Deep Learning of Hydrogen Trapping Sites in Tungsten for Atomistic Plasma-Wall Simulations

Seiki SAITO[1)]*, Keisuke TAKEUCHI[1)], Hiroaki NAKAMURA[2,3)], Yasuhiro ODA[4)], Keiji SAWADA[5)], Kazuo HOSHINO[6)], Yuki HOMMA[7)], Shohei YAMOTO[8)], Yuki UCHIDA[9)]

[1)] *Graduate School of Science and Engineering, Yamagata University, Yonezawa 992-8510, Japan*
[2)] *Department of Research, National Institute for Fusion Science, National Institutes of Natural Sciences, Toki 509-5292, Japan*
[3)] *Graduate School of Engineering, Nagoya University, Nagoya 464-8603, Japan*
[4)] *Simulation Engineering Division, Toyota Technical Development Corporation, Toyota 470-0334, Japan*
[5)] *Graduate School of Science and Technology, Shinshu University, Nagano 380-8553, Japan*
[6)] *Department of Applied Physics and Physico-Informatics, Keio University, Yokohama 223-8522, Japan*
[7)] *National Institutes for Quantum Science and Technology, Rokkasho 039-3212, Japan*
[8)] *National Institutes for Quantum Science and Technology, Naka 311-0193, Japan*
[9)] *National Institutes of Technology, Nagaoka College, Nagaoka 940-0817, Japan*

(Received 29 September 2025 / Accepted 22 December 2025)

Understanding hydrogen trapping in tungsten is crucial for accurate modeling of plasma-wall interactions in fusion devices. In this study, we developed a deep learning model to predict hydrogen trapping sites, which serve as essential input parameters for kinetic Monte Carlo (kMC) simulations. The model employs a U-Net–based convolutional neural network that directly maps three-dimensional potential energy distributions to trapping site positions. Training data were generated from atomistic calculations using the embedded atom method (EAM) potential, and ground-truth trapping sites were systematically identified by force relaxation. The trained model achieved an F1 score of approximately 0.76, with most predicted sites coinciding with the true minima within ± 1 voxels (1 voxel = 0.1 Å per side). Visual comparisons confirmed the ability of the model to capture both global and local features of the potential energy landscape. In terms of efficiency, the proposed approach reduced prediction time by more than three orders of magnitude compared with conventional force-relaxation searches, enabling predictions in less than one second on a GPU. These results demonstrate that deep learning provides an accurate and computationally efficient method for identifying trapping sites. Future extensions include incorporating more complex defect structures and integrating the model into molecular dynamics-kMC hybrid frameworks for large-scale plasma-wall interaction simulations.

Keywords: deep learning, U-Net, dice loss, atomistic simulations, hydrogen trapping sites, tungsten, plasma-wall interactions, MD, kMC, fusion materials

DOI: 10.1585/pfr.21.2403017

1. Introduction

Fusion power generation is one of the most promising solutions to the global energy problem. In magnetic confinement fusion devices, the plasma-facing wall is continuously bombarded by high-energy plasma particles, leading to significant changes in its physical and chemical properties. Since these modifications strongly affect both plasma behavior and fuel recycling, elucidating plasma-wall interactions is a key issue [1–6] for realizing steady-state operation of future reactors.

To investigate such complex interactions, we are developing a hybrid computational approach that combines molecular dynamics (MD) and kinetic Monte Carlo (kMC) simulations. MD is suitable for describing the short-timescale injection process of incident plasma particles into the wall, whereas kMC can efficiently simulate long-timescale diffusion processes within the wall. However, the application of the hybrid method has been hindered by the extremely high computational cost required to identify trapping sites and migration barriers, which are indispensable inputs for kMC calculations.

To overcome this difficulty, we are trying to develop a deep learning technique to accelerate plasma-material interaction simulations [7–13]. In our previous work, we developed a deep learning model based on the pix2pix framework to predict the three-dimensional spatial distribution of binding energy from atomic configurations [14]. The previous study

**Corresponding author's e-mail: saitos@yz.yamagata-u.ac.jp*
This article is based on the presentation at the Joint Conference of the 22nd International Conference on Atomic Processes in Plasmas (APiP 2025) and 1st NIFS Conference on Atomic and Molecular Processes in Plasmas.

demonstrated that the model can reproduce binding energy landscapes at a significantly reduced computational cost, establishing a basis for the rapid identification of both potential minima and migration barriers. Nevertheless, additional steps are still necessary to directly connect the predicted energy distribution to the determination of trapping sites used in kMC simulations.

In the present study, we extend this line of research by constructing a deep learning model that directly predicts the positions of hydrogen trapping sites in tungsten lattices. By employing a U-Net architecture and training it on potential energy distributions obtained from atomistic calculations, the proposed model achieves both high prediction accuracy and remarkable computational efficiency. The ultimate goal is to integrate this model into MD-kMC hybrid frameworks to realize realistic plasma-wall interaction simulations under fusion-relevant conditions.

2. Deep-Learning Model

In this study, a deep learning framework was developed to predict the locations of hydrogen trapping sites in tungsten. The prediction task was formulated as a voxel-to-voxel mapping problem, where the three-dimensional potential energy distribution of the tungsten-hydrogen system is provided as input, and the output is a binary map indicating the spatial positions of trapping sites (Fig. 1).

2.1 Network architecture

We employed a U-Net–based convolutional neural network architecture (Fig. 2), which has proven effective in pixel-wise prediction tasks [15, 16]. The model consists of two main parts: an encoder and a decoder. The encoder extracts hierarchical features from the input 3D potential energy distribution through repeated convolution and down-sampling operations, thereby compressing the spatial information into a latent representation. The decoder restores the original spatial resolution by successive up-sampling operations, while skip connections from the encoder to the decoder are incorporated to preserve fine-scale spatial features. This design allows the model to learn both global and local structures of the potential energy landscape that are relevant for determining trapping sites.

Each convolutional layer was followed by a Leaky Rectified Linear Unit (Leaky ReLU) activation, and dropout layers were inserted in the deeper encoder blocks to prevent overfitting. The final output layer applies a sigmoid activation function to provide voxel-wise probabilities, which were thresholded to generate the binary trapping site map.

2.2 Loss function

Since the dataset was highly imbalanced, with the number of non-trapping voxels greatly exceeding that of trapping voxels, a weighted loss function was adopted to improve prediction stability. The total loss L is defined as a linear combination of binary cross-entropy (BCE) loss and Dice loss [17]:

$$L = aL_{\mathrm{BCE}} + (1 - a)\,L_{\mathrm{Dice}}, \tag{1}$$

where the weight parameter a controls the trade-off between voxel-wise accuracy and overall shape similarity of predicted trapping sites. Several values of a were tested to evaluate the sensitivity of the model to this parameter. As shown in Fig. 3, the balance between precision and recall varied with a, and the F1 score was maximized for intermediate values, indicating the importance of appropriate weighting in the loss function. In this study, we used $a = 0.3$.

2.3 Implementation details

The model was implemented in Keras (TensorFlow backend). Training was performed with the Adam optimizer, with an initial learning rate of 1.0×10^{-5}, and the batch size was set to 1 due to GPU memory constraints. Early stopping was applied based on the validation loss to avoid overfitting. All training processes were carried out on an NVIDIA V100 GPU. The model training required 23.48 hours in the computational environment.

3. Preparation of Training Data

To develop the deep learning model, we prepared training data consisting of input-output pairs, where the input is the three-dimensional potential energy distribution of a tungsten-hydrogen system and the output is the spatial positions of trapping sites.

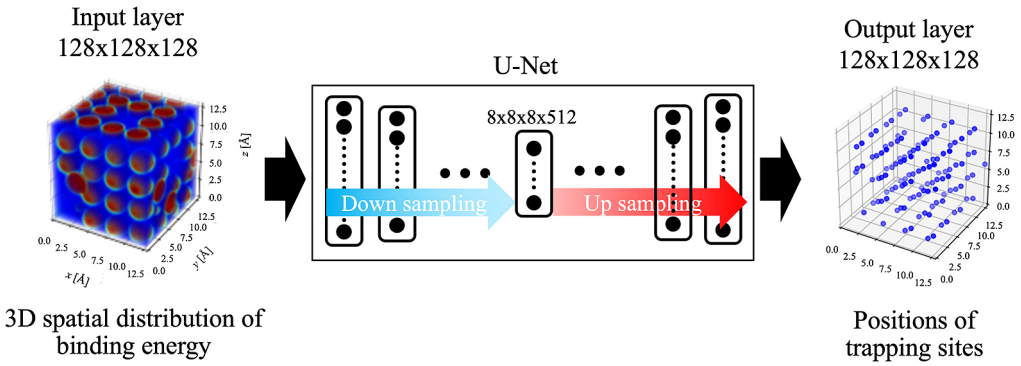

Fig. 1. Voxel-to-voxel mapping from 3D potential energy distribution to hydrogen trapping site positions.

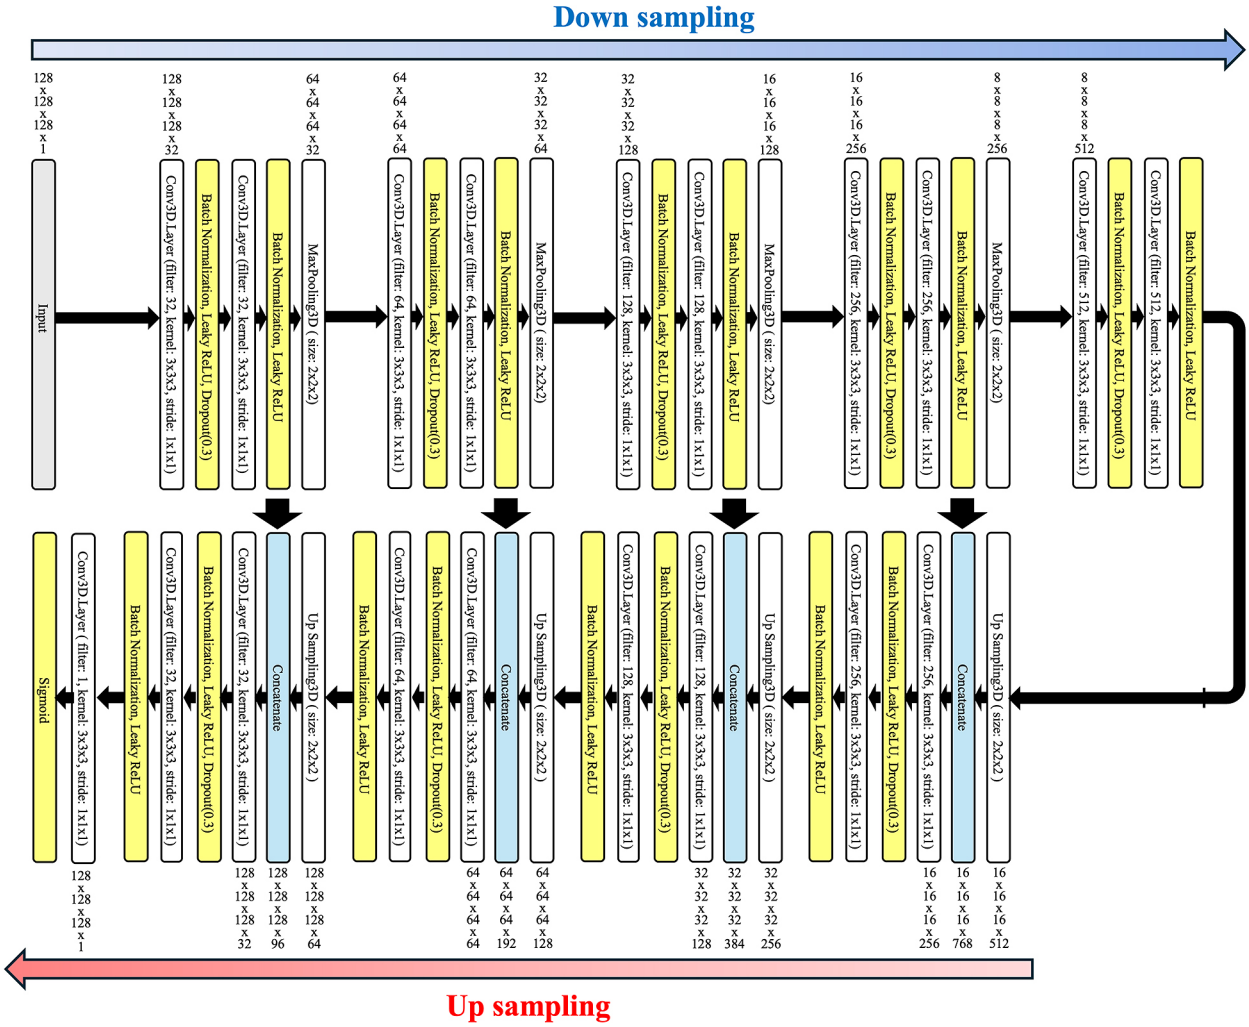

Fig. 2. U-Net architecture of the deep learning model used in this study.

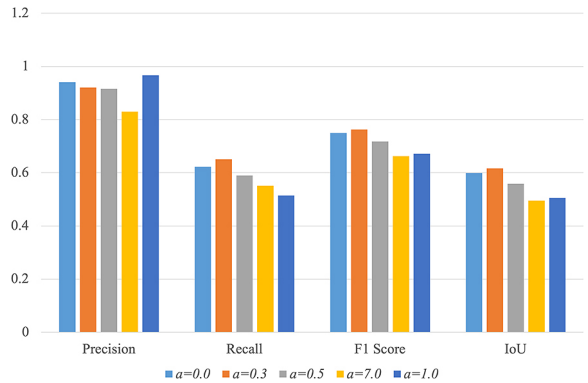

Fig. 3. Effect of weighting parameter a on precision, recall, F1 score, and IoU.

3.1 Calculation of potential energy distributions

We constructed a cubic cell of 12.8 Å on each side. Tungsten atoms were initially arranged in a body-centered cubic (bcc) lattice, but random tungsten atoms were removed to introduce defects (i.e., vacancy-type defects). No self-interstitial atoms (SIAs), dislocations, grain boundaries, or free surfaces were intentionally introduced in the present dataset. To account for thermal vibrations, the positions of tungsten atoms were displaced from their ideal lattice points according to the Boltzmann distribution. In addition, hydrogen atoms were randomly placed within the cell to reflect realistic atomic environments; specifically, 0–4 hydrogen atoms were placed in/near the vacancy defects, and additional hydrogen atoms were also randomly distributed in the bcc lattice. Accordingly, the scope of the present training/test systems is bulk bcc tungsten containing vacancy-type defects and thermal displacements.

The potential energy distributions were calculated using the embedded atom method (EAM) potential proposed by Wang *et al.* [18], which has been widely used for tungsten-hydrogen systems. The total energy of the system is expressed as

$$U = \sum_i F_i\left[\sum_{i\neq j}\rho_j(r_{ij})\right] + \sum_{i,j>i}\phi_{i,j}(r_{ij}), \tag{2}$$

where $F_i(\rho)$ and ρ_i are the embedding energy and effective electron density of ith atom, respectively. $\phi_{i,j}(r)$ and r_{ij} are the pairwise interatomic potential and interatomic distance

between *i*th and *j*th atom, respectively. This formulation allows the many-body effects of metallic bonding to be incorporated efficiently into atomistic simulations.

For each prepared structure, the interaction energy between the configuration and a test hydrogen atom was evaluated. The binding energy of the test hydrogen atom located at position $\boldsymbol{r}_{\mathrm{H}}$ was calculated using the following expression,

$$U^{\mathrm{bind}}_{\mathrm{M+H}}(\boldsymbol{r}_{\mathrm{H}}) = U_{\mathrm{M+H}}(\boldsymbol{r}_{\mathrm{H}}) - (U_{\mathrm{M}} - U_{\mathrm{H}}), \tag{3}$$

where U_{M} is the potential energy of the material, $U_{\mathrm{M+H}}$ is the potential energy of the material together with the test hydrogen atom, and U_{H} is the potential energy of the isolated hydrogen atom. The test hydrogen atom was systematically moved on a three-dimensional grid with 0.1 Å spacing, and the binding energy was computed at each grid point. To improve the predictive performance of the neural network, the binding energy values were normalized as follows:

$$\tilde{U}^{\mathrm{bind}}_{\mathrm{M+H}}(\boldsymbol{r}_{\mathrm{H}}) = \begin{cases} 1 & \text{if } U^{\mathrm{bind}}_{\mathrm{M+H}}(\boldsymbol{r}_H) \geq 10\ \mathrm{eV}, \\ \dfrac{U^{\mathrm{bind}}_{\mathrm{M+H}}(\boldsymbol{r}_H)}{10\ \mathrm{eV}} & \text{if } U^{\mathrm{bind}}_{\mathrm{M+H}}(\boldsymbol{r}_H) < 10\ \mathrm{eV}. \end{cases} \tag{4}$$

3.2 Identification of trapping sites

To generate the ground-truth labels, the positions of local minima in the potential distribution were determined. A test hydrogen atom was initially placed at a random position, and only the position of this atom (while keeping all other atoms fixed) was displaced step-wise along the direction of the acting force. When the force magnitude fell below a threshold of 10^{-5} eV/Å, a local minimum was identified. By repeating this procedure with different random initial positions of the test hydrogen atom, trapping sites were comprehensively explored. To remove redundancy, minima within 0.001 Å of each other were merged into a single representative trapping site by averaging their coordinates.

Figure 4 shows the dependence of the number of distinct trapping sites on the number of sampled test hydrogen atoms. The curve indicates that the number of unique trapping sites saturates as the sampling density increases, confirming the reliability of the method in capturing all relevant minima. From this figure, it can be seen that a sampling size of 10,000 points is sufficient; therefore, in this study, 10,000 test positions were used to identify local minima and to generate the training dataset.

3.3 Construction of training pairs

Finally, input–output pairs were constructed by voxelizing both the calculated potential distribution (input) and the trapping site positions (output). The output data were represented as binary 3D maps, in which voxels corresponding to trapping sites were assigned a value of 1 and all other voxels were set to 0. These pairs provide the essential data for training the U-Net model described in Sec. 2.

Although the present 3D input is constructed for a 4 × 4 × 4 bcc supercell, this does not restrict the method to that size.

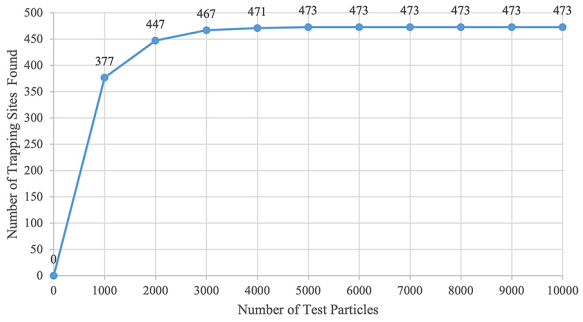

Fig. 4. Number of distinct trapping sites as a function of sampled test hydrogen atoms.

The trained model can be applied to a larger simulation cell by a sliding-window (tiling) strategy, where overlapping 4 × 4 × 4 subcells are extracted while shifting the window across the system. Predictions from all subcells are then aggregated to identify trapping-site candidates in the entire domain. The total computational cost scales approximately linearly with the number of evaluated subcells.

4. Prediction Results

4.1 Visual comparison of predictions

The performance of the proposed deep learning model was first evaluated by comparing the predicted trapping site distributions with the ground truth obtained from atomistic calculations. Figure 5 shows two-dimensional slices taken along the *z*-axis from randomly selected three-dimensional voxelized training data. The slicing coordinate was chosen so that the slice contained a large number of trapping sites. The normalized interaction energy is displayed as a grayscale contour plot, while the ground-truth trapping sites are shown as red filled circles and the predicted positions are marked with green crosses. For both the ground-truth and predicted sites, markers are plotted together with their surrounding ranges of ± 1 voxel (± 0.1 Å). The predicted sites generally coincide with the true minima, with most discrepancies limited to regions within ± 1 voxel of the true positions. These results demonstrate that the model can effectively capture the spatial features of the potential energy landscape that are essential for identifying hydrogen trapping sites.

4.2 Quantitative evaluation metrics

To evaluate prediction accuracy, four metrics were calculated: precision, recall, F1 score, and intersection-over-union (IoU). Here, predictions within ± 1 voxels in *x*-, *y*-, *z*-directions of the ground-truth positions were considered correct to account for discretization errors in voxelization. The results are summarized in Fig. 6. The model achieved a maximum F1 score of approximately 0.76, demonstrating balanced performance between precision and recall. The IoU values also confirmed that the overlap between predicted and true sites was substantial.

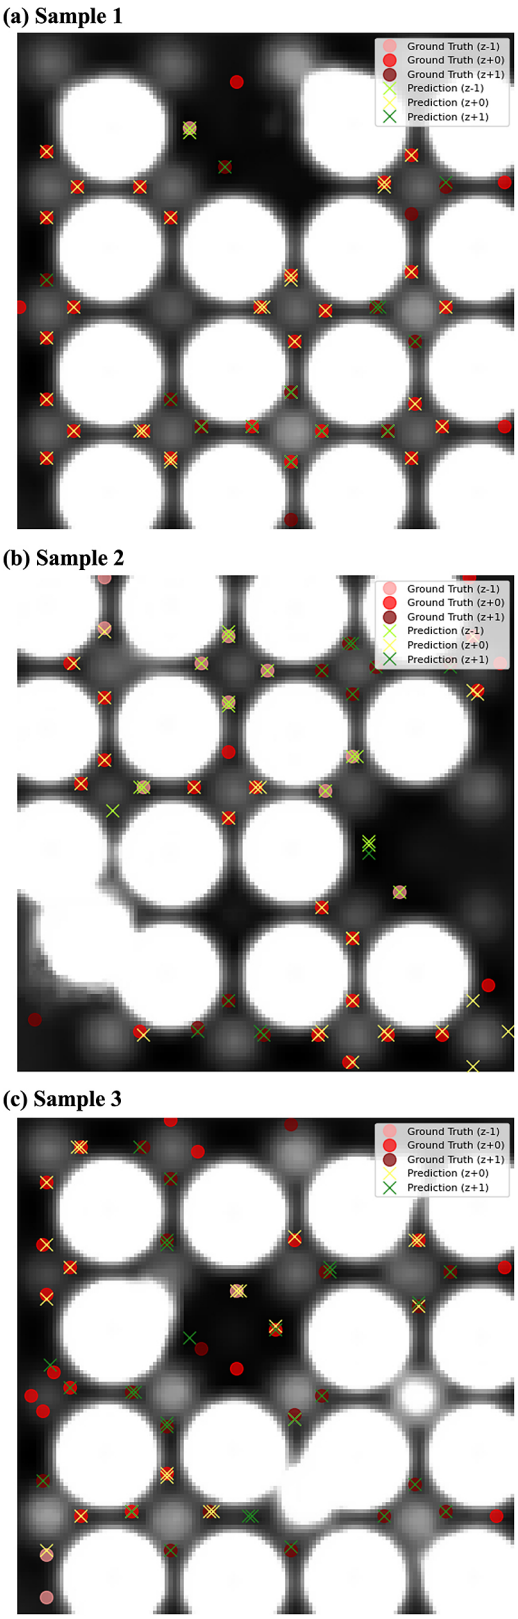

Fig. 5. Example 2D slice showing normalized interaction energy (gray-scale), ground-truth sites (red circles), and predictions (green crosses).

4.3 Computational efficiency

Table 1 compares the computation times required for conventional atomistic calculations and for the proposed deep learning model. The reference time $t_{\text{true}}^{\text{1CPU}}$ corresponds to the direct force-relaxation calculation of trapping sites, equivalent to Eq. (3), performed using a single core of an Intel Xeon Gold 6130 processor. The prediction time $t_{\text{pred}}^{\text{1CPU}}$ represents the inference of the trained deep learning model executed on the same single CPU core, while $t_{\text{pred}}^{\text{1CPU+GPU}}$ represents the inference time when the model is executed on an NVIDIA Tesla V100 GPU (16 GB) in combination with the same CPU processor. All times were measured for detecting trapping sites in single spatial distribution of interaction energy, represented on a 128 × 128 × 128 voxels. The conventional method requires several hundred seconds per configuration, whereas the trained model predicts the trapping sites in less than one second when accelerated by a GPU. This corresponds to a speedup of more than three orders of magnitude, clearly demonstrating that the proposed approach is computationally efficient and suitable for integration into MD-kMC hybrid simulations.

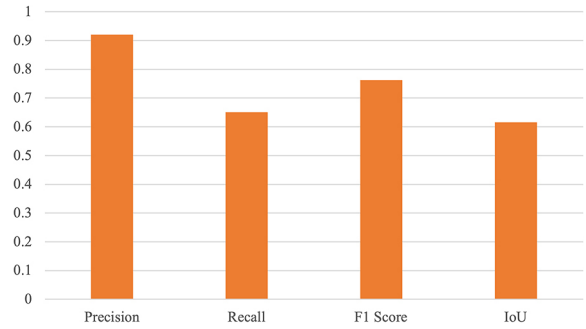

Fig. 6. Prediction performance metrics: precision, recall, F1 score, and IoU.

Table 1. Prediction time using the developed deep learning model.

	$t_{\text{true}}^{\text{1CPU}}$	$t_{\text{pred}}^{\text{1CPU}}$	$t_{\text{pred}}^{\text{1CPU+GPU}}$
Time [s]	690.0	8.15	0.21
Speed-up ratio	1	85	3,286

5. Summary

In this study, we developed a deep learning model for predicting hydrogen trapping sites in tungsten, which are essential input parameters for kMC simulations of plasma-wall interactions in fusion devices. By employing a U-Net–based architecture, the model directly maps three-dimensional potential energy distributions to trapping site positions.

The training dataset was constructed using potential distributions obtained with the EAM potential, and the corresponding local minima were systematically identified as ground truth. The trained model successfully reproduced the spatial distribution of trapping sites, achieving an F1 score of approximately 0.76. Visual inspections confirmed that most predicted sites coincided with the true minima, with deviations limited to within ± 1 voxels.

In terms of computational performance, the proposed

approach reduced the prediction time by more than three orders of magnitude compared with conventional force-relaxation searches, demonstrating its applicability to large-scale plasma-wall interaction simulations.

Future work will focus on extending the present framework to incorporate more complex atomic configurations, including defects and hydrogen accumulation, and on integrating the model into MD-kMC hybrid simulations. Such developments are expected to significantly accelerate realistic modeling of plasma-facing materials under fusion-relevant conditions.

Acknowledgments

The research was partially supported by Grant-in-Aid for Scientific Research Nos. 22K03572 and 23K03362 from the Japan Society for the Promotion of Science, by the NINS program of Promoting Research by Networking among Institutions (01322301), and by the NIFS Collaborative Research Program NIFS24KIPT013, NIFS25KSPT009, and NIFS22KIGS002. The computations were performed using the JFRS-1 supercomputer system at Computational Simulation Centre of International Fusion Energy Research Centre (IFERC-CSC) in Rokkasho Fusion Institute of QST (Aomori, Japan), using Research Center for Computational Science, Okazaki, Japan (Project: 24-IMSC099), and Plasma Simulator of NIFS (Toki, Gifu, Japan).

[1] G. Janeschitz *et al.*, J. Nucl. Mater. **220-222**, 73 (1995).
[2] G.F. Matthews, J. Nucl. Mater. **220-222**, 104 (1995).
[3] B. Lipshulz *et al.*, J. Nucl. Mater. **220-222**, 967 (1995).
[4] C.S. Pitcher and P.C. Stangeby, Plasma Phys. Control. Fusion **39**, 779 (1997).
[5] S.I. Krasheninnikov *et al.*, Phys. Lett. A **214**, 285 (1996).
[6] K. Ohsawa *et al.*, J. Nucl. Mater. **458**, 187 (2015).
[7] S. Saito *et al.*, Nucl. Fusion **64**, 126067 (2024).
[8] S. Saito *et al.*, Nucl. Mater. Energy **43**, 101942 (2025).
[9] S. Saito *et al.*, Contrib. Plasma Phys. e201900152 (2020).
[10] S. Saito *et al.*, Plasma Fusion Res. **15**, 2403073 (2020).
[11] H. Nakamura *et al.*, Jpn. J. Appl. Phys. **61**, SA1005 (2022).
[12] S. Saito *et al.*, Jpn. J. Appl. Phys. **60**, SAAB08 (2021).
[13] K. Sawada *et al.*, Contrib. Plasma Phys. e201900153 (2020).
[14] S. Saito *et al.*, Jpn. J. Appl. Phys. **63**, 09SP03 (2024).
[15] O. Ronneberger *et al.*, MICCAI 2015, LNCS **9351**, 234 (2015).
[16] Ö. Çiçek *et al.*, MICCAI 2016, LNCS **9901**, 424 (2016).
[17] F. Milletari *et al.*, *Proceedings of 2016 Fourth International Conference on 3D Vision (3DV)*, (IEEE, 2016), 565–571.
[18] L.F. Wang *et al.*, J. Phys. Condens. Matter. **29**, 435401 (2017).