility to cross-osmosis phenomena.^{1,2}

In particular, the generation of oxygen gas bubbles at the anode of the catalytic layer is a significant factor inhibiting the access of water to the reactive sites, particularly under

conditions of high current densities.^{2,3} Consequently, several studies focus on optimizing the two-phase flow behavior as a means to increase PEMWE performance. For instance, several works focus on exploring the effects of flow field channel geometry on cell efficiency, comparing parallel and serpentine flow fields has been conducted. Majasan and coworkers found that under a range of relevant operating conditions, including variations in water inflow rate, fluid dynamics, and temperature, a parallel channel configuration is more efficient than a serpentine one.⁴

The relationship between flow regimes and cell current density is another critical area of study. Dedigama *et al.* provided evidence that the presence of larger bubbles can displace more water, facilitating the removal of smaller bubbles adhered to the electrode surface, and ultimately improving water transport.⁵ Aubras *et al.* demonstrated that the transition from effervescent flow of smaller bubbles to slug flow dominated by larger bubbles enhances overall mass transport, diminishes ohmic resistance, and contributes to improved cell efficiency.⁶ Generally, the dynamics of bubble coverage, along with the concurrent processes of bubble growth/detachment and electrochemical reactions, are recognized as crucial to cell performance. Su *et al.* proposed that a high bubble coverage can be

^a Theory and Computation of Energy Materials (IEK-13), Institute of Energy and Climate Research, Forschungszentrum Jülich GmbH, 52425 Jülich, Germany.
E-mail: m.eslamibidgoli@fz-juelich.de

^b Centre for Advanced Simulation and Analytics (CASA), Simulation and Data Science Lab for Energy Materials (SDL-EM), Forschungszentrum Jülich GmbH, 52425 Jülich, Germany

^c Electrochemical Process Engineering (IEK-14), Institute of Energy and Climate Research, Forschungszentrum Jülich GmbH, 52425 Jülich, Germany

^d Chair of Theory and Computation of Energy Materials, Faculty of Georesources and Materials Engineering, RWTH Aachen University, 52062 Aachen, Germany

[†] Equal contribution.



indicative of a deficient water supply to the cell, leading to a higher voltage and, decreased efficiency.⁷ These findings collectively underscore the complex interplay between flow design, bubble dynamics, and cell efficiency.

Understanding the processes of bubble formation, growth, and detachment is pivotal in the context of PEMWE, where their interplay determines the two-phase flow dynamics.⁶ A gap exists in terms of characterization techniques, necessitating advanced methodologies to elucidate the complex behavior of bubbles in PEMWE. Conventionally, the evaluation of bubble dynamics has relied on the analysis of extensive datasets acquired *via* indirect techniques such as acoustic emission (AE), electrochemical impedance spectroscopy (EIS), and pressure drop measurement. AE serves as a non-destructive, *operando* diagnostic tool utilizing a piezoelectric sensor to capture mechanical disturbances emanating from an object.⁸ This methodology facilitates the determination of bubble size distribution and the elucidation of diverse gas-flow patterns through the analysis of acoustic waves generated by bubbles.^{9,10}

However, AE has limited spatial resolution and susceptibility to background noise, which can impede the accurate localization and size determination of bubbles. EIS offers insights into the individual contributions to an electrochemical system's total impedance, encompassing processes such as interfacial charge transfer and mass transport.¹¹ Nonetheless, interpreting impedance data, especially in systems like PEM electrolyzers with evolving bubbles, is challenging. Especially, distinguishing between the various contributing factors, *e.g.* charge transfer resistance, mass transport limitations, and bubble dynamics, in the impedance spectrum. Conversely, pressure drop measurement, utilizes pressure sensor probes to detect signals that reflect bubble behavior, allowing for the correlation of these signals to dynamic bubble phenomena, as discussed by Zhang *et al.*¹² Therefore, the integration of indirect techniques with direct observational methods can offer a more holistic understanding of bubble dynamics.

Optical photography facilitates the observation and quantification of bubble behavior.¹³ While this technique is advantageous for transparent cells, offering unobstructed visualization, the quantification of important parameters can still pose significant challenges. The manual processing of extensive video data sets is not only laborious but also prone to subjective interpretations that can vary with the analyst's bias. The advent of sophisticated image analysis coupled with the evolution of artificial intelligence (AI) offers a compelling solution to these limitations. These technological advances have laid the groundwork for automated, objective extraction of data from image sequences, effectively circumventing the issues of manual analysis.^{14–17}

Recently, Sun *et al.* performed a study where a multi-task deep learning (DL) network was employed for instance segmentation to elucidate fission gas bubbles within nuclear fuel.¹⁸ Anderson *et al.* utilized DL to autonomously identify helium bubbles in irradiated micrographs and to extract their radii and cumulative volumes. The model exhibited a high accuracy, achieving a 93% success rate in the detection of bubbles on

high-magnification micrographs, and maintained robust performance in analyzing lower magnification samples.¹⁹ Kim and Park introduced an instance segmentation model designed to autonomously discern and delineate the contours of bubbles across diverse flow conditions, attaining an average precision (AP50) of 98%.²⁰ Nevertheless, DL for bubble analysis in PEMWE remains an under-explored field. The closest work found in the literature employs a state-of-the-art object detection framework, YOLOv7, for the identification of anodic oxygen bubbles within a transparent PEMWE system.² While this method facilitates the extraction of bubble features, including the area, their count, and the extent of bubble coverage across the electrolyzer, it relies on bounding boxes to approximate bubble detection, a strategy that is prone to inaccuracies, especially when confronted with bubbles that have coalesced or present irregular shapes. Hence, there is a widely recognized need for methodologies capable of delineating the actual contours of individual bubbles during their formation, growth and release times, to enhance the precision of feature extraction and analysis.

2 Methods

This manuscript presents a detailed, automated deep learning-based workflow for analyzing video footage of oxygen bubbles in proton exchange membrane water electrolyzers (PEMWE). The proposed workflow is twofold: firstly, it involves the creation of a meticulously annotated dataset, which is then used to train models for the semantic segmentation of bubbles; secondly, it facilitates the automated extraction of bubble size and shape characteristics, which are crucial for the analysis of their distribution properties. The practicality of this approach is exemplified by successfully training models using a relatively small dataset of only 35 images, demonstrating that even limited data can yield insights for specific use cases. The dimensions of the dataset were determined based on our prior research, in which we empirically established the optimal quantity of images required to achieve saturation in the performance of the U-Net model. This investigation facilitated a balance between model efficacy and the time expended on annotations.¹⁶ Furthermore, this study offers a flexible framework, allowing users to replace the provided model with a bespoke one tailored to their own datasets, leveraging the robust computer-vision capabilities developed in this research for extraction of time-resolved size and shape analysis of the regions of interest. This adaptability makes the workflow a powerful tool for researchers seeking to enhance their analysis of electrolyzer performance through advanced image processing techniques.

To demonstrate the application of deep learning for the localization of gas bubbles in the flow field of a PEM electrolyzer, we employed a commercial reversible fuel cell (FCSU-023, Horizon Fuel Cell Europe, Czech Republic) operated in the electrolysis mode. The housing and the flow field of the cell consisted of green-blueish transparent plastic allowing for *operando* optical inspection of bubble evolution (Fig. 1a).






Fig. 1 (a) Schematic experimental setup for the recording of the dataset. (b) Exemplification of the obtained data from the experiment from the PEMWE transparent cell and the oxygen bubbles.

The cell was driven by a power supply in constant voltage mode (2651A, Keithley, USA), which also served for measuring the current. Before starting each electrolysis run, the anode compartment of the cell was filled with deionized water, which was then consumed during the process as no water circulation was applied. Images of the anode, where the oxygen bubbles were released, were recorded by a digital camera (Lumix TZ 18, Panasonic, Japan) with a frame rate of 30 fps and a resolution of approximately 25 px mm^{-1} . The images were cropped to the active electrode area of approximately $25 \times 25 \text{ mm}^2$ as shown in Fig. 1b, with the water inlet visible in the bottom right and the gas outlet in the top left part. The cyclic voltammogram of the cell obtained at a speed of 10 mV s^{-1} is shown in Fig. 2a. The onset of the electrolysis process can be seen as steep increase of the current at 1.5 V . Hence, the images were recorded at four different voltages of 1.5 , 1.6 , 1.7 , and 1.8 V while the voltage was applied for approximately 400 s each.

Our methodology for analyzing bubble dynamics in videos encompasses four principal steps, as illustrated in Fig. 3: (1) manual annotation of oxygen bubbles to prepare the training dataset; (2) employing supervised learning for the semantic segmentation of bubbles; (3) extracting features automatically from the regions of interest (ROI); and (4) statistical analysis and visualization of the findings. Initially, we meticulously annotate the oxygen bubbles, as the precision of segmentation relies heavily on the quality of this labeled data. To facilitate this process, we used Label Studio,²¹ a comprehensive and intuitive annotation tool. This software facilitates efficient annotation and the automated export of masks in the desired format.

For the training and validation of our model, we annotated thirty-five optical frames. Originally, each frame measured

640×630 pixels, but for the training process, we resized them to 512×512 pixels. This resizing was crucial to optimize the model's performance, as it specializes in analyzing squared cells. Furthermore, maintaining a minimum resolution of 512×512 pixels is essential to prevent any potential resolution-related issues. These frames were carefully chosen from various parts of the experiment to represent a broad range of bubble scenarios. This deliberate choice is aimed at improving the model's generalization abilities. An expert, using visual cues, performed the detailed task of annotation, with a special focus on bubbles that were difficult to differentiate. To guarantee the reliability of the annotations, two additional reviewers examined them to achieve a unified consensus on the bubble selection.

In this study, utilizing the provided images and their associated annotated masks as ground truth, we focused on the pixel-wise identification of oxygen bubbles through the implementation of several U-Net-based architectures. These architectures are based on fully convolutional neural networks, which conduct a series of convolutions and down-sampling operations to identify the key patterns present within the images in a latent space. The encoding process progresses iteratively until the entirety of the information is consolidated into a singular vector. Subsequently, a reverse decoding process is initiated: up-sampling, coupled with the application of transpose convolutions, restores the data to its original size. A salient feature of the U-Net architecture is the concatenation of the up-sampled decoder feature map with the encoder feature maps of a matching resolution. This technique facilitates the discernment of object boundaries and edges, harnessing both low-level and high-level features, culminating in a refined segmentation output.²²

Within the scope of our research, we experimented with the canonical U-Net framework along with two advanced variations. The first is a U-Net model augmented with a ResNeXt101 network as its backbone, integrating residual connections. The second is the more recent attention U-Net model, which capitalizes on attention mechanisms to further enhance segmentation precision.

The ResNeXt101 U-Net architecture fuses the ResNeXt101 model's capabilities with the U-Net structure. ResNeXt101, a 101-layer variant of the ResNeXt series, utilizes the split-transform-merge strategy. Its unique feature is "cardinality" – parallel paths in a block. Instead of increasing depth or width, elevating cardinality enhances performance. By using ResNeXt101 as U-Net's backbone, the network captures intricate image features more effectively.²³

Attention U-Net augments the traditional U-Net with an attention mechanism, enabling the model to concentrate on specific image regions. This is vital for detecting various object scales or minor yet crucial image sections. Attention gates, applied before each decoder concatenation step, weigh encoder features, intensifying key pixels and suppressing irrelevant ones. This heightens segmentation precision, especially when target regions are surrounded by noise.²⁴

In our study, we evaluated the performance of our trained model using standard evaluation metrics, notably the




Fig. 2 (a) Cyclic voltammogram of the electrolyzer. (b) Current as function of time during the application of different constant voltages.






Fig. 3 Methodological pipeline for the video analysis of oxygen bubbles in PEM electrolyzers.

intersection over union (IoU) threshold. IoU is a commonly used metric in image segmentation, providing a quantitative assessment of the overlap between the model's predictions and the actual annotated masks. It calculates the proportion of overlapping pixels to the total number of pixels present in both the predicted and true masks for the target class.

Typically, a prediction is classified as accurate if the IoU score surpasses 0.5. Under this paradigm, a true positive (TP) signifies a scenario where a pixel is accurately identified as a bubble, evident from an IoU score greater than 0.5. Conversely, a true negative (TN) marks a pixel correctly identified as part of the background. Furthermore, a false positive (FP) arises when a pixel, presumed by the model to represent a bubble, lacks any corresponding annotation in the ground truth. On the other hand, a false negative (FN) emerges when a pixel, indicative of a bubble in the ground truth, is either overlooked or misclassified by the model.

Fig. 4 shows examples of TP, TN, FP, and FN classifications. Following the determination of these fundamental metrics, we

can proceed towards more advanced evaluation metrics like precision, recall, and the F1 score, defined as:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

$$\text{F1} = \frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$$

Utilizing the segmentation maps generated from model predictions, we performed further statistical analysis. By employing elementary mathematical procedures, we calculated the pixel count representing bubbles in each frame, offering an estimation of frame-wise bubble coverage. This analysis facilitated the creation of plots illustrating time-resolved visual coverage and bubble area distributions.

Furthermore, by overlaying these masks, we derived density maps that provide insights into the spatial distribution of bubbles over the course of the experiment. Focusing on the distinct contours of bubbles in the masks, we formulated computer vision algorithms using the OpenCV library.²⁵ These algorithms were adept at identifying individual bubbles and subsequently extracting a plethora of shape characteristics, enhancing our understanding of bubble dynamics and morphology.

3 Results and discussion

In the study of proton exchange membrane water electrolyzers (PEMWE) videos, the accurate segmentation of bubbles is of




Fig. 4 Demonstration of an optical image along with the ground truth and predicted masks. Examples of true-positive (TP), true-negative (TN), false-positive (FP), and false-negative (FN) predictions are shown on the predicted mask.



great importance for detailed analysis. Our methodology for this segmentation task begins by choosing an appropriate neural network architecture tailored for optimal feature extraction and segmentation capability. To enhance the model's performance, we initialize the model's weights using a pre-trained model, capitalizing on previously learned features. The learning rate, which plays a crucial role in the convergence of the training process, is then optimally selected to determine the magnitude of steps taken toward the minimum of the loss function during training. The primary objective during training is to minimize a loss function that quantifies the discrepancy between manually annotated masks and the model's predictions. A reduced loss indicates superior model performance. After the training process, we critically assess the model's effectiveness using a validation set to determine its robustness and generalization capability.

In our research, we undertook a benchmarking study, comparing three distinct variations of the U-Net model. Firstly, we evaluated the conventional U-Net 2D model, which was trained using the ZeroCostDL4Mic approach.²⁶ This was compared with a custom U-Net 2D model that integrated a pre-trained ResNeXt101 from the ImageNet dataset as its backbone. The third model was the more contemporary attention U-Net. The models were trained on the same dataset consisting of 28 images for training and 7 images for validation for 100 epochs with a batch size of 8. This training involved segmenting the images into patches and employing a dynamic learning rate which was adjusted in real time based on model performance across epochs.

As illustrated in Table 1, we present a comprehensive array of metrics for these models. In terms of F1 scores, the models showcased similar performance levels; however, the U-Net model augmented with the ResNeXt101 backbone slightly outperformed the others. It was observed that the standard U-Net architecture achieved lower precision metrics but with a compensatory spike in recall compared to the other two architectures. In the context of our bubble segmentation task, the model's precision is especially relevant. A diminished precision hints at the introduction of non-existent bubble pixels, which reduce the interpretability and accuracy when analyzing features extracted from videos.

In addition to the models' performances, the decisive step remains the visual evaluation of the segmentation outcomes. To this end, Fig. 5 offers a comparative visualization with sub-figure (a) depicting a frame randomly sourced from the validation set, and sub-figure (b) showcasing the segmentation results as generated by the different models under study.

Table 1 Comparison of various U-Net models trained on an identical dataset, evaluated using the metrics: precision, recall, and F1-score

Model	Precision [%]	Recall [%]	F1-score [%]
U-Net 2D	81	89	85
U-Net with ResNeXt101 backbone	95	78	86
Attention U-Net	95	75	84

A close inspection of these frames bolsters our conviction in the superior performance of the U-Net model incorporating the custom ResNeXt101 backbone. Notably, this model yields the most seamless contours around the segmented bubbles and achieves good results in separating bubbles in close proximity—a feat not as pronounced in the other models in our dataset.

To provide a more granular view of our results, we have rendered videos of the segmentation outputs. These videos, available on our GitHub repository, allow for a detailed examination and facilitate a deeper appreciation of the differences in each model's performance.

The implementation of automatic video segmentation paves the way for an enhanced software utility capable of extracting intrinsic data from the bubble dynamics unfolding within the cell. These parameters are instrumental for parametrization in simulation sciences and hold the potential to expedite investigations within experimental cohorts. Further, model predictions reveal multifaceted insights pertaining to the bubble dynamics of the utilized PEM electrolyzer.

Fig. 6 delineates the time-resolved bubble ratio and presents its corresponding histogram for four distinct voltages but retaining an identical experimental configuration. A comparative analysis reveals differences across the different experiments. As anticipated, a reduced voltage corresponds to decelerated gas generation, marked by merely four significant gas releases and a diminished mean bubble coverage. Conversely, an increased voltage shows an augmented visual bubble coverage of the electrode, and the time-sequence plots show a higher periodicity in gas releases.




Fig. 5 (a) Original image and the corresponding ground truth. (b) Prediction output for the three bench-marked models. Gray pixels indicate true positives for correctly detected bubbles, blue pixels represent false negatives where bubble pixels were missed, and red pixels signify false positives for erroneously detected non-bubble pixels. The U-Net model with the ResNeXt101 backbone achieves the smoothest predictions and the highest model performance in comparison to the other tested models. A clear video comparison is available under the project repository at <https://github.com/andyco98/UTILE-Oxy>.






Fig. 6 Time-resolved bubble ratio evolution alongside the total bubble area distribution for four distinct voltages, all conducted under the same experimental setup. The total bubble area distribution is further analyzed with a Gaussian fit, providing insights into the mean visual bubble coverage and its standard deviation.

An interesting result obtained through our methodology is the spatiotemporal distribution of bubbles, visualized as a density map, showing bubble position probabilities over a time series. These maps are constructed using model predictions, where each frame's pixels are classified into two categories: '0' representing the background and '1' for the bubbles. Upon aggregating all these predictions, the regions most consistently occupied by bubbles across the time series attain a higher value, which is depicted in the map with warmer hues. Conversely, regions with infrequent bubble presence are represented with cooler colors. Intermediary zones, represented in green, elucidate regions within the cell where bubbles tend to coalesce and accumulate prior to release. Such maps offer a correlation with bubble nucleation points, shedding light on predominant bubble generation zones and their distribution across the cell.

Fig. 7 exemplifies the utility of these density maps. Presented therein are four distinct maps corresponding to a single experiment across varied voltage levels. At lower voltage settings, the map reveals fewer warm-colored regions, indicative of bubble nucleation, and a relatively random distribution pattern across the cell. As the voltage is incremented, there's a palpable increase in these nucleation regions, complemented by a clearer depiction of bubble accumulation zones—the latter exhibiting a subtle proclivity towards the cell's right side. Another salient observation gained from these heat maps is the remarkable consistency in nucleation points across experiments, especially at higher voltage levels.

Our methodology facilitates an individual analysis of bubbles, focusing on their unique morphological attributes. Analyzing the predicted masks, bubbles fully encapsulated by background pixels are classified as distinct entities. Employing specialized computer vision techniques, we derive a series of




Fig. 7 Density maps depicting the bubble spatial distribution probability throughout the experiment at four distinct voltages, all conducted under the same experimental setup. Areas with warmer colors indicate positions frequently occupied by bubbles, while colder colors suggest less frequented areas.

attributes elucidating the structural properties of individual bubbles detected throughout the experimental phase. These commonly calculated properties in the realm of materials science include:^{27–29}

- Area: denoted by the pixel count of an individual bubble.
- Diameter: inferred under the assumption that the bubble's area aligns with a perfect circle.



• Aspect ratio: derived by inscribing a bounding box around the bubble and calculating the ratio of its longest to shortest side.

• Solidity: a metric indicative of the bubble's convexity. It is defined for a given shape as the ratio of the area of the shape to the area of its convex hull, represented as:

$$\text{Solidity} = \frac{A_{\text{shape}}}{A_{\text{convex hull}}}$$

Being the convex hull of a shape the smallest convex set that contains the shape.

- Orientation: determined by superimposing an ellipse on the bubble and calculating the tilt angle of this fitted ellipse.
- Perimeter: represents the contour length of a bubble.
- Extent: defines the bubble's squareness.
- Roundness: assesses the bubble's resemblance to an ideal circle.

This set of parameters, when used together, provides a detailed geometric characterization of bubbles, aiding in a more thorough understanding of their behavior and interaction with surrounding environments.

Given that our data acquisition is frame-centric, it permits temporal analysis. As illustrated in Fig. 8, time-series data for a




Fig. 8 Time-resolved analysis of bubble property evolution for one experiment at 1500 mV. Results are visualized using a moving average calculated every hundred frames, with the standard deviation shown as a shaded region around the line.



specific experiment conducted at four distinct voltage levels are plotted. Initially, bubbles are individually analyzed and subsequently averaged on a per-frame basis. Given the volatility in individual data points, which can make the visualization of emergent patterns difficult, we opted for data pre-treatment prior to visualization. We employed a moving average approach, wherein an average spanning 100 frames was taken, which then informed the central line of our plot. The accompanying standard deviation is depicted as a shaded region around this line.

This single bubble data provides a comprehensive perspective on the morphological trends of bubbles across temporal intervals. A particular revelation from this data is related to bubble uniformity. For example, smaller standard deviations in the area signify periods of systemic uniformity, indicating a preponderance of similar bubbles. Conversely, pronounced standard deviations, *e.g.* with respect to area, indicate a bifurcation between small and large bubbles, suggestive of system heterogeneity.

The results shown in this work, establish a basis for a detailed, time-resolved characterization of bubble system dynamics within the cell, facilitated through automation.

Additionally, combining the various characteristics previously mentioned can result in a solid manner to characterize particular systems of interest under examination. From the time-series bubble coverage data, we can discern vital information about cell cycling, the periodicity of prominent gas releases, average bubble coverage, and its inherent volatility. The density maps are instrumental in visually pinpointing nucleation sites and frequent bubble aggregation zones, thereby shedding light on the distribution of active sites and the efficacy of the gas flow field. The contours and characteristics of individual bubbles further provide insights into the bubble formation dynamics and how the environment influences specific bubble shapes. For instance, comparing the average bubble area and its associated standard deviation with the average coverage metrics from the histograms and the active zones from the density maps enables elucidation of the underlying bubble dynamics.

4 Conclusions

This research marks a crucial advancement for the autonomous analysis of bubble dynamics in proton exchange membrane water electrolyzers, utilizing optical videos as a rich source of data. The autonomous software developed yields substantial results, including detailed visualizations of time-resolved bubble ratio evolution, bubble area distribution, bubble position probability density maps, and detailed analyses of individual bubble shapes. These results provide vital information on the active area of the cell at different voltage levels, facilitating a deeper understanding of bubble dynamic developments across varied experimental settings and voltage ranges.

We envisage that our software will significantly enhance the depth and complexity of bubble dynamics analysis within a

PEMWE cell. This is achieved by enabling the interpretation of data-rich yet easily conducted experimental recordings. Creating an accessible path to obtain such detailed data is essential for augmenting the sophistication of simulation models aimed at understanding these phenomena and resulting in insightful diagnostic techniques. For experimentalists, this tool offers the dual advantage of rapid yet robust characterization, facilitating iterative testing of diverse cell geometries and setups. Furthermore, the workflow is able to analyze several thousands of frames in a matter of minutes enabling an in-depth analysis, unfeasible to obtain manually in a reasonable time frame or cost.

Looking forward, this work will be employed to facilitate further research aimed at deciphering the complexities of bubble dynamics in relation to varying voltages and gas flow fields. Combining this approach with physical modelling, it is possible to unfold the full potential in terms of mechanistic interpretation and prediction of bubble dynamics.

Data availability

The code and demonstration of the model predictions can be found in the repository UTILE-Oxy at <https://github.com/andyco98/UTILE-Oxy>. The employed annotated dataset and the trained model are accessible at Zenodo: <https://doi.org/10.5281/zenodo.10184579>.

Author contributions

André Colliard-Granero: Conceptualization, methodology, software, visualization, writing – original draft, writing – review & editing. Keusra A. Gompou: conceptualization, methodology, software, visualization, writing – original draft, writing – review & editing. Christian Rodenbücher: visualization, writing – original draft, writing – review & editing. Michael H. Eikerling: conceptualization, supervision, writing – original draft, writing – review & editing. Kourosh Malek: conceptualization, supervision, writing – original draft, writing – review & editing. Mohammad J. Eslamibidgoli: conceptualization, methodology, supervision, visualization, writing – original draft, writing – review & editing.

Conflicts of interest

The authors declare no competing financial interest.

Acknowledgements

The authors acknowledge the financial support from the Helmholtz Imaging Platform HIP (a platform of the Helmholtz Incubator on Information and Data Science) grant number DB002249, as well as the Federal Ministry of Science and Education (BMBF) under the German-Canadian Materials Acceleration Centre (GC-MAC) grant number 01DM21001A. The authors also acknowledge the Gauss Centre for Supercomputing



e.V. (<https://www.gauss-centre.eu>) for funding this project by providing computing time on the GCS Supercomputer JUWELS³⁰ at Jülich Supercomputing Centre (JSC) of Forschungszentrum Jülich. The authors thank K. Klafki for technical support and C. Korte for fruitful discussions.

Notes and references

- 1 Y. Guo, G. Li, J. Zhou and Y. Liu, *IOP Conference Series: Earth and Environmental Science*, 2019, p. 042022.
- 2 I. Sinapan, C. Lin-Kwong-Chon, C. Damour, J.-J. A. Kadjo and M. Benne, *Hydrogen*, 2023, **4**, 556–572.
- 3 J. K. Lee and A. Bazylak, *Joule*, 2021, **5**, 19–21.
- 4 J. O. Majasan, J. I. Cho, I. Dedigama, D. Tsaoulidis, P. Shearing and D. J. Brett, *Int. J. Hydrogen Energy*, 2018, **43**, 15659–15672.
- 5 I. Dedigama, P. Angeli, N. van Dijk, J. Millichamp, D. Tsaoulidis, P. R. Shearing and D. J. Brett, *J. Power Sources*, 2014, **265**, 97–103.
- 6 F. Aubras, J. Deseure, J.-J. Kadjo, I. Dedigama, J. Majasan, B. Grondin-Perez, J.-P. Chabriot and D. Brett, *Int. J. Hydrogen Energy*, 2017, **42**, 26203–26216.
- 7 X. Su, L. Xu and B. Hu, *Int. J. Hydrogen Energy*, 2022, **47**, 20027–20039.
- 8 M. Maier, Q. Meyer, J. Majasan, R. E. Owen, J. B. Robinson, J. Dodwell, Y. Wu, L. Castanheira, G. Hinds and P. R. Shearing, *et al.*, *Front. Energy Res.*, 2020, **8**, 582919.
- 9 A. Pandit, J. Varley, R. Thorpe and J. Davidson, *Chem. Eng. Sci.*, 1992, **47**, 1079–1089.
- 10 G. G. Yen and H. Lu, *ISA Trans.*, 2002, **41**, 273–282.
- 11 I. Dedigama, D. J. Brett, T. J. Mason, J. Millichamp, P. R. Shearing and K. E. Ayers, *ECS Trans.*, 2015, **68**, 117.
- 12 Y. Zhang, J. Zhang, Y. Zhao, X. Zhang, X. Yang, E. Zhou, C. Duan, G. Wang and L. Dong, *Chem. Eng. J.*, 2020, **395**, 125129.
- 13 G. Besagni, P. Brazzale, A. Fiocca and F. Inzoli, *Flow Meas. Instrum.*, 2016, **52**, 190–207.
- 14 M. J. Eslamibidgoli, F. P. Tipp, J. Jitsev, J. Jankovic, M. H. Eikerling and K. Malek, *RSC Adv.*, 2021, **11**, 32126–32134.
- 15 A. Colliard-Granero, J. Jitsev, M. H. Eikerling, K. Malek and M. J. Eslamibidgoli, *ACS Nanosci. Au*, 2023, 398–407.
- 16 A. Colliard-Granero, M. Batool, J. Jankovic, J. Jitsev, M. H. Eikerling, K. Malek and M. J. Eslamibidgoli, *Nanoscale*, 2022, **14**, 10–18.
- 17 A. Malek, M. J. Eslamibidgoli, M. Mokhtari, Q. Wang, M. H. Eikerling and K. Malek, *ChemPhysChem*, 2019, **20**, 2946–2955.
- 18 S. Sun, F. Xu, L. Cai, D. Salvato, F. Dilemma, L. Capriotti, M. Xian and T. Yao, *An Efficient Instance Segmentation Approach for Extracting Fission Gas Bubbles on U-10Zr Annular Fuel*, 2023.
- 19 C. M. Anderson, J. Klein, H. Rajakumar, C. D. Judge and L. K. Beland, *Ultramicroscopy*, 2020, **217**, 113068.
- 20 Y. Kim and H. Park, *Sci. Rep.*, 2021, **11**, 8940.
- 21 M. Tkachenko, M. Malyuk, A. Holmanyuk and N. Liubimov, *Label Studio: Data labeling software*, 2020–2022, <https://github.com/heartexlabs/label-studio>, Open source software available from <https://github.com/heartexlabs/label-studio>.
- 22 O. Ronneberger, P. Fischer and T. Brox, *U-Net: Convolutional Networks for Biomedical Image Segmentation*, 2015.
- 23 S. Xie, R. Girshick, P. Dollár, Z. Tu and K. He, *Aggregated Residual Transformations for Deep Neural Networks*, 2017.
- 24 O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, B. Glocker and D. Rueckert, *Attention U-Net: Learning Where to Look for the Pancreas*, 2018.
- 25 G. Bradski, *Dr Dobb's Journal: Software Tools for the Professional Programmer*, 2000, **25**, 120–123.
- 26 L. von Chamier, R. F. Laine, J. Jukkala, C. Spahn, D. Krentzel, E. Nehme, M. Lerche, S. Hernández-Pérez, P. K. Mattila and E. Karinou, *et al.*, *Nat. Commun.*, 2021, **12**, 2276.
- 27 T. Yan, Y. Liu, D. Wei, X. Sun and Q. Liu, *Environ. Sci. Pollut. Res.*, 2023, **30**, 62803–62814.
- 28 M. Hentschel and N. Page, *Part. Part. Syst. Charact.*, 2003, **20**, 25–38.
- 29 Y. Yu, P. Yang, H. Meng, H. Liu, J. Zhang and H. Yu, *Ind. Eng. Chem. Res.*, 2022, **61**, 18574–18587.
- 30 D. Krause, *J. Large-Scale Res. Facilities*, 2019, **5**, A135–A135.

