

PAPER • OPEN ACCESS

Neural network-based surrogate model for 3D edge-plasma transport in the standard configuration of W7-X

To cite this article: Y. Luo *et al* 2026 *Nucl. Fusion* **66** 016038

View the [article online](#) for updates and enhancements.

You may also like

- [Core plasma fueling by fast inward particle transport after hydrogen pellet injection in Wendelstein 7-X](#)
H. Damm, J. Baldzuhn, R.C. Wolf *et al.*
- [Feedforward equilibrium trajectory optimization with GSPulse](#)
J.T. Wai, M.D. Boyer, D.J. Battaglia *et al.*
- [On the possibility of reducing the size of a fusion reactor by increasing the plasma density](#)
C. Angioni, E. Fable and H. Zohm



HIDEN
ANALYTICAL
*Trusted in Research
for over 40 years*

www.HidenAnalytical.com

Ultra-High Resolution Fusion Gas Analysis for H/He isotopes, light gases, and complex vapour mixtures

DLS Series <ul style="list-style-type: none">• Real-time ultra-high resolution• ppm-level isotope sensitivity• Built for fusion environments• Dual-zone operation• Remote mounting capability	HAL 101X <ul style="list-style-type: none">• For tokamak and torus gas analysis• No radiation shielding required• TIMS mode for real-time H/He isotope quantification
--	--

Find Solutions for Your Research

Neural network-based surrogate model for 3D edge-plasma transport in the standard configuration of W7-X

Y. Luo^{1,2} , S. Xu^{1,*} , Y. Liang^{1,2,*} , E. Wang¹ , J. Cai¹, A. Knieps¹ , R. Pei^{1,2}, M. Tan^{1,2}, Y. Feng³ , D. Reiter², S. Brezinsek^{1,2} , D. Harting¹, M. Krychowiak³ , D. Gradic³ , M. Jakubowski³  and the W7-X Team^a

¹ Forschungszentrum Jülich GmbH, Institute of Fusion Energy and Nuclear Waste Management—Plasma Physics, Partner of the Trilateral Euregio Cluster (TEC), 52425 Jülich, Germany

² Faculty of Mathematics and Natural Science, Heinrich Heine University Düsseldorf, 40225 Düsseldorf, Germany

³ Max Planck Institute for Plasma Physics, 17491 Greifswald, Germany

E-mail: s.xu@fz-juelich.de and y.liang@fz-juelich.de

Received 25 June 2025, revised 15 October 2025

Accepted for publication 17 November 2025

Published 27 November 2025



Abstract

This paper presents a neural-network surrogate model for Wendelstein 7-X (W7-X) edge transport simulations, trained on an EMC3-EIRENE dataset that spans nearly the entire operating-parameter space currently explored in the W7-X standard configuration. The model uses an autoencoder to compress EMC3-EIRENE outputs into a low-dimensional latent space representation, then a neural regressor maps EMC3-EIRENE input parameters to those latent vectors, and finally both components are fine-tuned together to predict outputs directly from inputs. In benchmark tests, the improved surrogate outperforms a traditional multilayer perceptron model, most notably in predicting detached regime. Leave-one-value-out evaluation indicates high accuracy within the interpolation domain, with minor degradation when extrapolating. Relative to full EMC3-EIRENE runs, the surrogate provides over a 10^8 - fold speedup, enabling large-scale parameter scans or real-time feedback control based on 3D transport simulations to become feasible in the future.

Keywords: machine learning, EMC3-EIRENE, edge plasma transport, Wendelstein 7-X

(Some figures may appear in colour only in the online journal)

^a See Grulke *et al* 2024 (<https://doi.org/10.1088/1741-4326/ad2f4d>) for the W7-X Team.

* Authors to whom any correspondence should be addressed.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1. Introduction

The Wendelstein 7-X (W7-X) [1–3] is an advanced stellarator optimized for neoclassical transport, fast particle confinement and magnetohydrodynamic stability, whose intrinsic three-dimensional (3D) magnetic field structure imposes high demands on simulation. The 3D edge transport code EMC3-EIRENE [4], which couples the fluid model EMC3 [5] with the kinetic neutral particle transport code EIRENE [6], is well suited for handling such geometries and its application has substantially deepened our understanding of edge plasma transport, radiation dissipation, divertor heat load, particle exhaust and so on [7–18]. However, the complexity of the EMC3-EIRENE simulation leads to significant CPU resource demands, limiting its application in many scenarios, such as guiding ongoing experiments directly from the control room. For example, even when applying stellarator symmetry to reduce the toroidal range of the entire simulation space to one-tenth of the complete edge range, a single simulation on W7-X meeting satisfactory grid resolutions typically demands several thousand core-hours on a supercomputer. In recent years, although simplified codes such as the diffusive field line tracing [19] and EMC3-Lite [20] improved computational efficiency, their neglect of impurity transport and particle-neutral interactions makes them fall far short of replacing EMC3-EIRENE, particularly in high-density plasma simulations including detached plasmas, which represent a potential scenario in future nuclear fusion reactors.

To address this problem, machine learning technology provides a promising solution by learning input-output relationships from a precomputed dataset, enabling rapid output predictions. In fusion research, machine learning has been widely applied to tasks such as plasma control [21–23], disruption predictions [24–26], and simulation accelerations [27–31]. For simulation accelerations, one strategy is to replace full solvers with direct input-output mappings, as shown by one-dimensional codes DIVID [32, 33], and the two-dimensional code SOLPS-ITER [29, 34, 35]. Another approach focuses on accelerating specific time-consuming steps of a workflow [28, 36, 37]. As is well known, if the upfront cost of constructing a simulation dataset is acceptable and cost-effective, machine learning methods that can rapidly and accurately provide predictive results will offer significant value.

Here, neural-network surrogates for the computationally intensive 3D EMC3-EIRENE code are successfully developed for the standard configuration of W7-X. Due to the complex island divertor configuration of W7-X [38, 39], a fully connected multilayer perceptron (MLP) we initially trained on converged EMC3-EIRENE simulations, cannot adequately capture the nonlinear input-output relationship, resulting in non-negligible errors when predicting detached-plasma cases. Therefore, an improved architecture is proposed: an autoencoder first compresses EMC3-EIRENE outputs into latent space, an MLP then maps the inputs into these latent vectors, and the two networks are finally fine-tuned together. The resulting surrogate achieves markedly higher accuracy than

the MLP, with relative errors compared to EMC3-EIRENE simulations remaining at satisfactory levels. Meanwhile, the improved model can provide results extremely quickly, making it feasible to perform extensive parameter scans or real-time feedback control based on 3D transport simulations in the future.

The remainder of this paper is organized as follows. Dataset construction, neural network architecture, and the hyperparameter optimization are detailed in section 2. Section 3 evaluates the fine-tuned surrogate’s accuracy, its interpolation and extrapolation performance, and its computational efficiency. Finally, section 4 concludes the paper.

2. Data and methods

2.1. Dataset generation

A comprehensive dataset was generated using EMC3-EIRENE code. In these simulations, carbon sputtered from divertor targets or baffles was specified as the sole radiation source, with its source strength controlled by the total radiated power. Cross-field particle diffusivities for hydrogen (bulk plasma) and carbon were assumed equal and spatially uniform, as were the electron and ion thermal diffusivities. Five primary input parameters were systematically scanned in the EMC3-EIRENE simulations under the standard configuration of W7-X: input heating power across the innermost boundary surface (P_{in}), radiation power fraction (the ratio of radiation power to total heating power, $f_{rad} = \frac{P_{rad}}{P_{in}}$), electron density at the separatrix ($n_{e,sep}$), cross-field particle diffusivity D_{\perp} and cross-field thermal diffusivity χ_{\perp} . The ranges of these parameters, as shown in table 1, were selected to cover present operational window in W7-X standard configuration. As we know, at constant heating power, increasing the plasma density drives the edge plasma from low to high recycling and eventually into detachment. Figure 1 illustrates this trend in EMC3-EIRENE by tracking the total recycling flux during a density ramp. To ensure the surrogate captures the roll-over of recycling flux, the separatrix electron density was sampled more densely in the range $3\text{--}6 \times 10^{19}\text{m}^{-3}$. Conducting EMC3-EIRENE simulations with exhaustive sampling across the entire parameter space defined in table 1 would be prohibitively expensive computationally and unnecessary in practice. We found that, beyond a certain dataset size, adding further samples yielded only marginal improvements in model performance. Thus, 463 input sets (about 8% of the full parameter space) were chosen at random from this parameter space and each was run to numerical convergence with EMC3-EIRENE. Training the neural network model on this converged simulation dataset ensures that physical relationships and self-consistency are effectively preserved throughout the entire 3D edge region. The suitability of this sample size for neural-network training is assessed later.

EMC3-EIRENE produces a wide array of quantities, including electron temperature, electron density, ion




Figure 1. Total recycling flux versus separatrix electron density for simulations with $P_{in} = 10$ MW, $D_{\perp} = 0.5$ m²s⁻¹, $\chi_{\perp} = 1.5$ m²s⁻¹, and a carbon sputtering rate of 0.04.

Table 1. Input parameter ranges used to generate the EMC3-EIRENE dataset.

Parameter	Range									
P_{in} (MW)	4		6		8		10			
f_{rad}	0.2		0.4		0.6		0.8			
$n_{e,sep}$ (10^{19} m ⁻³)	2	3	3.5	4	4.5	5	5.5	6	7	
D_{\perp} (m ² /s)	0.1	0.3	0.5	0.7	0.9	1.1	1.5	2		
χ_{\perp} (m ² /s)	0.2		0.6		0.9		1.2			1.5

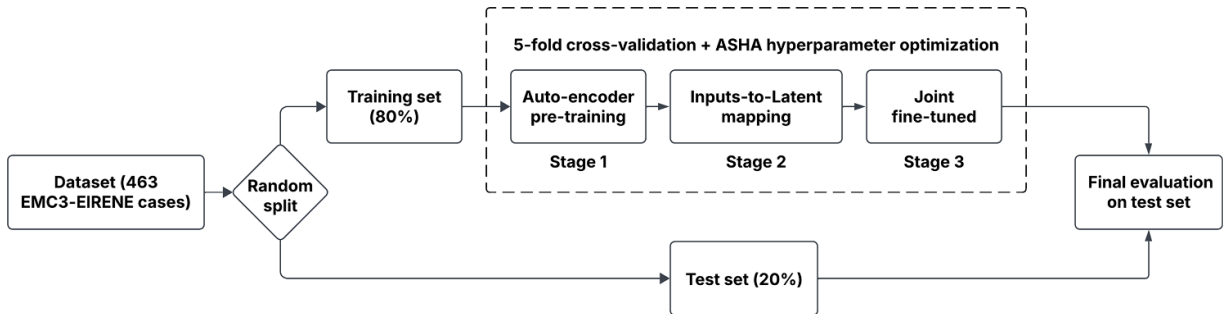


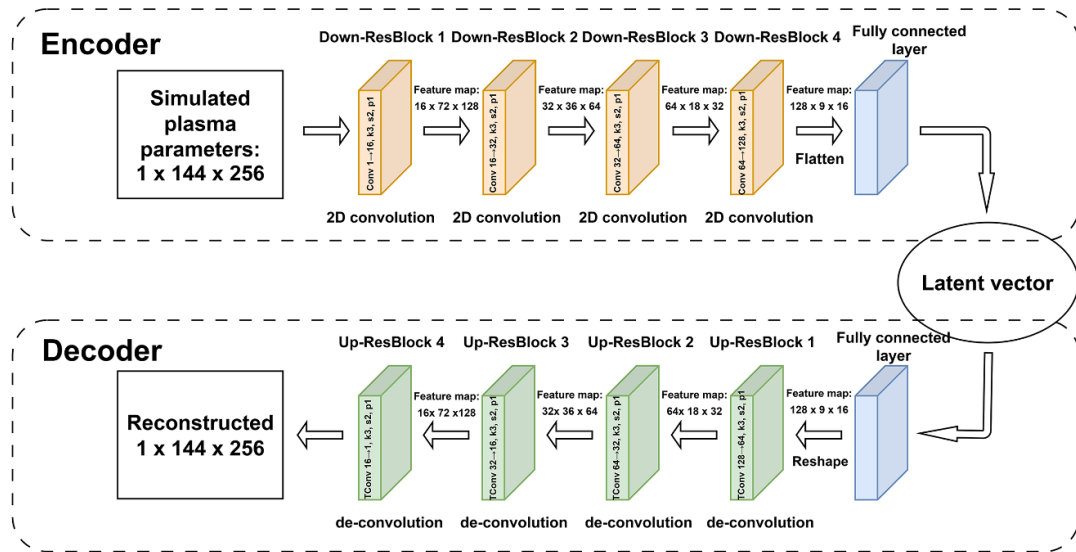
Figure 2. Training workflow for the fine-tuned surrogate with 5-fold cross-validation and ASHA early-stopping strategy.

temperature, divertor heat load and so on. In the present study, a separate neural network is trained for each individual output. The EMC3-EIRENE simulation dataset was constructed using a 3D numerical mesh with a resolution of $144 \times 256 \times 144$ cells in the radial, poloidal, and toroidal directions, respectively. To facilitate interpretation, a toroidal slice was extracted, producing a 2D 144×256 (radial \times poloidal) cross-section at the bean-shaped location on W7-X, which serves as the representative output discussed below. All inputs and outputs were scaled with independent min-max normalization to promote stable and rapid training. For every case, each of the five input parameters and every grid cell of outputs was transformed by subtracting its global minimum (over the 463 simulations) and dividing by the corresponding range (maximum minus minimum).

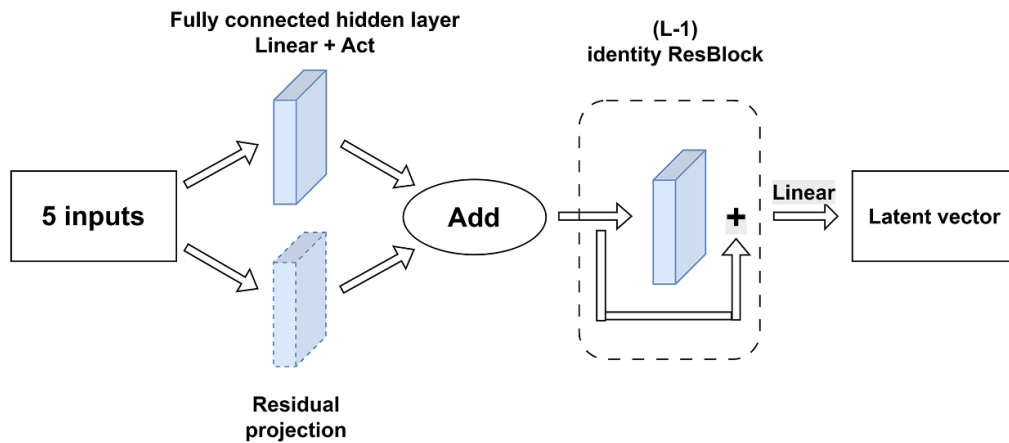
2.2. Training workflow

Figure 2 display the training workflow of the fine-tuned surrogate. After min-max normalization, the 463 converged EMC3-EIRENE cases are randomly divided into a training set (80%) and an independent test set (20%). The training set is used in all subsequent network-training stages, while the test subset is reserved for the final performance assessment. The surrogate network training includes three stages: (1) autoencoder pre-training, (2) inputs-to-latent mapping, and (3) joint fine-tuned of the combined networks. Hyperparameters are selected by searching the training set with five-fold cross-validation and the ASHA early-stopping scheduler, and the optimal hyperparameter settings were then applied to complete the training. The next subsections describe

Stage 1: Auto-encoder



Stage 2: Input-to-latent mapping



Stage 3: Joint fine-tuned




Figure 3. Architecture of the three-stage (a) autoencoder pre-training, (b) input-to-latent mapping, (c) joint fine-tuned.

the network architecture and the hyperparameter optimization in detail.

2.2.1. Neural network architecture. The neural network architecture for each phase is illustrated in figure 3. The surrogate aims to efficiently map five input parameters of EMC3-EIRENE to the corresponding simulation outputs.

Stage 1. Autoencoder pre-training

The autoencoder consists of an encoder-decoder pair. The encoder compresses high-dimensional simulation outputs into a low-dimensional latent representation, whereas the decoder reconstructs these outputs from a compact latent representation. The aim is to learn an information-rich latent space that captures the dominant spatial structure of the outputs. With the learned representation, Stage 2 predicts a compact latent vector from the five inputs instead of full outputs, simplifying the




Figure 4. Schematic of the reference MLP used for direct input-output mapping.

complex mapping task. The loss function used for autoencoder pre-training is the mean squared error (MSE).

In this work, the encoder contains four down-sampling residual blocks. Each block applies a 3×3 convolution with stride 2 and padding 1, followed by batch normalization and a two-layer residual refinement before the shortcut is added. After the four blocks, the spatial resolution is reduced from $144 \times 256 \rightarrow 9 \times 16$, while the channel width increases from 1 to 128. The resulting $128 \times 9 \times 16$ tensor is flattened and transformed into a compact latent vector by a fully-connected layer, whose size is optimized through hyperparameter tuning. This latent vector significantly compresses the data dimensions while preserving the essential spatial structure.

The decoder can be regarded as the inverse transformation of the encoder. It first reshapes the latent vector back to $128 \times 9 \times 16$ with a fully connected layer, then uses four up-sampling residual blocks. In each block a transposed convolution doubles the spatial dimensions and decreases the channel width, restoring the resolution to 144×256 while reducing the channels from 128 to 1.

Stage 2. Input-to-latent mapping

The mapping network is a fully connected feed-forward neural network translating the five EMC3-EIRENE input parameters into the latent vector produced by the encoder. The encoder is frozen during this stage. A residual branch projects the original 5 inputs directly to the first hidden representation and is added to the activated output, retaining the linear part of the mapping. The subsequent hidden layers use identity skip connections, which help gradients flow and stabilize training. A final linear layer produces the latent vector. Key hyperparameters such as the number of hidden layers L , the number of neurons in each layer, and the choice of activation function are selected through the hyperparameter search described below. The loss function is a latent space MSE between the mapping output and the encoder-derived latent vector.

Stage 3. Joint fine-tuned

In the final stage, the inputs-to-latent mapping network and the decoder are chained to form an end-to-end surrogate that

predicts the EMC3-EIRENE outputs from the 5 input parameters. Specifically, the mapping network generates latent vectors based on the 5 EMC3-EIRENE inputs, which subsequently serve as inputs to the decoder. The decoder then reconstructs the simulation outputs through its up-sampling structure. Training minimizes the MSE between the surrogate prediction and the reference simulation, updating both the mapping network and the decoder.

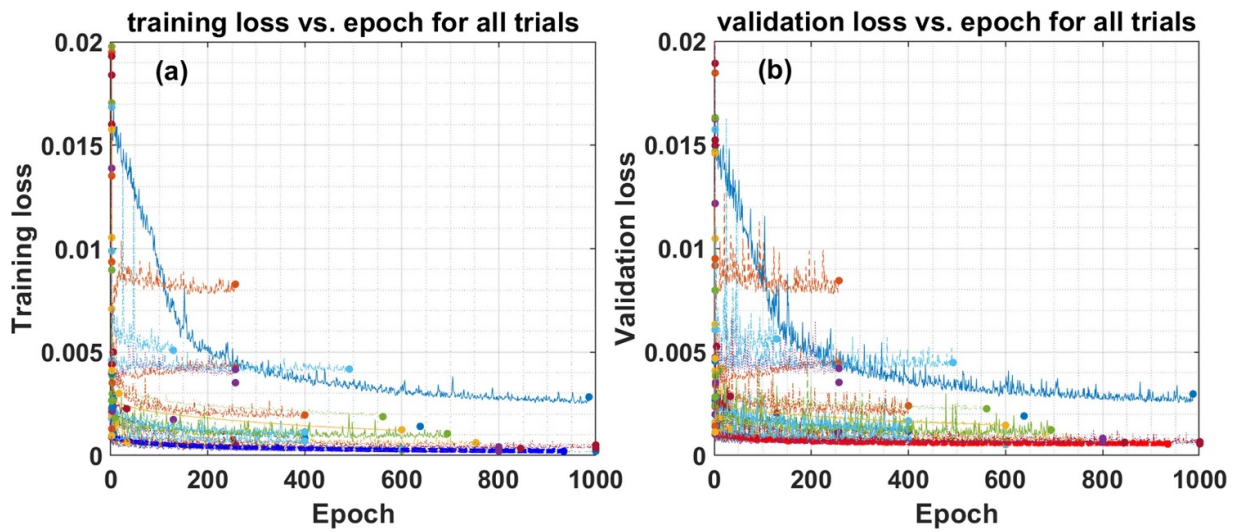
For comparison, we replace the three stages with a single MLP and train it using the same procedure. The architecture of this reference MLP, illustrated in figure 4, comprises an input layer, hidden layers, and an output layer. Every neuron in a layer connects to every neuron in the next, and nonlinear activation functions supply the required nonlinearity. The MLP's hyperparameters are also determined through the same optimization process.

2.2.2. Hyperparameter optimization. This section illustrates the hyperparameter optimization strategy using plasma temperature prediction as an example. To maximize data utilization and minimize sampling errors, we employ 5-fold cross-validation. Specifically, the training dataset is divided into five equal subsets, with each four subsets used for training and the remaining subset reserved for validation. Table 2 lists the ranges explored for the main hyperparameters. A total of 100 random hyperparameter sets were selected from this search space, and each set underwent 5-fold cross-validation to calculate the average training and validation loss. The set with the lowest average validation loss was selected as optimal. To efficiently carry out hyperparameter optimization and reduce computational costs associated with numerous trials, the asynchronous successive halving algorithm (ASHA) early stopping strategy [40] was employed.

Figure 5 shows the average training and validation losses from the 5-fold cross-validation across these 100 hyperparameter sets. In panel (a) the dark blue curve traces the average training loss for the best set, while in panel (b) the red curve traces the corresponding average validation loss. Both panels illustrate that the early-stopping strategy clearly terminate

Table 2. Search space of the main hyperparameters.

Hyperparameters	Search space
Learning rate	0.0001–0.01
Number of hidden layers	2, 3, 4
Number of units per hidden layer	64, 128, 256
Latent vector dimension	32, 64, 128, 256
Batch size	8, 16, 32
L2 regularization	10^{-7} – 10^{-3}
Optimizer	Adam
Activation function	ELU, GELU, LeakyReLU
Scheduler	ASHA
Auto-encoder epoch	100, 200, 300, 400, 500, 600
Mapping epoch	1000, 1500, 2000, 2500, 3000
Fine-tuned epoch	400, 600, 800, 1000

**Figure 5.** Average training loss (a) and validation loss (b) over 100 hyperparameter trials using 5-fold cross-validation.

poorly performing trials at an early stage and reallocates computational resources to more promising sets. With the optimal hyperparameter settings, the average training and validation losses during the joint fine-tuned stage converge to approximately 2.0×10^{-4} and 5.6×10^{-4} , respectively. These low converged losses indicate that the surrogate model has accurately captured the nonlinear mapping between EMC3-EIRENE input parameters and outputs.

Figure 6 summarizes the optimal hyperparameters. Across all training stages (autoencoder pre-training, mapping network training, and joint fine-tuned), the Adam optimizer was employed with a learning rate of 6.2×10^{-5} , batch size of 8, ELU activation functions [41], and L2 regularization [42] of 3.7×10^{-7} . The autoencoder compresses simulation outputs into a 256-dimensional latent representation. The mapping network consists of 3 hidden layers, each containing 256 neurons. Training involved 100 epochs for autoencoder pre-training, 2500 epochs for the mapping network, and 1000 epochs for the joint fine-tuned stage. In practice, due to the ASHA early stopping strategy, the actual number of epochs

may be fewer than initially configured. Following this hyperparameter optimization, the final model was trained on the complete training dataset using these optimal settings and applied the same settings to all subsequent model training. Figure 7 shows the hyperparameter importance and their correlations with the validation loss. The importance metric [43] quantifies how each hyperparameter impacts the validation loss. Higher importance values means stronger impact. The correlation indicates the linear relationship between the value of hyperparameters and the validation loss. A positive correlation (green) indicates that as a hyperparameter's value increases, the validation loss also tends to increase. Conversely, lower values tend to yield smaller loss. As shown, the learning rate, batch size, and latent vector dimension are the most significant parameters. All of them demonstrate positive correlations with validation loss, indicating that within the explored parameter space, lower these three parameters generally enhance model performance.

The reference MLP model underwent a similar hyperparameter optimization procedure, resulting in the following

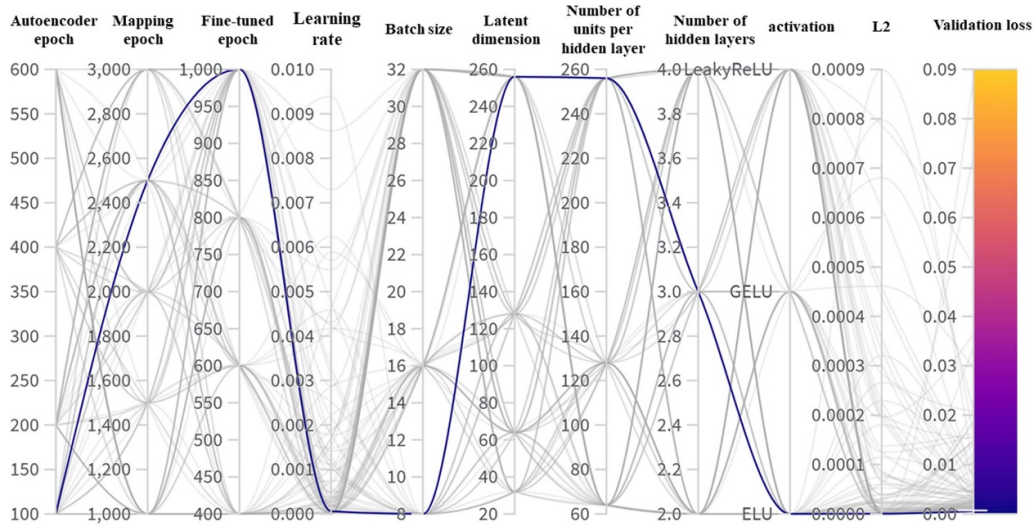


Figure 6. Parallel coordinates plot of hyperparameter settings for 100 trials. Each grey line represents a trial's hyperparameter combination and its corresponding validation loss, while the blue line highlights the optimal trial.

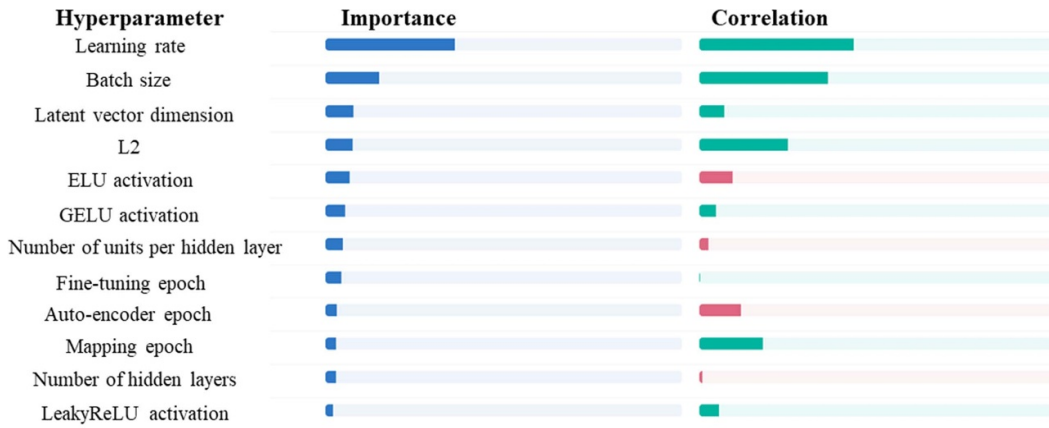


Figure 7. Hyperparameter importance and their correlations with validation loss. Left: importance quantifying each hyperparameter's contribution to the variability of the validation loss. Right: correlation with the validation loss.

optimal configuration: 200 epochs, three hidden layers with 256 neurons each, a learning rate of 9.6×10^{-4} , batch size of 16, ELU activation function, and L2 regularization of 2.3×10^{-7} , also employing the Adam optimizer. Compared to the joint fine-tuned surrogate, the MLP model exhibited higher training and validation losses, approximately 9.1×10^{-4} and 1.3×10^{-3} , respectively.

To ensure the dataset provides sufficient information for accurately learning the relationship between EMC3-EIRENE inputs and outputs, an assessment of dataset adequacy was performed. We varied the fraction of data allocated to training from 0.1 to 0.9, using the remaining cases as the test set. Throughout, we maintained the same network architecture and hyperparameters and recorded the resulting test loss. As shown in figure 8, the test loss decreases as the trainset fraction increases, then remains nearly unchanged once the fraction exceeds around 0.7. This behavior indicates that the current dataset is already dense enough for accurate surrogate training, and adding more data would produce limited performance gains. This result also justifies using $k = 5$ for k -fold

cross-validation in hyperparameter tuning mentioned above. Each fold effectively trains on approximately 64% of the full dataset, close to the performance plateau. Therefore, $k = 5$ is deemed sufficient and increasing k further would mainly increase runtime with a negligible impact on hyperparameter selection.

3. Performance evaluation of the fine-tuned method

3.1. Accuracy assessment

The overall performance of the MLP and joint fine-tuned surrogates is evaluated using the mean absolute error (MAE) and mean squared error (MSE) for the three main outputs (T_e , n_e and P_{rad}). For n_e and P_{rad} , the MLP and joint fine-tuned surrogates were trained using the same network architecture and hyperparameter settings as for T_e . For each test case, we first average the per-cell errors over all grid cells to obtain a single MAE and a single MSE. Across the test set, we then compute the mean and standard deviation of these per-case metrics, as

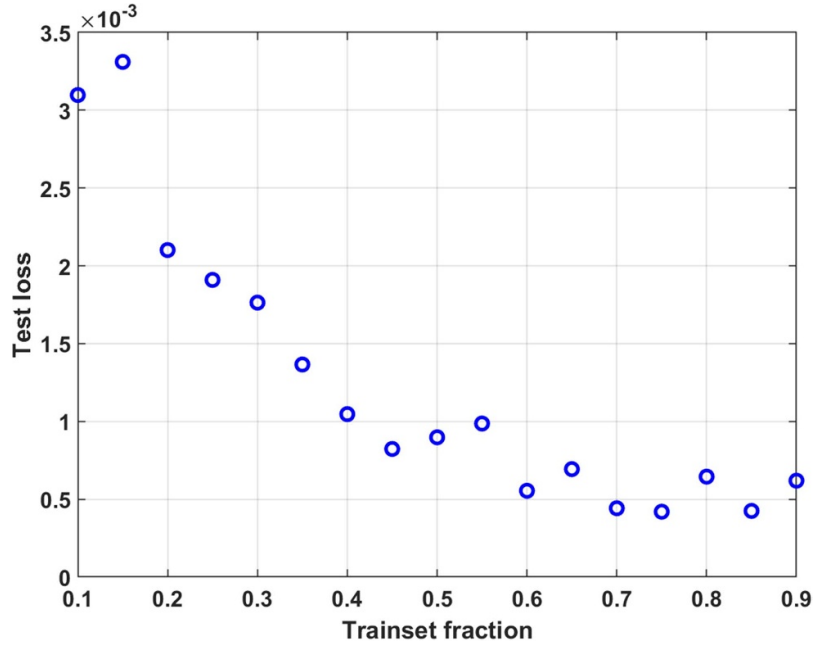


Figure 8. Test loss as the function of the trainset fraction. Each point corresponds to training on a given fraction of the full dataset, with the remaining cases held out as the test set.

Table 3. Quantitative error assessment on the test set: MAE and MSE.

Output	Surrogate model	MAE		MSE	
		mean	standard deviation	mean	standard deviation
T_e (eV)	MLP	3.56×10^{00}	2.42×10^{00}	6.36×10^{01}	1.07×10^{02}
	Joint fine-tuned	1.42×10^{00}	1.51×10^{00}	1.49×10^{01}	3.85×10^{01}
n_e (m^{-3})	MLP	1.90×10^{00}	1.32×10^{18}	3.37×10^{37}	1.12×10^{38}
	Joint fine-tuned	1.12×10^{18}	1.10×10^{18}	1.58×10^{37}	4.67×10^{37}
P_{rad} (W cm^{-3})	MLP	9.62×10^{-03}	8.48×10^{-03}	6.87×10^{-03}	1.46×10^{-02}
	Joint fine-tuned	1.04×10^{-03}	2.49×10^{-03}	1.01×10^{-03}	4.76×10^{-03}




Figure 9. Probability density distributions of T_e , n_e and P_{rad} relative error: MLP (blue), joint fine-tuned surrogate (red).

presented in table 3. Compared to the MLP, the joint fine-tuned surrogate achieves substantially lower MAE and MSE for T_e , n_e and P_{rad} , indicating that adopting the latter leads to significantly more accurate results. The smaller standard deviations further support the conclusion of more consistent performance across the test set.

Figure 9 complements table 3 by plotting relative error distributions over all grid cells from all cases in the test set. For

each output, the red curve representing the joint fine-tuned surrogate is narrower, implying a smaller proportion of large errors. The share of grid cells with absolute value of relative error exceeding 10% drops from 13.3% to 7.3% for T_e , from 8.3% to 4.9% for n_e , and from 5.3% to 2.9% for P_{rad} . The thinner tails of the red curves confirm the MAE/MSE reductions reported in table 3. Overall, on the test set, the joint fine-tuned surrogate outperforms the MLP.

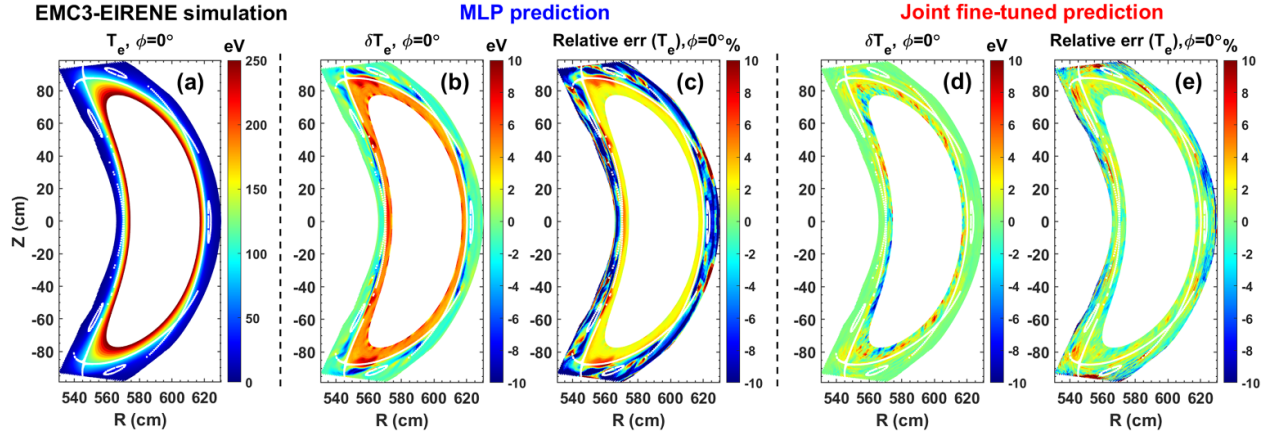


Figure 10. Comparison of electron temperature predictions for a low radiation fraction case. (a) EMC3-EIRENE reference, (b), (c) MLP raw and relative errors, (d), (e) joint fine-tuned surrogate raw and relative errors.




Figure 11. Comparison of electron temperature predictions for a high radiation fraction case. (a) EMC3-EIRENE reference, (b), (c) MLP raw and relative errors, (d), (e) joint fine-tuned surrogate raw and relative errors.

To identify where large errors concentrate, we compare both surrogates with EMC3-EIRENE for two representative test cases, one with low radiation fraction ($f_{rad} = 0.2$) and one with high radiation fraction ($f_{rad} = 0.8$), examining predictions for T_e , n_e and P_{rad} . Both cases are run with an input power of $P_{in} = 6\text{MW}$, cross-field particle diffusivity $D_{\perp} = 1.1\text{m}^2\text{s}^{-1}$, and cross-field thermal diffusivity $\chi_{\perp} = 1.5\text{m}^2\text{s}^{-1}$. The separatrix electron densities are $n_{e,sep} = 3.0 \times 10^{19}\text{m}^{-3}$ for the low radiation case and $5.5 \times 10^{19}\text{m}^{-3}$ for the high-radiation case.

Figure 10 compares T_e predictions from the two surrogate models with the EMC3-EIRENE simulation reference at the bean-shaped cross-section for the low radiation case. Panel (a) shows the EMC3-EIRENE electron-temperature distribution. Panels (b) and (c) give the MLP's raw error ($\delta = y_{pred} - y_{true}$) and relative error ($r = \frac{y_{pred} - y_{true}}{y_{true}}$), while panels (d) and (e) show the same errors for the joint fine-tuned surrogate, all measured against the EMC3-EIRENE reference. In each figure, the Poincaré plot is not derived from the neural-network output but rather refers to the visualization of the underlying magnetic equilibrium of the W7-X standard configuration on which the EMC3-EIRENE training data are based. It

is readily apparent that the joint surrogate achieves significantly higher accuracy. Its relative error remains low across most of the domain and rises only in very low-temperature areas, where the raw deviation stays modest at roughly 2–4 eV. Figure 11 presents the corresponding results for the high radiation fraction case. Accuracy decreases for both surrogates, and the MLP's relative error exceeds 10% in the scrape-off layer (SOL), making it unreliable for prediction. The joint surrogate's relative error approaches 10% only in a few low-temperature regions, where the raw error remains small at roughly 1 eV, while relative errors elsewhere stay modest.

For n_e predictions, the MLP shows large relative errors in the SOL in both the low and high radiation cases, with errors further amplified in the high radiation case, as shown in figures 12 and 13. The joint fine-tuned surrogate markedly reduces these SOL errors. In the low radiation case, only a very narrow edge region exhibits noticeable relative error, but it remains below 10%. In the high radiation case, overall performance degrades somewhat, with relative errors tending to be higher in high-density zones while remaining modest elsewhere.




Figure 12. Comparison of electron density predictions for a low radiation fraction case. (a) EMC3-EIRENE reference, (b), (c) MLP raw and relative errors, (d), (e) joint fine-tuned surrogate raw and relative errors.




Figure 13. Comparison of electron density predictions for a high radiation fraction case. (a) EMC3-EIRENE reference, (b), (c) MLP raw and relative errors, (d), (e) joint fine-tuned surrogate raw and relative errors.

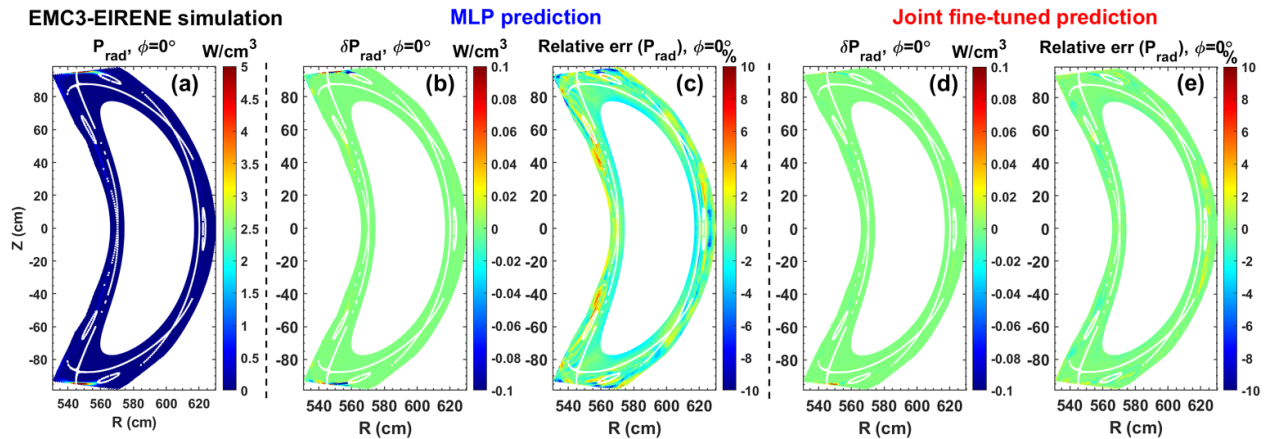


Figure 14. Comparison of radiation power predictions for a low radiation fraction case. (a) EMC3-EIRENE reference, (b), (c) MLP raw and relative errors, (d), (e) joint fine-tuned surrogate raw and relative errors.

For P_{rad} predictions shown in figures 14 and 15, the joint fine-tuned surrogate also demonstrates better performance. In the low radiation fraction case, radiation is concentrated near the strike point. The MLP shows visible errors there, whereas the joint fine-tuned surrogate is more accurate.

In the high radiation fraction case, the radiation front moves toward the plasma core. The MLP continues to struggle near the front, while the joint fine-tuned surrogate remains accurate, with only a modest increase in relative error around the front and small errors elsewhere.




Figure 15. Comparison of radiation power predictions for a high radiation fraction case. (a) EMC3-EIRENE reference, (b), (c) MLP raw and relative errors, (d), (e) joint fine-tuned surrogate raw and relative errors.

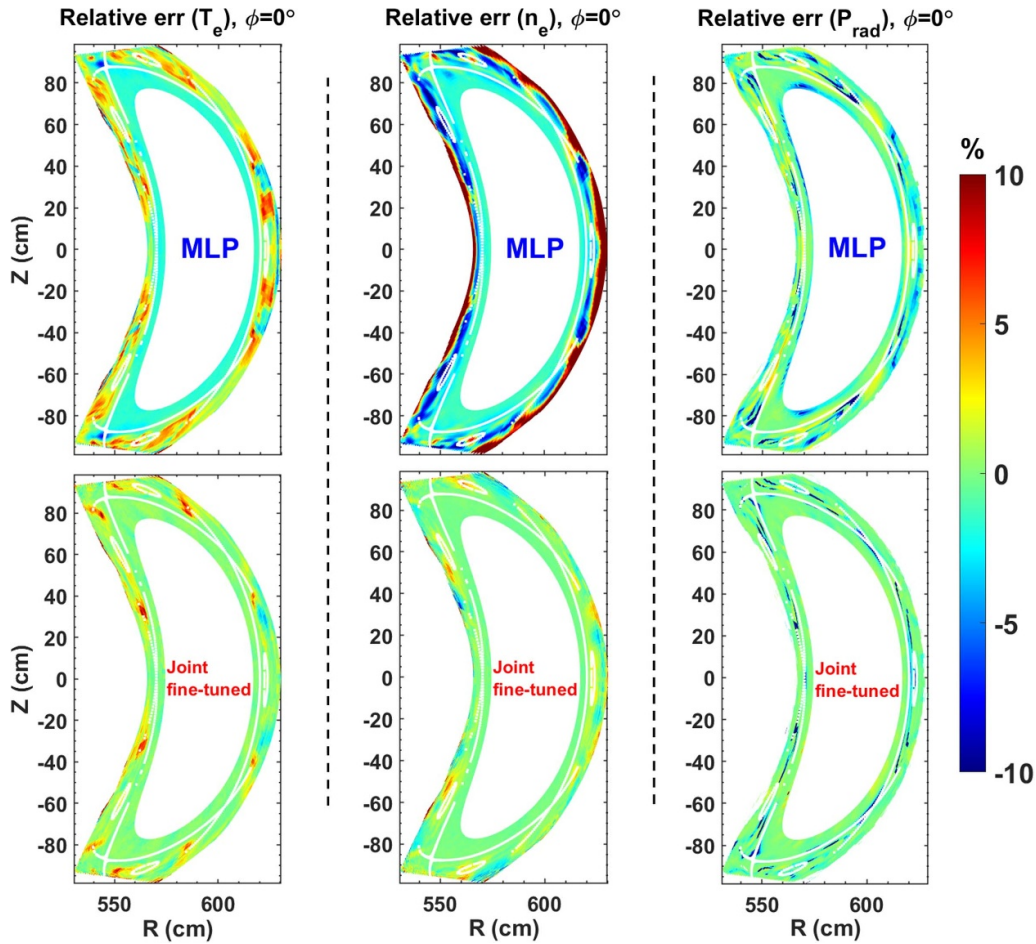


Figure 16. Spatial maps of relative error averaged over the test set for T_e , n_e and P_{rad} . Top: MLP. Bottom: joint fine-tuned surrogate.

To give a global view of the error distribution, we averaged the relative errors for each case across the entire test set and plot the results in figure 16 for T_e , n_e and P_{rad} . Overall, the joint fine-tuned surrogate reduces errors throughout the scrape-off layer compared with the MLP.

3.2. Coverage-gap robustness assessment

To evaluate robustness to gaps in the training data, we conducted a leave-one-value-out (LOVO) test. For each of the five input parameters, we removed every sample that shared one discrete value, trained on the remaining data, and then

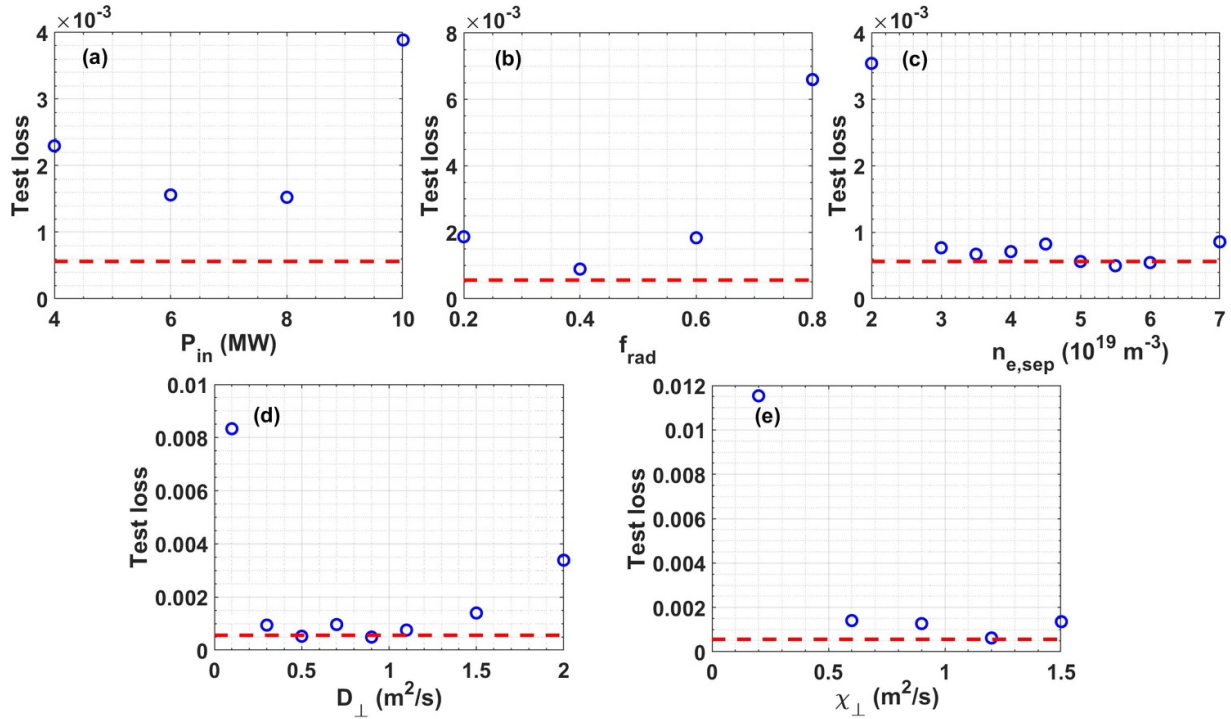


Figure 17. LOVO test loss as a function of the omitted input value. The red dashed line indicates the average validation loss from 5-fold cross-validation with the optimal hyperparameters.

used the removed cases as an independent test set. Each blue marker in figure 17 therefore shows the test loss when all training samples at that value are removed. The horizontal red dashed line indicates the average validation loss from 5-fold cross-validation obtained with the optimal hyperparameters. When the omitted value lies within the training bounds, the test loss remains close to this baseline, demonstrating strong in-domain interpolation despite the gap. However, test loss increases mainly for omissions near domain edges and for omissions that remove a large fraction of cases (panels (a), (b), and e; more than 20%). This analysis suggests that future dataset development should first extend parameter ranges and fill edge regions, and only then consider densifying interior regions. In addition, extrapolation beyond the training range remains challenging.

3.3. Sensitivity assessment of the surrogate's input-output mapping

To examine the input-output relations learned by the final joint fine-tuned surrogate, we performed a post-hoc sensitivity analysis of the three outputs T_e , n_e and P_{rad} , complemented by spatial maps that indicate where each input is most influential. We use partial dependence (PDP) to summarize how each spatial mean output varies with one input. The PDP is the average of the individual conditional expectation (ICE) curves, where each ICE curve is obtained by sweeping a single input while holding the other four fixed. A set of 128 background samples was randomly selected from the input parameter ranges listed in table 1. For each input value, we computed the central 95% pointwise interval across the ICE curves, forming the

shaded ‘95% ICE band’ in figure 18. The PDP for the mean T_e increases with heating power P_{in} and decreases with f_{rad} , $n_{e,sep}$, D_{\perp} , and χ_{\perp} . The mean n_e grows almost linearly with $n_{e,sep}$, drops rapidly and then becomes nearly flat as D_{\perp} increases, and shows only weak dependence on the remaining inputs when $n_{e,sep}$ is fixed. For the mean P_{rad} , the PDP rises with either P_{in} (at fixed f_{rad}) or f_{rad} (at fixed P_{in}), and is nearly flat with respect to the other inputs. Notably, these trends are consistent with physical expectations, further validating the reliability of the joint fine-tuned surrogate model. The width of the 95% ICE band reflects dependence on the other inputs: wider bands indicate stronger input-input interactions.

To provide a dimensionless global ranking of input importance for the mean outputs, normalized gradient sensitivities (elasticities) are used, defined as:

$$E_i = \left| \frac{\partial S}{\partial x_i} \frac{x_i}{S} \right| \quad (1)$$

where S is the spatial mean of each output and x_i is an input. Gradients $\frac{\partial S}{\partial x_i}$ were obtained by automatic differentiation at 256 randomly selected input sets within the range of table 1. Figure 19 plots the median E_i across those sets. For the mean T_e the dominant inputs are P_{in} and $n_{e,sep}$, followed by χ_{\perp} , f_{rad} , and D_{\perp} . For the mean n_e , $n_{e,sep}$ dominates with smaller contributions from P_{in} and f_{rad} . For the mean P_{rad} , P_{in} and f_{rad} are most influential, with the other inputs playing minor roles.

Finally, nondimensional sensitivity (elasticity) maps were constructed at each grid cell to localize sensitivity in space. The maps adopt the same elasticity as equation (1) but are evaluated pointwise by replacing the global mean with the local output value. The resulting per-sample sensitivities are then




Figure 18. Partial dependence of mean outputs (T_e , n_e and P_{rad}) on each input with 95% ICE band. The black curve is the PDP, and the blue shaded region is the 95% ICE band.




Figure 19. Global normalized gradient sensitivities (elasticities) of mean T_e , n_e and P_{rad} to the inputs. Bars show the median elasticities over 256 samples.

averaged at each grid cell over 128 background samples. The resulting maps are shown in figure 20. Bright regions mark locations where small fractional changes in an input lead to large fractional changes in the output. The maps indicate that P_{in} , $n_{e,sep}$, and χ_{\perp} affect T_e broadly across the domain, whereas f_{rad} and D_{\perp} act mainly in the SOL. For n_e , $n_{e,sep}$ controls the profile almost everywhere while the other inputs mostly modulate the SOL. For P_{rad} , the cross-field particle transport D_{\perp} can have a more global impact, because D_{\perp} alters the penetration depth of impurity carbon and thereby reshapes the radiation

distribution throughout. The remaining inputs primarily influence radiation in the SOL.

3.4. Efficiency assessment of the surrogate model

An illustrative example is presented here to show the computational efficiency of the surrogate model. The surrogate requires only about 3.0×10^{-6} core · h for each prediction, whereas a full EMC3-EIRENE run on the present $144 \times 256 \times 144$ mesh on W7-X typically exceeds 1000 core · h, providing a

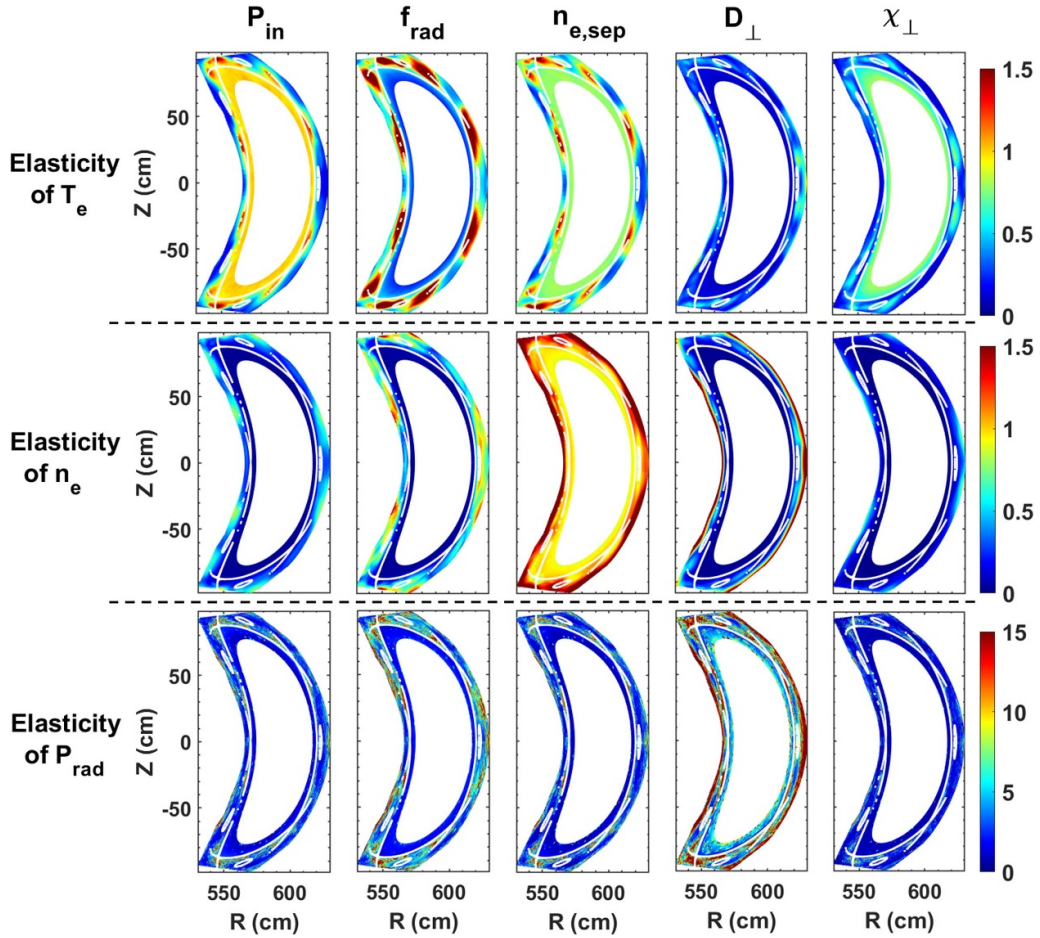


Figure 20. Spatial sensitivity maps of T_e , n_e and P_{rad} to each input, averaged over 128 background samples. Bright regions mark locations of large fractional response to inputs.

speed-up of over 10^8 . This performance makes the surrogate ideal for rapid studies. Tasks such as mapping divertor operating regimes against density, radiation fraction, or cross-field transport coefficients, which are prohibitively slow with the full EMC3-EIRENE code, can now be completed in milliseconds. Figure 21(a) plots the surrogate's recycling flux predictions versus separatrix electron density and f_{rad} at fixed $P_{in} = 6$ MW, $D_{\perp} = 0.5$, and $\chi_{\perp} = 1.5$ m² s⁻¹. The plot includes about 6000 cases, all evaluated on a standard office PC within a few seconds. The recycling flux is a key parameter. Its rollover is an indicator of the onset of divertor detachment. Figure 21(b) shows the probability density of the relative error in recycling flux predictions on the test set. In the vast majority of cases, the errors are quite small, further demonstrating that surrogate model is well-suited for fast large-scale data analysis in the future.

4. Conclusion

We have developed a neural-network surrogate combining an autoencoder and a mapping network to rapidly predict EMC3-EIRENE simulations on W7-X. The autoencoder compresses EMC3-EIRENE outputs into a low-dimensional latent

representation, and the mapping network transforms the five input parameters into this latent space. In the final joint fine-tuned stage, the model achieves the prediction from EMC3-EIRENE inputs to outputs. Hyperparameters were tuned via 5-fold cross-validation combined with the ASHA early-stopping strategy, which avoids biased performance estimates and greatly reduces search time.

The joint fine-tuned surrogate outperforms the simple MLP, especially for detachment cases prediction, and reproduces EMC3-EIRENE results with satisfactory accuracy. The sensitivity analysis conducted with the trained surrogate model has provided deeper insights into the interdependencies among edge plasma parameters. In robustness tests, the surrogate exhibits excellent interpolation capability but weaker extrapolation performance, providing guidance for future dataset construction and expansion. Compared to full EMC3-EIRENE runs, the surrogate achieves at least an 8-order-of-magnitude speed-up, making it well suited for rapid predictions and extensive parameter scans, and paving the way for real-time application of 3D edge transport simulations to optimize ongoing experiments.

In the future, we plan to expand the training dataset and incorporate mesh information as input, enabling the trained model to predict 3D edge parameters for different magnetic




Figure 21. (a) Surrogate-predicted recycling flux as a function of separatrix electron density and radiation fraction at $P_{in} = 6$ MW, $D_{\perp} = 0.5$, and $\chi_{\perp} = 1.5$ m²s⁻¹. (b) Probability density of relative error in recycling flux predictions on the test set.

configurations on W7-X, including the effects of magnetic topology changes induced by high plasma beta. In addition, this surrogate could be deployed in the control room for fast inter-shot analysis, quickly providing plasma parameters in regions not accessible to limited diagnostics, and supporting experimental efforts to understand edge transport characteristics under different plasma scenarios.

Acknowledgments

The project this report was supported with funds from the German Federal Ministry of Education and Research in the Förderprogramm ‘Fusion 2040—Forschung auf dem Weg zum Fusionskraftwerk’, contract number 13F1012C. The sole responsibility for the report’s contents lies with the authors. This work has also been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No. 101052200—EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

The authors gratefully acknowledge the computing time granted by the JARA-HPC Vergabegremium and VSR commission on the supercomputer JURECA at Forschungszentrum Jülich.

ORCID iDs

Y. Luo 0000-0003-2319-1172
 S. Xu 0000-0002-0033-3468
 Y. Liang 0000-0002-9483-6911
 E. Wang 0000-0001-7828-1162

A. Knieps 0000-0003-0083-7188
 Y. Feng 0000-0002-3846-4279
 S. Brezinsek 0000-0002-7213-3326
 M. Krychowiak 0009-0001-4141-5558
 D. Gradic 0000-0002-6109-9345
 M. Jakubowski 0000-0002-6557-3497

References

- [1] Renner H., Boscary J., Greuner H., Grote H., Hoffmann F.W., Kisslinger J., Strumberger E. and Mendelevitch B. 2002 *Plasma Phys. Control. Fusion* **44** 1005
- [2] König R. et al 2002 *Plasma Phys. Control. Fusion* **44** 2365
- [3] Pedersen T.S. et al 2022 *Nucl. Fusion* **62** 042022
- [4] Feng Y. et al 1997 *J. Nucl. Mater.* **241** 930
- [5] Feng Y., Sardei F. and Kisslinger J. 1999 *J. Nucl. Mater.* **266** 812
- [6] Reiter D., Baelmans M. and Börner P. 2005 *Fusion Sci. Technol.* **47** 172
- [7] Winters V.R. et al 2024 *Nucl. Fusion* **64** 056042
- [8] Bold D., Reimold F., Niemann H., Gao Y., Jakubowski M., Killer C., Winters V.R. and Maaziz N. 2024 *Nucl. Fusion* **64** 126055
- [9] Effenberg F. et al 2017 *Nucl. Fusion* **57** 036021
- [10] Winters V.R. et al 2024 *Nucl. Fusion* **64** 126047
- [11] Effenberg F. et al 2019 *Nucl. Fusion* **59** 106020
- [12] Xu S. et al 2023 *Nucl. Fusion* **63** 066005
- [13] Feng Y. et al 2024 *Nucl. Fusion* **64** 086027
- [14] Feng Y. et al 2021 *Nucl. Fusion* **61** 086012
- [15] Stephey L. et al 2018 *Phys. Plasmas* **25** 062501
- [16] Boeyaert D. et al 2024 *Plasma Phys. Control. Fusion* **66** 015005
- [17] Xu S. et al 2018 *Nucl. Fusion* **58** 106008
- [18] Xu S. et al 2020 *Nucl. Fusion* **60** 056006
- [19] Bozhenkov S.A., Geiger J., Grahl M., Kießlinger J., Werner A. and Wolf R.C. 2013 *Fusion Eng. Des.* **88** 2997
- [20] Feng Y. et al 2022 *Plasma Phys. Control. Fusion* **64** 125012
- [21] Seo J., Na Y.-S., Kim B., Lee C.-Y., Park M.S., Park S.J. and Lee Y.H. 2021 *Nucl. Fusion* **61** 106010

- [22] Degrave J. et al 2022 *Nature* **602** 414
- [23] Seo J., Kim S., Jalalvand A., Conlin R., Rothstein A., Abbate J., Erickson K., Wai J., Shousha R. and Kolemen E. 2024 *Nature* **626** 746
- [24] Kates-Harbeck J., Svyatkovskiy A. and Tang W. 2019 *Nature* **568** 526
- [25] Fu Y.C., Eldon D., Erickson K., Kleijwegt K., Lupin-Jimenez L., Boyer M.D., Eidietis N., Barbour N., Izacard O. and Kolemen E. 2020 *Phys. Plasmas* **27** 022501
- [26] Cai J.Q. et al 2024 *Plasma Sci. Technol.* **26** 055102
- [27] Meneghini O. et al 2017 *Nucl. Fusion* **57** 086034
- [28] van de Plassche K.L., Citrin J., Bourdelle C., Camenen Y., Casson F.J., Dagnelie V.I., Felici F., Ho A. and Van Mulders S. 2020 *Phys. Plasmas* **27** 022310
- [29] Dasbach S. and Wiesen S. 2023 *Nucl. Mater. Energy* **34** 101396
- [30] Luo Y. et al 2025 *Nucl. Fusion* **65** 096016
- [31] Burby J.W., Tang Q. and Maulik R. 2021 *Plasma Phys. Control. Fusion* **63** 024001
- [32] Derks G.L., Frankemölle J.P.K.W., Koenders J.T.W., van Berkel M., Reimerdes H., Wensing M. and Westerhof E. 2022 *Plasma Phys. Control. Fusion* **64** 125013
- [33] Poels Y., Derks G., Westerhof E., Minartz K., Wiesen S. and Menkovski V. 2023 *Nucl. Fusion* **63** 126012
- [34] Wiesen S. et al 2015 *J. Nucl. Mater.* **463** 480
- [35] Wiesen S. et al 2024 *Nucl. Fusion* **64** 086046
- [36] Felici F., Citrin J., Teplukhina A.A., Redondo J., Bourdelle C., Imbeaux F. and Sauter O. 2018 *Nucl. Fusion* **58** 096006
- [37] Meneghini O. et al 2021 *Nucl. Fusion* **61** 026006
- [38] Pedersen T.S. et al 2019 *Plasma Phys. Control. Fusion* **61** 014035
- [39] Pedersen T.S. et al 2019 *Nucl. Fusion* **59** 096014
- [40] Li L. et al 2018 arXiv:1603.06560
- [41] Clevert D.-A. et al 2015 arXiv:1511.07289
- [42] Krogh A. et al 1991 *5th Int. Conf. on Neural Information Processing Systems (Denver, 2–5 December)* vol 4 (available at: <https://dl.acm.org/doi/proceedings/10.5555/2986916>)
- [43] Biewald L. *Weights & Biases* 2020 (available at: <https://docs.wandb.ai/guides/app/features/panels/parameter-importance/>)