Construction of generalized quasilinear diffusion coefficient using neural networks with physical restrictions

Gyeonghun Pyeon[1,2,*], *Gregory* Wallace[2], *Paul* Bonoli[2], and *John* Wright[2]

[1]Department of Nuclear Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[2]Plasma Science and Fusion Center, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Abstract. The quasilinear diffusion coefficient (D_{QL}) derived from our machine learning framework shows comparable trends with the ground truth D_{QL} obtained from GENRAY-CQL3D simulations. Additionally, for the strong absorption cases, the radial current drive profiles generated using the D_{QL} from our model exhibit consistent behavior with those obtained from the original simulation. These findings indicate the potential of our surrogate modeling approach with physical restrictions to replicate key wave–plasma interaction characteristics while reducing computational costs. Traditionally, calculating D_{QL} for wave–particle interactions relies on computationally intensive wave simulations coupled with Fokker–Planck solvers. To address this challenge, we developed a machine learning-based surrogate model with physical restrictions derived from cold plasma theory and bounce-averaged damping effects. First, we establish the propagation domain of Lower Hybrid Waves in the $(N_{\|}, \rho)$ space by identifying the accessibility limit and determining the upper and lower bounds of $N_{\|}$ using the Potential Power Deposition (PPD) method. Subsequently, leveraging a database constructed using Latin hypercube sampling alongside the underlying physical restrictions (*e.g.* PPD), machine learning methods including U-Net and Recurrent Neural Networks are employed to design a physics-restricted machine learning framework capable of reconstructing D_{QL}.

1 Introduction

Current drive, a technique used to generate and sustain plasma current in a tokamak, is essential for achieving long-pulse, steady-state operation. Among various current drive methods, Lower Hybrid Current Drive (LHCD) stands out due to its high current drive efficiency [1]. The current drive by LHCD is achieved through wave–particle interactions. In quasilinear theory these interactions appear as a diffusion in velocity space, governed by a quasilinear-diffusion coefficient (D_{QL}).

However, calculating D_{QL} for wave–particle interactions requires computationally intensive wave propagation simulations coupled with Fokker–Planck solvers. Traditionally, there are two approaches to simulating wave propagation. The first is full-wave simulation, which directly solves Maxwell's equations (*e.g.* TORIC [2]). The second approach is ray tracing, which approximates waves as rays and computes their propagation paths (*e.g.* GENRAY [3]). Both approaches are computationally demanding; full-wave simulations require substantial computation time, while ray tracing, although relatively faster, still incurs high computational costs when many simulations are needed. After calculating the wave propagation, this information is then passed to a Fokker–Planck solver (*e.g.*, CQL3D [4]) to compute the D_{QL}. In the case of the coupled GENRAY-CQL3D code, GENRAY first simulates the wave propagation and evaluates the power deposition profile across the plasma. This power deposition information is subsequently passed to CQL3D, which performs bounce- and flux-surface-averaging to obtain initial D_{QL}. Using this D_{QL}, CQL3D updates the distribution function based on Fokker-Planck equation. After the distribution function is updated, a new power deposition profile is recalculated, leading to a revised D_{QL}. This iterative loop continues until both the distribution function and D_{QL} converge, resulting in high computational costs. In the case of full wave simulations, the computational costs become even higher. Unlike the coupled GENRAY-CQL3D code, which assumes that the ray path remains fixed and computes it only once, the coupled TORIC-CQL3D code, CQL3D returns the updated distribution function to TORIC, which then recalculates the wave propagation during each iteration. Consequently, the calculation of D_{QL} is computationally demanding in both ray-tracing and full-wave approaches, since it involves computationally intensive wave propagation simulations coupled with an iterative Fokker–Planck solver. This further underscores the necessity of developing efficient and accurate surrogate models.

To address the computational challenges associated with D_{QL} calculations in wave–particle interactions, Sanchez-Villar *et al.* [5, 6] developed machine learning based surrogate models to reproduce Ion Cyclotron Range of Frequencies (ICRF) heating results. In parallel, Wallace *et al.* [7] focused on predicting the final current drive value. Extending this work, the present work aims to

*Corresponding author: pyeon924@mit.edu

develop a machine learning–based surrogate model for efficient and accurate prediction of the fully converged D_{QL}, from which the current drive can then be calculated. The proposed surrogate model uses physical restrictions derived from cold plasma theory and bounce-averaged damping effects. By leveraging machine learning techniques, this research seeks to predict D_{QL} while significantly reducing computational costs (speeding up time to compute and reducing the number of iterations required) and simultaneously maintaining the accuracy of traditional numerical simulations.

This study was conducted using the following approach. First, a comprehensive database was constructed using coupled CQL3D–GENRAY simulations, systematically exploring various parameter spaces. The Latin hypercube method [9] was employed to ensure efficient and uniform sampling of the high-dimensional parameter space, enabling a well-distributed dataset for subsequent analysis. Next, the propagation domain of Lower Hybrid Waves in the $(N_{\|}, \rho)$ space was identified by analyzing the accessibility limit and Landau damping limit, and by determining the upper and lower bounds of $N_{\|}$ based on the Stix cold plasma model [10], where $N_{\|}$ denotes the parallel refractive index and ρ represents the square root of the normalized toroidal flux, here $\sqrt{\psi}$ for simplicity. Based on this propagation domain, the Potential Power Deposition (PPD) regions [11] in velocity space—where D_{QL} could exist—were derived by examining the effects of bounce-averaging and flux-surface averaging within the identified $(N_{\|}, \rho)$ domain. Building upon the aforementioned database and physical insights, various machine learning techniques, including Recurrent Neural Networks (RNN) and U-Net [12], were employed. These methods were integrated into a machine learning framework designed to efficiently reconstruct the converged value of D_{QL}.

The remainder of this paper is structured as follows. Section 2 details the construction of the dataset and simulation framework used to generate the data. Section 3 introduces the surrogate model for D_{QL} prediction. Finally, Section 4 presents the performance of the model in terms of D_{QL} prediction accuracy, as well as simulation results obtained using the reconstructed D_{QL} in GENRAY-CQL3D.

2 Simulation Framework and Dataset Generation

2.1 Overview of GENRAY-CQL3D Simulation Code

To construct the dataset for D_{QL}, the coupled GENRAY-CQL3D simulation framework has been used. It determines the radial profile of bounce-averaged 2D ion and electron distribution functions in momentum space, accounting for RF power deposition. In the computation of D_{QL}, GENRAY first computes the wave trajectories and electric fields along those trajectories. This data is then coupled with the CQL3D code to calculate the quasilinear diffusion coefficient.

The update of the distribution function due to the interaction between the injected wave and plasma in pitch-angle coordinates is given by:

$$\left.\frac{\partial f}{\partial t}\right|_{QL} = \frac{1}{u^2}\frac{\partial}{\partial u}\left[\left(\delta B_0 \frac{\partial}{\partial u} + \delta C_0 \frac{\partial}{\partial \theta}\right) f\right] + \frac{1}{u^2 \sin\theta}\frac{\partial}{\partial \theta}\left[\left(\delta E_0 \frac{\partial}{\partial u} + \delta F_0 \frac{\partial}{\partial \theta}\right) f\right]. \quad (1)$$

Here, f represents distribution function. u represents the relativistic velocity, and θ denotes the pitch angle. The terms δB_0, δC_0, δE_0, and δF_0 refer to the bounce- and volume-averaged local quasilinear diffusion coefficients in their respective directions. Since δB_0 determines the other local diffusion coefficients (*i.e.* δC_0, δE_0, and δF_0) through expressions provided by Kerbel and McCoy [13], this research specifically aims to reconstruct the δB_0 coefficient. For simplicity, we denote δB_0 as D_{QL} throughout the remainder of this paper.

2.2 Dataset Generation

The input parameters were selected to span ranges relevant to LHCD, based on their direct influence on wave propagation and power deposition characteristics, particularly in the context of the EAST tokamak. The selected input parameters include the starting parallel refractive index at the launcher $N_{\|}^{\text{start}}$ (ranging from 1.5 to 6.5), the magnetic field at the plasma center B_0 (1–4 T), the central electron temperature T_{e0} (1–7 keV), and the central electron density n_{e0} (from 1×10^{19} to 4×10^{19} m^{-3}).

While the GENRAY and CQL3D codes require a significantly larger set of input parameters for execution, many were held fixed using standard values representative of the EAST tokamak. In particular, parameters related to the tokamak geometry (*e.g.* the major radius R_0, minor radius a, etc.) and the lower hybrid wave system (*e.g.*, the LH wave frequency, etc.) were configured based on the established specifications of EAST. In addition, default settings provided by the GENRAY-CQL3D package were used for less critical inputs to simplify the simulation workflow.

A total of 10,000 simulations were conducted under the prescribed input conditions, with sampling performed using the Latin hypercube method [9]. Each simulation produced a corresponding D_{QL}, defined in velocity space with pitch-angle coordinates and the radial coordinate $(u/v_{\text{norm}}, \theta, \sqrt{\psi})$. Here, u denotes the relativistic velocity, v_{norm} is the maximum momentum divided by the rest mass on the mesh, θ is the pitch angle, and $\sqrt{\psi}$ is the square root of the normalized toroidal flux. Consequently, the resulting dataset spans a three-dimensional phase-space grid. The mesh is discretized into 320 points in u/v_{norm} over the range $[0, 1]$, 240 points in θ over the range $[-\pi, \pi]$, and 23 ψ indices, yielding a dimensionality of $320 \times 240 \times 23$ per simulation. Here, $\rho \approx \sqrt{\psi}$, defined over the range 0.025–0.95. The complete dataset was subsequently partitioned into training (80%), validation (10%), and test (10%) subsets.

3 Surrogate Model for D_{QL} Prediction

In this section, we introduce the surrogate architecture employed to predict D_{QL} as well as the physical restrictions applied to post-process final D_{QL} predictions. As noted in the previous section, D_{QL} is not a one-dimensional variable, but a three-dimensional quantity defined over the velocity–pitch–radial phase space. Flattening this 3D output to directly predict hundreds of thousands of elements using a fully connected neural network poses a significant challenge, as it requires an enormous number of parameters, resulting in high memory consumption and slow training times. To address this challenge, we introduce a specialized surrogate architecture.

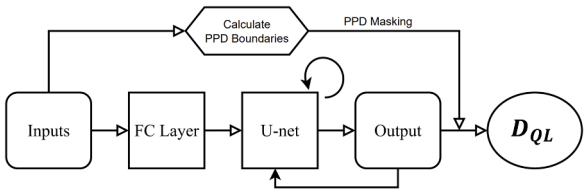

Figure 1. Machine learning network architecture diagram integrating U-Net and physical post-processing for D_{QL} prediction.

3.1 Surrogate Architecture

For our surrogate architecture, we employed U-Net and RNN components based on two key observations. First, when examining the structure of D_{QL} in two-dimensional velocity space, we found that it resembles image-like patterns with coherent spatial structures. Second, from a physical perspective, wave injection in tokamak devices occurs from the outer boundary toward the plasma core. This implies that the values of D_{QL} reconstructed at outer flux surfaces (*i.e.*, larger $\sqrt{\psi}$) contain predictive information useful for estimating D_{QL} at inner surfaces. Based on these insights, we adopted the following neural network strategy: to reconstruct D_{QL} in velocity space, we employed a U-Net [12]. Its encoder–decoder structure, along with skip connections, enables the capture of both localized features and global context, which is effective for reconstructing complex spatial patterns in D_{QL}. In addition, to model the sequential nature of wave propagation across flux surfaces, we incorporated an RNN across $\sqrt{\psi}$: at each step, the prediction at $\sqrt{\psi}^{\text{index}} = n$ informs the prediction at the adjacent inner surface $\sqrt{\psi}^{\text{index}} = n - 1$, enabling the model to capture inter-surface dependencies and enhance overall predictive accuracy across the radial domain.

Using a U-Net and an RNN, we designed a surrogate architecture for predicting D_{QL}. The overall flow is presented in Fig. 1, and Fig. 2 details the layer configuration of the U-Net block. We now explain how the architecture operates: First, a fully connected layer expands the 4 scalar input parameters to a vector of size 76,800. This vector is reshaped into a tensor of dimensions 1×320×240, corresponding to the velocity space $(u/v_{\text{norm}}, \theta)$ at a given flux surface. The reshaped tensor serves as the input to a U-Net, which extracts spatial features and produces an output of the same shape, representing the predicted D_{QL} at the outermost radial position, $\sqrt{\psi}^{\text{index}} = n_\psi$ (where n_ψ denotes the total number of radial indices, which is 23 in our case as explained in Section 2). The raw U-Net output is then processed by a physics-based masking procedure that applies restrictions derived from the input parameters to remove non-physical regions of phase space. Here, the physically meaningful region is defined by the so-called PPD boundaries, which will be described in detail in Sec. 3.2. The masked output is treated as the final prediction of D_{QL} at $\sqrt{\psi}^{\text{index}} = n_\psi$, while the unmasked output (*i.e.*, the raw U-Net output prior to masking) together with an internal hidden state is utilized to guide the prediction at the next inner flux surface, $\sqrt{\psi}^{\text{index}} = n_\psi - 1$. This procedure is repeated recursively until the innermost flux surface, $\sqrt{\psi}^{\text{index}} = 1$, is reached. A schematic representation of this algorithm is provided in Algorithm 1.

Algorithm 1: Recursive Prediction of D_{QL} Using U-Net with Physics-Based Masking

Input : Input parameters: $N_{\parallel}^{\text{start}}, B_0, T_{e0}, n_{e0}$
Output: Predicted D_{QL} across all ψ indices
Compute physical boundaries from input parameters (PPD mask);
Expand input parameters using a fully connected layer;
Reshape the result to a tensor of shape $1 \times 320 \times 240$;
for $\psi = n_\psi \to 1$ **do**
  Pass reshaped input and hidden state $h_{\psi+1}$ to U-Net to obtain raw $\hat{D}_{QL}^{(\psi)}$;
  Apply PPD masking to $\hat{D}_{QL}^{(\psi)}$ (post-processing) to obtain final $D_{QL}^{(\psi)}$;
  Update hidden state $h_\psi \leftarrow \text{HiddenUpdate}(h_{\psi+1}, \hat{D}_{QL}^{(\psi)})$;
  Accumulate masked $D_{QL}^{(\psi)}$ into the output tensor;
return All masked $D_{QL}^{(\psi)}$

Remark on enforcement. The PPD mask enforces physics as a post-processing (physics-guided) restriction on the network output. Loss is evaluated on the unmasked D_{QL} only; masked D_{QL} are excluded from the loss and are not used to propagate gradients. The derivation of the PPD regions and their dependence on the input parameters is presented next.

3.2 Physical Restrictions Applied to Predict D_{QL}

In this section, we discuss the physical restrictions to enhance the prediction of D_{QL}. The section is structured as follows: we first explain how the dispersion relation limits the parallel refractive index, $N_{\parallel}$, during wave propagation in the plasma and identify the range in which power absorption could occur. Then, we describe how these restrictions define permissible regions in velocity space where D_{QL} can exist.

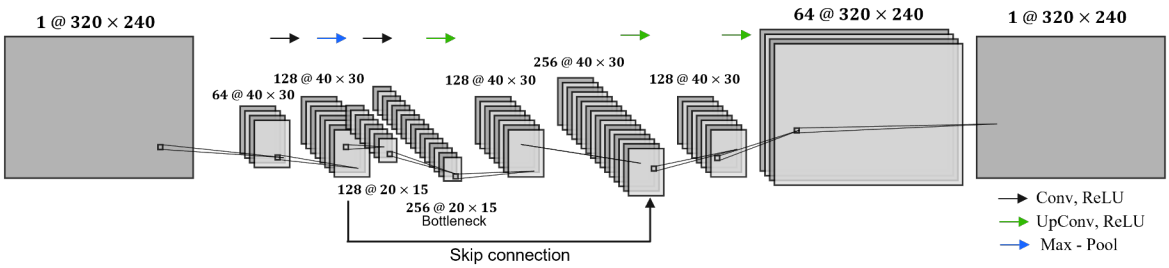

Figure 2. Overall configuration of the U-Net block used for D_{QL} prediction. Each block shows the number of channels and the feature map size ($C@H \times W$), and the arrows denote the operations applied when moving between stages (convolution, pooling, or up-convolution).

3.2.1 *Power Absorption Range of $N_\parallel$ During LH Wave Propagation*

During the propagation of lower hybrid waves in plasma, the parallel refractive index $N_\parallel$ is bounded by the following relation:

$$N_\parallel = \frac{N_{\varphi,\mathrm{axis}}}{1 \mp \sqrt{\alpha}\sqrt{n_e}\frac{B_\theta}{|\mathbf{B}|}}, \tag{2}$$

where the negative and positive signs correspond to the upper and lower bounds of $N_\parallel$, respectively [14]. Here, $\alpha = \frac{\omega_{pe}^2}{S\omega^2 n_e}$, with S being the "sum" component of the Stix dielectric tensor for cold plasmas [10], ω_{pe} the plasma frequency, ω the LH wave frequency, n_e the electron density, $\mathbf{B}$ the total magnetic field, B_θ the poloidal magnetic field component, and $N_{\varphi,\mathrm{axis}}$ the toroidal refractive index at the magnetic axis. The Landau damping condition is given by [15]:

$$N_{\parallel,L} \sim \frac{5.7}{\sqrt{T_e\,(\mathrm{keV})}}, \tag{3}$$

and the accessibility condition is [16]:

$$N_{\parallel,\mathrm{acc}} \gtrsim \frac{\omega_{pe}}{|\omega_{ce}|} + \sqrt{S}, \tag{4}$$

where ω_{ce} is the electron cyclotron frequency. Since n_e and B_θ vary with ρ, these restrictions can be represented in the $(N_\parallel, \rho)$ space. For a specific ρ, if the Landau damping condition lies within the accessible $N_\parallel$ range, LH wave power can potentially be deposited. This region is referred to as the PPD region. In this study, we use these restrictions to enforce physical limits on the output of our model.

3.2.2 *Potential Power Absorption Region in Velocity Space*

We now examine how the range of $N_\parallel$ obtained in the previous subsection determines the velocity space region where LH wave power may be deposited.

For particles following bounce trajectories, the midplane parallel velocity can be expressed, based on energy and magnetic moment conservation, as $v_\parallel^2 = v_{\parallel 0}^2 - \delta b\, v_{\perp 0}^2$, where $v_{\parallel 0}$ and $v_{\perp 0}$ are the parallel and perpendicular components of velocity at the midplane, and δb is defined as $\delta b(\ell) \equiv \frac{B(\ell)}{B_0} - 1$, with ℓ representing the bounce path length. The parallel velocity $v_\parallel$ of a resonant particle is related to the parallel refractive index $N_\parallel$ by $v_\parallel = \frac{c}{N_\parallel}$. Therefore, restrictions on $N_\parallel$ translate into restrictions on $v_{\parallel 0}$ and $v_{\perp 0}$ at the midplane.

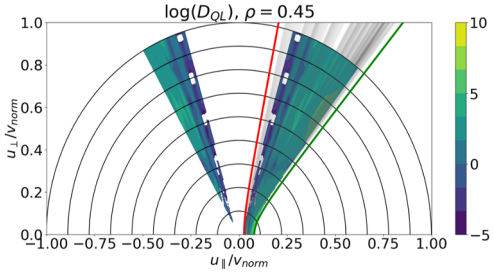

Figure 3. Permissible midplane velocities and corresponding $\log(D_{QL})$ values in velocity space at $\rho = 0.45$, with $N_\parallel^{start} = 3.0$, $n_{e0} = 2.0 \times 10^{19}$ m^{-3}, $T_{e0} = 2.3$ keV, and $B_0 = 2.5$ T. In the figure label, $u_\parallel$ denotes the parallel component of the relativistic velocity, and $u_\perp$ denotes the perpendicular component of the relativistic velocity. Gray lines indicate mid-plane velocities corresponding to all possible $N_\parallel$ combinations. The red and green lines indicate the left and right boundaries of PPD, corresponding to the upper limit of $N_\parallel$ and the accessibility limit, respectively

To calculate the permissible velocity region for our simulation's D_{QL} at a specific radial index, we must take into account how GENRAY-CQL3D performs its calculations. In our case, the GENRAY-CQL3D simulation applies both bounce-averaging and flux-surface-averaging when computing D_{QL} at each radial location. As a result, the determination of the permissible region in midplane velocity space requires considering the union of all valid $N_\parallel$ ranges, as well as the corresponding combinations of temperature, density, and magnetic field, within the spatial region covered by the averaging procedures.

Fig. 3 shows the resulting permissible D_{QL} regions in midplane velocity space for an example case. The gray lines represent possible midplane velocities corresponding to the spatial region covered by the averaging procedures. Among the regions outlined by these gray lines, we select the leftmost red line, corresponding to the upper limit of $N_\parallel$, and the rightmost green line, corresponding to the accessibility limit, as the PPD boundary. Since D_{QL} varies significantly across the velocity space, we present D_{QL} on

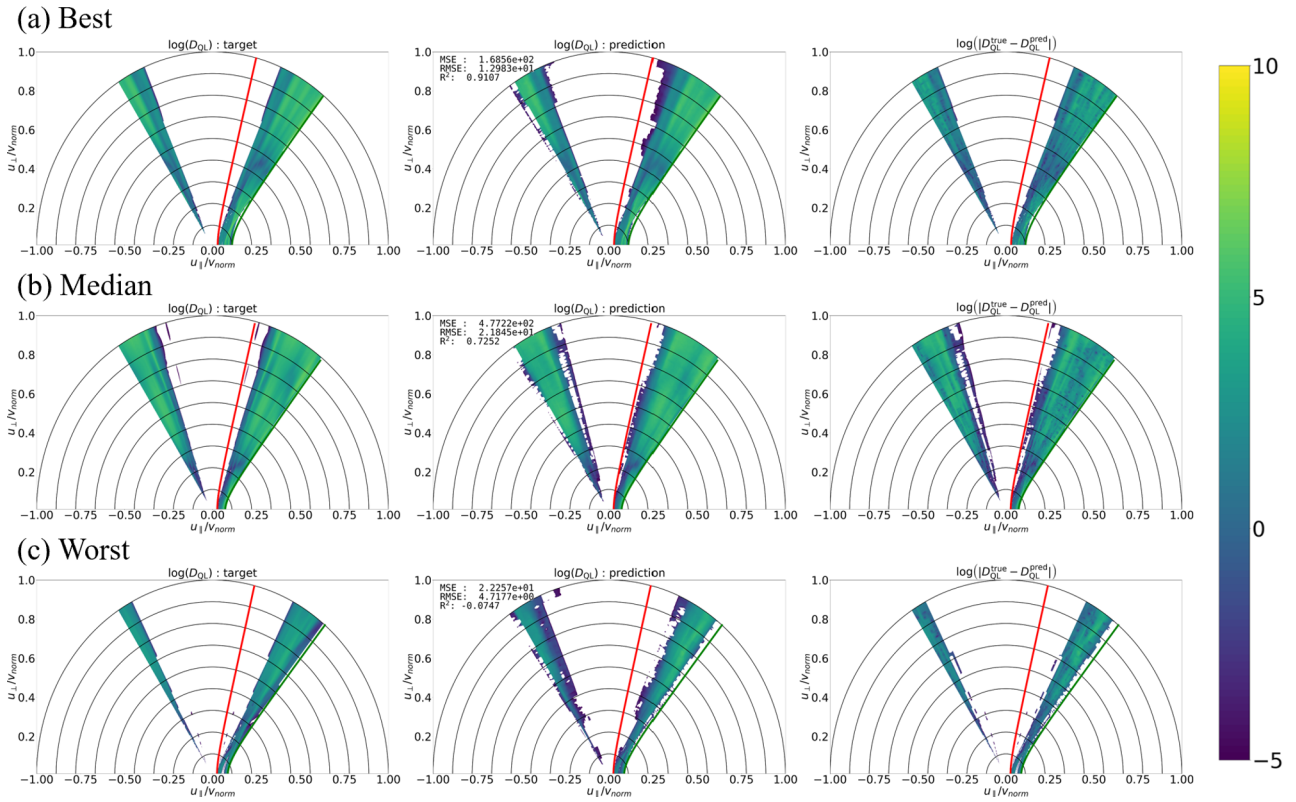

Figure 4. Comparison of simulated ground truth (first column), U-Net prediction (second column), and their difference (third column) for D_{QL} from the test dataset, where (a) the top row shows the sample with the maximum R^2, (b) the middle row shows the sample with the median R^2 , and (c) the bottom row shows the sample with the minimum R^2 . The red and green lines indicate the left and right boundaries of the PPD region, corresponding to the upper limit of $N_\parallel$ and the accessibility limit, respectively.

a logarithmic scale. For better visibility, values close to zero, *i.e.*, $\log(D_{QL}) < -5$, are rendered in white, and the same convention is applied to the subsequent plots.

As shown in the figure, the diffusion coefficients D_{QL} from the simulation are well-contained within these boundaries in velocity space. This boundary consistently aligns with simulation results across most tested cases. These PPD boundaries are subsequently used to mask the U-Net output, thereby enforcing physical restrictions in the final predictions.

3.3 Training Procedure

The proposed surrogate architecture was trained using supervised learning with four normalized input parameters. Prior to training, the input variables were scaled to the range $[0, 1]$, and the target values were log-transformed as $\log(D_{QL})$ to stabilize the dynamic range. The surrogate model was implemented in `PyTorch` and optimized using the Adam optimizer with a learning rate of 1×10^{-3}. The loss function was defined as the mean squared error (MSE), expressed as $\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$, where y_i and $\hat{y}_i$ denote the true and predicted values, respectively.

4 Results

In this section, we first evaluate how well our model predicts D_{QL}, and subsequently examine the impact of integrating the reconstructed D_{QL} into the GENRAY–CQL3D simulation framework.

4.1 Evaluation of D_{QL} Prediction Model Performance

We first assess the performance of the proposed architecture for predicting D_{QL}. The model was trained using the training dataset, while the validation dataset was employed to monitor whether overfitting occurred. As shown in Fig. 5, the training loss steadily decreased as the number of epochs increased, whereas the validation loss stopped improving beyond a certain point. To avoid overfitting, we limited the training to 75 epochs.

We present representative examples of the predicted $\log(D_{QL})$ at a specific radial index ($\rho = 0.75$) from the test dataset, selected according to their R^2 scores. As shown in Fig. 4, the left panel displays the ground truth $\log(D_{QL})$ obtained from simulation, the center panel shows the corresponding prediction by our neural network, and the right panel illustrates the difference between the ground truth and the prediction. The top row corresponds to the sample with the maximum $R^2 = 0.9107$, the bottom row to

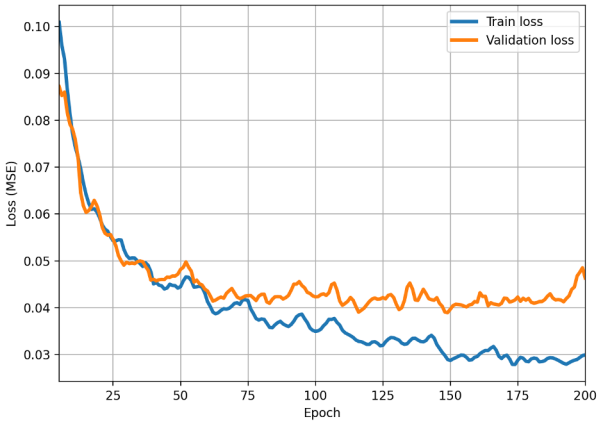

Figure 5. Training and validation loss curves over epochs. To improve visibility, figure is plotted from epoch 5 to 200. The blue curve denotes the training loss, while the orange curve represents the validation loss. MSE is computed on $\log(D_{QL})$.

the sample with the minimum $R^2 = -0.0747$, and the middle row to the sample with the median $R^2 = 0.7252$ within the test dataset. The top-left corner of each panel reports MSE, the root mean squared error (RMSE, defined as $\sqrt{\text{MSE}}$), and the coefficient of determination R^2, computed as

$$R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2},$$

where y_i, $\hat{y}_i$, and $\bar{y}$ denote the ground truth, the prediction, and the mean of the ground truth, respectively. As mentioned above, the median R^2 score over the test dataset was 0.7252, indicating that the model captures the overall trends reasonably well, although some variability remains inaccurate. All metrics were computed using the original scale of D_{QL}, without logarithmic transformation.

In addition to the R^2 score, we evaluated the Structural Similarity Index (SSIM) [17] across the test set. The SSIM is defined as

$$\text{SSIM}(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)},$$

where μ_x and μ_y denote the mean values of the images x and y, σ_x^2 and σ_y^2 their variances, and σ_{xy} their covariance. The constants are given by $C_1 = (K_1 L)^2$ and $C_2 = (K_2 L)^2$, where L is the dynamic range of the data and K_1, K_2 are small scalars. In our implementation, the raw network outputs representing $\log(D_{QL})$ were first converted back to the original scale and then normalized to the range $[0, 1]$ prior to calculating the SSIM. Accordingly, we set $L = 1$ and adopted $K_1 = 0.01$ and $K_2 = 0.03$, consistent with the original definition by Wang [17]. The data were subsequently transformed from polar coordinates $(u/v_{\text{norm}}, \theta)$ into Cartesian coordinates $(u_\parallel/v_{\text{norm}}, u_\perp/v_{\text{norm}})$. During the SSIM calculation, regions affected by coordinate conversion artifacts and non-physical domains masked by the PPD were excluded. The average SSIM score across the test set was 0.909, suggesting that the surrogate model predictions share a generally strong structural similarity with the reference simulation results. However, it should be noted that in the D_{QL}, there exist regions where only certain parts exhibit relatively large values, while most of the remaining areas are characterized by small values. As a result, the SSIM score may be overestimated.

Consequently, the model reproduced the general spatial trends and distribution patterns of the target, despite the high-dimensional and highly nonlinear nature of D_{QL}. However, the accuracy of the predicted values remained limited.

4.2 Validation Through GENRAY-CQL3D Simulation

We further tested our model by integrating the reconstructed D_{QL} of the test set into the GENRAY-CQL3D simulation framework. Specifically, we replaced the original D_{QL} values used in the simulation with those predicted by our neural network, and analyzed the resulting power absorption and current drive profiles. The presented results correspond to the stage where the D_{QL} predicted by our neural network is supplied to CQL3D and executed until sufficient convergence.

Left figure of Fig. 6 shows the ray trajectory for a representative case from the test set, in which rapid absorption of ray power is observed. The corresponding power deposition profile as a function of ρ is shown in top right figure of Fig. 6. In the original GENRAY-CQL3D simulation, the power deposition gradually converges to a peak of 0.5 W/cm^3 at $\rho = 0.5$. When the machine-learned D_{QL} is used, the deposition profile exhibits the same peak location and magnitude but converges significantly faster, demonstrating the computational efficiency of the ML-based approach. The bottom right figure of Fig 6 compares the resulting current profiles obtained from the original simulation and our machine learning model with PPD masking. The two profiles are nearly identical, indicating that the reconstructed D_{QL} preserves the physical fidelity required for accurate current drive prediction. It is worth noting that when the PPD masking is not applied, the resulting current profile shows an artificial increase. This is due to unphysical D_{QL} values outside the physically allowable region, which are filtered by the PPD-based restrictions. In other words, allowing finite D_{QL} to exist outside the PPD region broadens the quasi linear plateau, and the resulting artificial increase in the current profile is generally attributed to this broader plateau.

Fig. 7 shows the difference between the total current calculated by the original GENRAY–CQL3D method and the prediction using the ML model-based D_{QL} for 100 simulation cases from the test set: $\Delta I_{\text{LH}} = \left| I_{\text{LH}}^{\text{target}} - I_{\text{LH}}^{\text{ML}} \right|$, where I_{LH} denotes the total lower-hybrid driven current by wave. As shown in Fig. 7, the model performs well for cases with large $N_\parallel^{\text{start}}$, which typically result in strong lower hybrid wave absorption. In these cases, the predicted total current closely matches the target values, with ΔI_{LH} generally below 0.01 MA. The average R^2 score of the current profile over ρ for strong absorption cases is 0.865, indicating good agreement with the overall reference profile. However, for cases with low $N_\parallel^{\text{start}}$, the perfor-

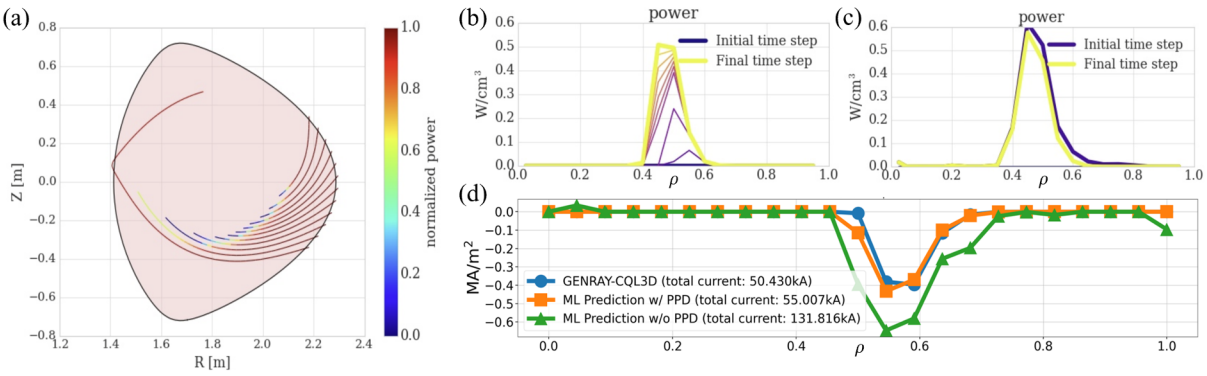

Figure 6. (a): Ray trajectories and corresponding power absorption along the paths. (b): Power deposition profiles along ρ using D_{QL} from the original code and (c): Power deposition profiles along ρ using D_{QL} from machine learning prediction. (d): Comparison of current profiles along ρ between the original GENRAY-CQL3D and the machine learning-based prediction.

mance degrades compared to expectations. In these cases, a possible reason for the low prediction performance might be that, since the wave is weakly absorbed, the ray continues circulating in the poloidal cross-section. In other words, our recursive prediction scheme, which predicts D_{QL} from the outer to the inner region, may not have been effective. The Pearson correlation coefficient was found to be $r = -0.628$, indicating a negative correlation such that ΔI_{LH} decreases as $N_{\parallel}^{\mathrm{start}}$ increases.

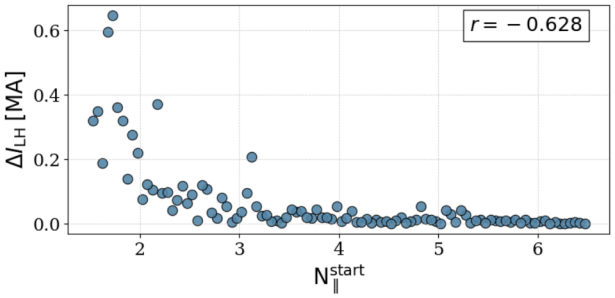

Figure 7. Scatter plot of the difference in total lower hybrid current (ΔI_{LH}) between GENRAY–CQL3D and ML-based predictions. The annotation r denotes the Pearson correlation coefficient. Here $r = \mathrm{cov}(X, Y)/(\sigma_X \sigma_Y)$, where cov is covariance and σ is standard deviation for each variable.

4.3 Comparison of Computational Time Between Conventional and Machine Learning Methods

In this section, we compare the computational efficiency and convergence time between the conventional GENRAY-CQL3D simulation and the proposed machine learning approach.

First, the time required to run GENRAY in the GENRAY-CQL3D simulation was measured for the test dataset. Generally, GENRAY is continuously updated until the injected power is sufficiently absorbed. Therefore, cases with smaller $N_{\parallel}^{\mathrm{start}}$ and lower T_{e0}, corresponding to weak absorption scenarios, require more time, whereas the opposite cases are calculated faster. On average for the test dataset, approximately 87.53 seconds were required. In contrast, the proposed machine learning method shows almost no dependency on input parameters and requires only about 3.46 seconds on average to compute D_{QL}. This demonstrates a significant advantage in estimating D_{QL} using the machine learning approach.

Second, the convergence time for the power deposition profile was evaluated. Convergence was assessed based on the relative change in the power deposition profile between successive iterations, defined as $\Delta = \|P^{(n)} - P^{(n-1)}\|_2 / \|P^{(n-1)}\|_2$, and was considered converged when $\Delta < 10^{-3}$. Here, P^n is power deposition profile at time step n. In the conventional approach, after D_{QL} is introduced at the first time step, it is iteratively updated along with the power absorption profile until convergence is achieved at each subsequent time step. However, in the proposed method, a pre-converged D_{QL} is directly provided at every time step. As a result, while the traditional CQL3D simulation typically required around 10 steps (averaging 77.12 seconds in the test dataset) to converge, convergence was usually achieved within 3 steps in the proposed method, taking approximately 26.77 seconds.

5 Conclusion

In this study, we developed a machine learning-based surrogate model to reconstruct the generalized quasilinear diffusion coefficient (D_{QL}) for Lower Hybrid Current Drive (LHCD) in a tokamak. By using physically motivated restrictions derived from cold plasma theory and bounce-averaged damping effects, it was able to reproduce overall trends of D_{QL} while offering a notable reduction in computational cost compared to traditional GENRAY–CQL3D simulations.

The evaluations indicated that our recurrent U-Net architecture tended to preserve the spatial structure and patterns inherent in D_{QL}. The model achieved a median R^2 score of 0.7304 on the test dataset, suggesting a reasonable level of predictive capability. In addition, an average

Structural Similarity Index (SSIM) score of 0.909 across the test dataset suggests that the model generally maintained structural fidelity, though some limitations remain.

Furthermore, for the strong absorption cases, validation through direct integration of the reconstructed D_{QL} into the GENRAY–CQL3D framework demonstrated consistent and accurate predictions of both power deposition and current profiles. In particular, the machine learning–based reconstructions closely reproduced the original simulation results, underscoring the efficacy of the proposed surrogate model as a computationally efficient alternative to conventional numerical simulations.

While the model produced promising results for strong absorption cases, further investigation is needed to improve performance in weak absorption scenarios. Future work will explore enhancements in network architecture and training methodologies to address these limitations and further generalize the model's predictive capabilities.

Moreover, we plan to investigate a hybrid approach that integrates machine learning-predicted values of D_{QL} into full-wave simulations. In this scheme, the ML-predicted D_{QL} will be used as the initial guess for the full-wave–Fokker–Planck collision-coupled code. We will examine whether the solution converges within a small number of iterations (*e.g.*, three time steps), as observed in this paper. Since full-wave simulations need to be executed repeatedly at each time step during the Fokker–Planck evolution, achieving fast convergence of D_{QL} using ML is expected to significantly reduce the overall computational cost.

Acknowledgements

This work was supported by the U.S. Department of Energy, Office of Fusion Energy Sciences under Award Number DE-SC0021202.

References

[1] N. J. Fisch,, Theory of current drive in plasmas, Reviews of Modern Physics **59**, 175 (1987). https://doi.org/10.1103/RevModPhys.59.175

[2] M. Brambilla, Numerical simulation of ion cyclotron waves in tokamak plasmas, Plasma Physics and Controlled Fusion **41**, 1 (1999). https://doi.org/10.1088/0741-3335/41/1/002

[3] A. P. Smirnov, and R. W. Harvey, The GENRAY ray tracing code, CompX Report CompX-2000-01 (2001).

[4] R. W. Harvey and M. G. McCoy, The CQL3D Fokker–Planck code, IAEA Technical Committee Meeting on Simulation and Modeling of Thermonuclear Plasmas, 489–526 (1992).

[5] Á. Sánchez-Villar, Z. Bai, N. Bertelli, E. W. Bethel, J. Hillairet, T. Perciano, S. Shiraiwa, G. M. Wallace and J. C. Wright, Machine learning enhanced predictions of ICRF heating: Overcoming numerical limitations via data curation, Physics of Plasmas **32**, 062504 (2025), https://doi.org/10.1063/5.0268341

[6] Á. Sánchez-Villar, Z. Bai, N. Bertelli, E. W. Bethel, J. Hillairet, T. Perciano, S. Shiraiwa, G. M. Wallace and J. C. Wright, Real-time capable modeling of ICRF heating on NSTX and WEST via machine learning approaches, Nuclear Fusion **64**, 096039 (2024). https://doi.org/10.1088/1741-4326/ad645d

[7] G. M. Wallace, Z. Bai, R. Sadre, T. Perciano, N. Bertelli, S. Shiraiwa, E. W. Bethel, and J. C. Wright, Towards fast and accurate predictions of radio frequency power deposition and current profile via data-driven modelling: applications to lower hybrid current drive, Journal of Plasma Physics, **88**, 895880401 (2022). https://doi.org/10.1017/S0022377822000708

[8] R. Prater, D. Farina, Yu. Gribov, R. W. Harvey, A. K. Ram, Y.-R. Lin-Liu, E. Poli, A. P. Smirnov, F. Volpe, E. Westerhof, et al., Benchmarking of codes for electron cyclotron heating and electron cyclotron current drive under ITER conditions, Nuclear Fusion **48**, 035006 (2008). https://doi.org/10.1088/0029-5515/48/3/035006

[9] R. L. Iman, Latin hypercube sampling, John Wiley & Sons, Ltd (2008).

[10] T. H. Stix, Waves in plasmas, Springer Science & Business Media (1992).

[11] X. M. Zhai et al., Theoretical analysis of key factors achieving reversed magnetic shear q-profiles sustained with lower hybrid waves on EAST, Plasma Physics and Controlled Fusion **61**, 045002 (2019). https://doi.org/10.1088/1361-6587/aaffe5

[12] O. Ronneberger et al., U-Net: Convolutional networks for biomedical image segmentation. In: MICCAI 2015, Part III, Lecture Notes in Computer Science **9351**, 234 (2015). https://doi.org/10.1007/978-3-319-24574-4_28

[13] G. D. Kerbel, and M. G. McCoy, Kinetic theory and simulation of multispecies plasmas in tokamaks excited with electromagnetic waves in the ion cyclotron range of frequencies, Physics of Fluids **28**, 3629 (1985). https://doi.org/10.1063/1.865319

[14] F. Paoletti et al., Lower hybrid current drive accessibility study with reconstructed magnetic equilibria in PBX-M, Nuclear Fusion **34**, 771 (1994). https://doi.org/10.1088/0029-5515/34/6/I01

[15] P. Bonoli, Linear theory of lower hybrid heating, *IEEE Transactions on Plasma Science* **12**, 95–107 (1984). https://doi.org/10.1109/TPS.1984.4316302

[16] M. Brambilla, Waveguide launching of lower hybrid waves, Nuclear Fusion **19**, 1343 (1979). https://doi.org/10.1088/0029-5515/19/10/006

[17] Z. Wang et al., Image quality assessment: from error visibility to structural similarity, IEEE Transactions on Image Processing **13**, 600–612 (2004). https://doi.org/10.1109/TIP.2003.819861