Lawrence Berkeley National Laboratory

LBL Publications

Title

Physics-constrained neural networks for real-time reconstruction of radiated power in DIII-D plasma control system

Permalink

https://escholarship.org/uc/item/8c07x49r

Journal

Review of Scientific Instruments, 97(11)

ISSN

0034-6748

Authors

Odstrcil, T

Bai, Z

Farre-Kaga, HJ

et al.

Publication Date

2026-11-01

DOI

10.1063/5.0348668

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NoDerivatives License, available at https://creativecommons.org/licenses/by-nd/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

ACCEPTED MANUSCRIPT
Review of Scientific Instruments
AIP Publishing
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.
PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0348668

Physics-Constrained Neural Networks for Real-Time Reconstruction of Radiated Power in DIII-D Plasma Control System

T. Odstrcil,[1] Z. Bai,[2] H. J. Farre-Kaga,[3] W. Choi,[1] D. Eldon,[1] B. Sammuli,[1] and A. Moser[1]
[1)]*General Atomics, San Diego, CA, USA*
[2)]*Lawrence Berkeley National Laboratory, Berkeley, CA, USA*
[3)]*Princeton Plasma Physics Laboratory, Princeton, NJ, USA*

(*Electronic mail:odstrcilt@fusion.gat.com)

(Dated: 29 September 2026)

A physics-constrained neural network is developed to predict radiated power across multiple plasma regions from foil bolometer measurements using a limited set of real-time equilibrium features. The surrogate model enforces a linear mapping between measured brightness and radiated power, enabling efficient inference with microsecond latency, suitable for real-time deployment. Model training is performed on a combination of large synthetic and real experimental datasets, allowing the neural network to learn the mapping between the measured signal and radiated power across varying plasma shapes while accounting for experimental noise, thereby improving its robustness. The model is optimized to estimate near-axis radiated power, enabling early detection of ongoing impurity accumulation. The approach has been implemented in the DIII-D plasma control system, demonstrating its capability for real-time operation.

I. INTRODUCTION

Early detection of impurity accumulation is crucial for maintaining stable plasma operation, as timely intervention such as electron heating or other countermeasures[1,2] can mitigate the onset of deleterious effects. Although impurity accumulation is not presently a major concern in the DIII-D tokamak with a carbon wall, the planned transition to a metal wall is expected to significantly increase the risk of core impurity buildup, necessitating reliable and rapid detection for future operation. Bolometric diagnostics provide the most direct measurements of ongoing impurity accumulation, as near-axis radiation, rather than impurity density alone, serves as the primary indicator of this process. However, bolometer signals are line-integrated, and resistive foil bolometer signals require deconvolution in time to obtain instantaneous radiated power. As a result, reconstructing local emissivity or region-resolved radiated power constitutes a challenging, ill-posed inverse problem that is highly sensitive to experimental noise[3]. Althought the other diagnostics, like soft X-ray also provides measurement of impurity emission, their interpretation strongly depends on the electron temperature and their are thus less suitable for realtime plasma control.

Conventional tomographic reconstruction methods[4–6] and more recent approaches based on Physics-Informed Neural Networks (PINNs)[7] are computationally expensive, making them unsuitable for real-time plasma control due to the high cost of the underlying inverse problem. Alternative approaches based on Convolutional Neural Networks (CNNs)[8–10] can reconstruct the full two-dimensional radiation profile below 10 ms. However, these methods typically require additional post-processing using magnetic equilibrium information to determine the radiated power within regions of interest. Another approach is the linear surrogate model[11], which learns a linear mapping from bolometer measurements to the radiation image. For a fixed spatial mask, this mapping reduces to a vector inner product, enabling real-time monitoring of plasma radiation. However, none of these fast approaches[8–11] explicitly incorporates magnetic equilibrium information into the reconstruction, despite its importance for obtaining reliable results from the limited number of bolometer lines of sight.

In this work, we propose a novel method based on a physics constrained neural network (PCNN) that directly estimates the radiated power from major plasma regions without explicit tomographic reconstruction. This proposed method achieves microsecond-level latency, improved noise robustness and and agreement with the conventional tomography code *pyTomo*[6], enabling practical deployment for real-time plasma control in tokamak operation.

This article is structured as follows. Section II describes the PCNN architecture, including the incorporation of physics constraints and the associated training procedures. Section III describes the construction of the training databases used for supervised and unsupervised learning. Finally, Section IV evaluates the proposed approach by comparing the radiated powers estimates from the PCNN with those obtained from offline conventional tomography reconstruction and a legacy method based on fixed geometry-derived weighting.

II. NEURAL NETWORK FOR RADIATED POWER ESTIMATION

As outlined in the introduction, the primary objective of the proposed PCNN is to directly estimate the radiated power within predefined regions of interest (ROIs), bypassing the computationally intensive reconstruction of a full two-dimensional emissivity profile. This reduced, low-dimensional mapping enables the PCNN architecture to utilize an efficient multilayer perceptron[12] (MLP) architecture that satisfies real-time control latency requirements. However, the model must also meet additional constraints arising from both the underlying physics of this problem, as well as the demands of real-time plasma control.

The forward tomographic problem is governed by a Fredholm integral equation of the first kind, which maps the total

10 November 2026 16:03:57

ACCEPTED MANUSCRIPT
Review of Scientific Instruments
AIP Publishing
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.
PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0348668
10 November 2026 16:03:57

plasma emissivity ε_{tot} to the diagnostic measurements:

$$b_i = \frac{4\pi}{\langle A\Omega \rangle_i} \int_{\text{LOS}_i} \varepsilon_{\text{tot}}(R,Z)\, dV, \tag{1}$$

where b_i represents the line-integrated brightness (W/m^2) measured along the i-th line-of-sight (LOS_i), $\langle A\Omega \rangle_i$ denotes the geometric etendue of the respective detector. The radiated power within a specified region of interest (P_{ROI}) is subsequently defined as the volume integral of this local emissivity:

$$P_{\text{ROI}} = \int_{\text{ROI}} \varepsilon_{\text{tot}}(R,Z)\, 2\pi R\, dR\, dZ. \tag{2}$$

The boundaries of the ROI, such as the core, and the upper/lower divertor, ..., depend on the magnetic equilibrium of plasma. Consequently, a linear relationship exists between the measured $\mathbf{b}$ and power P_{ROI}, but only for fixed plasma geometry. Furthermore, in the absence of experimental noise, brightness and power P_{ROI} are both non-negative. These geometry-dependent linearity and non-negativity constraints form the foundational physics that motivate the architecture of the PCNN developed in this work.

A. Network architecture

The model is designed to satisfy both physics constraints (i–iii) and real-time deployment requirements (iv–vii):

i) Radiated power is a linear function of the line-integrated brightness
ii) Radiated power is non-negative, i.e., $P_{\text{ROI}} \geq 0$
iii) Power estimates are consistent across overlapping regions
iv) Robustness to experimental noise
v) Dependence only on brightness $\boldsymbol{b}$ and real-time equilibrium parameters
vi) Low memory footprint ($< 200\,\text{kB}$)
vii) Ability to exclude broken bolometer channels without retraining

Requirements i) and ii) are mutually incompatible; therefore, positivity is enforced via the loss function rather than by the PCNN architecture. Linearity at a fixed plasma geometry implies that the PCNN depends only on plasma shape parameters, while radiated power is obtained as a linear combination of brightness with weights predicted by the PCNN. Plasma shape parameters are provided by realtime equilibrium reconstruction code rtEFIT[13]. A low memory footprint is required because keras2c[14], tool converting Keras neural network models into standalone C99 code is used to integrate the PCNN inside the DIII-D e plasma control system[15] (PCS), and the network is stored in stack memory[16]. This is not a fundamental limitation and can be overcome in the future for more complex models.

Let $\boldsymbol{x} \in \mathbb{R}^{N_{\text{EFIT}}}$ denote rtEFIT shape parameters and $b \in \mathbb{R}^{N_{\text{ch}}}$ the bolometer brightness. The radiated power $P \in \mathbb{R}^{N_{\text{power}}}$ is modeled as

$$P(\boldsymbol{x},\boldsymbol{b}) = \mathbf{W}(\boldsymbol{x})\,\tilde{\boldsymbol{b}}, \quad \tilde{\boldsymbol{b}} = \mathbf{A}b, \tag{3}$$

where $\mathbf{A} \in \mathbb{R}^{N_{\text{low}} \times N_{\text{ch}}}$ is a low-rank projection matrix and $\tilde{b} \in \mathbb{R}^{N_{\text{low}}}$ is a reduced representation of the measurements.

The network predicts a coupling matrix $\mathbf{W}(x) \in \mathbb{R}^{N_{\text{power}} \times N_{\text{low}}}$ using an MLP with a linear offset:

$$\mathbf{W}(x) = W_0 + W_{\text{lin}}x + \text{MLP}(x), \tag{4}$$

where W_0 is a constant offset and W_{lin} a linear mapping. The MLP is a fully connected feed-forward network with ReLU activations defined as $\text{ReLU}(x) = \max(0,x)$, and three hidden layers, with output reshaped to $\mathbb{R}^{N_{\text{power}} \times N_{\text{low}}}$. The full model reads

$$P(\boldsymbol{x},\boldsymbol{b}) = \left[W_0 + W_{\text{lin}}x + \text{MLP}(\boldsymbol{x})\right](\mathbf{A}\boldsymbol{b}). \tag{5}$$

Fig. 1 illustrates the complete network structure. The projection matrix $\mathbf{A}$ improves robustness to noise and reduces the size of the last hidden layer. It also enables handling of missing channels gracefully: by taking a pseudo-inverse, zeroing the corresponding rows, and recomputing the pseudo-inverse, a new projection operator is obtained that yields the least-squares optimal low-dimensional representation without retraining. The motivation behind the use of the linear residual branch is to improve extrapolation beyond the training set regime. All network parameters, including $\mathbf{A}$, are learned during the training process.

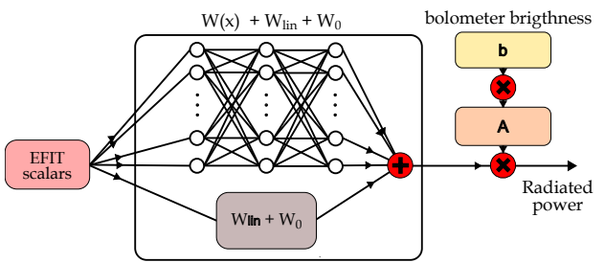


FIG. 1. Schematic of the neural network architecture for predicting radiated power from plasma shape parameters and line-integrated brightness $\boldsymbol{b}$.

B. Physics-based constraints

The PCNN can be trained for multiple different objectives by adding appropriate terms in the loss function. For example, positivity ii) and consistency iii) are not enforced through the network architecture, but are instead imposed as training constraints during optimization. The composite loss function $\mathscr{L}$ for training is thus defined as follows:

$$\mathscr{L} = \mathscr{L}_{\text{data}} + \alpha\,\mathscr{L}_W + \beta\,\mathscr{L}_{\text{neg}} + \gamma\,\mathscr{L}_{\text{con}}, \tag{6}$$

where α, β, γ are weighting coefficients and the individual

ACCEPTED MANUSCRIPT
Review of Scientific Instruments
AIP Publishing
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.
PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0348668

terms are defined as

$$\begin{aligned}
\mathscr{L}_{\text{data}} &= \text{MSE}\left(\frac{P(\boldsymbol{x},\boldsymbol{b}) - P_{\text{train}}}{P_{\text{mean}}}\right), \\
\mathscr{L}_{W} &= \text{MSE}(\mathbf{W}(x)), \\
\mathscr{L}_{\text{neg}} &= \text{MSE}(\max\{-P(\boldsymbol{x},\boldsymbol{b}),0\}), \\
\mathscr{L}_{\text{con}} &= C_{\text{ldiv}} + C_{\text{udiv}} + C_{\text{tot}}. \\
C_{\text{ldiv}} &= \text{MSE}(P_{\text{ldiv}} - P_{\text{ldivi}} - P_{\text{ldivi}}) \\
C_{\text{udiv}} &= \text{MSE}(P_{\text{udiv}} - P_{\text{udivi}} - P_{\text{udivi}}) \\
C_{\text{tot}} &= \text{MSE}(P_{\text{tot}} - P_{\text{ldiv}} - P_{\text{udiv}} - P_{\text{core}} - P_{\text{SOL}})
\end{aligned}$$

Here, $\mathscr{L}_{\text{data}}$ is mean square error (MSE) of power normalized by mean value over the database P_{mean}, $\mathscr{L}_W$ regularizes the coupling matrix $\mathbf{W}$, $\mathscr{L}_{\text{neg}}$ enforces positivity, and $\mathscr{L}_{\text{con}}$ enforces consistency. Specifically, C_{ldiv} and C_{udiv} penalize discrepancies between inner/outer divertor components and total divertor power, while C_{tot} enforces consistency between total power and the sum of the core, scrape-of-layer (SOL), and both divertors. Location of these radiative regions is defined in Sec. III A.

Predicting aggregated quantities (P_{tot}, P_{ldiv}, P_{udiv}) directly, rather than calculating them as a sum of individual terms, improves accuracy, as larger regions can be estimated more reliably than smaller, localized contributions.

III. TRAINING OF ARTIFICIAL NEURAL NETWORK

The PCNN was trained simultaneously on labeled and unlabeled datasets. Supervised training used a synthetic database generated from radiation patterns reconstructed by an offline tomography algorithm, as described below. In parallel, training was also performed on an unlabeled dataset consisting of experimentally measured line-integrated brightness. Although these data lack ground-truth labels, they still provide constraints on the solution through positivity and consistency requirements and, importantly, enable the model to learn realistic channel-specific noise and characteristics across a large set of experimental discharges.

A. Definition of radiative regions

The PCNN predicts radiated power in predefined regions rather than reconstructing a full 2D emissivity profile, which is more suitable for real-time control purposes. The regions depend on the plasma equilibrium and are illustrated in Fig. 2. The near-axis region power P_{axis} (normalised poloidal flux $\Psi_N < 0.2$) is used for the detection of impurity accumulation. The core region P_{core} ($\Psi_N < 0.9$) includes the axis and is relevant for power-over-separatrix estimate and core radiative fraction control. The edge region P_{SOL} ($0.9 < \Psi_N < 1.2$), dominated by low-Z radiation, is treated separately as the SOL. The lower (upper) divertor P_{ldiv}, (P_{udiv}) is defined as the region below (above) a horizontal line 20 cm from the corresponding X-point. Inner and outer strike point powers are obtained by splitting the divertor along a line connecting the X-point and magnetic axis. The total radiated power P_{tot} corresponds to emission within the vessel.

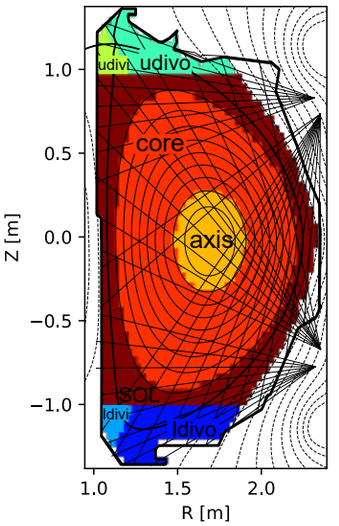


FIG. 2. Definition of radiative regions for an example upper single-null discharge. The plasma cross-section with poloidal flux contours (dashed lines) is overlaid with DIII-D bolometer lines of sight.

B. Synthetic training database

The synthetic database was constructed from reconstructions obtained using the fast 2D tomography code *pyTomo*[6] with bolometer regularization developed for JET[17]. A set of 149 representative discharges between 193831–205055 was selected using k-means clustering to span a broad range of plasma shapes without any additional constrains for selected scenarios. These reconstructions provided 35 006 time slices, which are, however, strongly correlated within each discharge. Plasma shape input for PCNN included lower and upper X-points location, magnetic axis location, lower and upper triangularity τ_{top}, τ_{bot}, elongation κ, minor radius a, and *dRsep* parameter mid-plane radial distance between the separatrixes of the upper and lower divertor X-points. All of these eleven parameters are provided by rtEFIT in real-time.

The diversity of plasma shapes spanned by the training database was increased by random stretching and shifting of the original tomograms, along with the corresponding poloidal flux matrices and proper transformation of plasma shape parameters. Each sample was augmented by random spatial transformations: horizontal and vertical shifts within ± 10 cm and scaling by a factor of 0.9-1.1, obtaining a new profile $\tilde{\varepsilon}$. Emission outside the vessel was masked to avoid unphysical profiles.

Since this PCNN is intended to be applied for radiated power measurement also with an anticipated metal-wall operation of DIII-D and there are no high-Z radiation profiles in the training database, we have decided to add an artificial tungsten radiation component ε_W:

$$\varepsilon_W = c_W n_e^2 L_W(T_e) \exp\left[\lambda(R^2 - R_0^2)\right], \tag{7}$$

$$c_W(\Psi_N) = c_{W,0}\left[1 + \eta \exp(-\Psi_N/\sigma^2)\right], \tag{8}$$

10 November 2026 16:03:57

ACCEPTED MANUSCRIPT
Review of Scientific Instruments
AIP Publishing
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.
PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0348668

with $c_{W,0} \sim \log\mathcal{N}(-9,1)$, $\eta \sim \log\mathcal{N}(0,1)$, and $\sigma \sim \log\mathcal{N}(0,1)$, where $\log\mathcal{N}$ denotes a log-normal distribution. Parameters $c_{W,0}$, η and σ were defined to produce realistically looking ε_W profiles. L_W is a W cooling factor evaluated at experimentally measured electron temperature T_e profile and the exponential term on the right side of Eq. (7) represents a centrifugal asymmetry[18], where the coefficient λ is function of measured plasma rotation, temperatures and effective charge.

In the final step of data augmentation, the total emissivity $\varepsilon_{\text{tot}} = \tilde{\varepsilon} + \varepsilon_W$ was independently scaled in each region (Fig. 2) by $\log\mathcal{N}(0,1)$. This augmentation was repeated $1000\times$ per each spatially transformed sample, resulting in $\sim 3.5\times10^7$ samples. Since every 1000 samples share identical plasma shape parameters, the MLP output remains the same, keeping training computationally efficient. Synthetic brightness was computed via the geometry matrix $\mathbb{T}$ from Cherab code[19], enabling efficient evaluation without storing full 2D emissivity profiles. Gaussian noise was added with a variance estimated from the difference between the measured and reconstructed brightness.

C. Experimental training database

To better capture real noise characteristics, the training set was extended with an unlabeled experimental dataset comprising all DIII-D discharges longer than 4 s with available real-time bolometer data (i.e. shots 193831–205055). The resulting dataset contains 4841 discharges and 4.8×10^6 samples. Obtaining radiated power via off-line tomography for this number of discharges is not practically feasible. The signals were smoothed using a first-order recursive filter with a 5 ms time constant and subsequently downsampled to 50 ms. Although unlabeled, these data contribute to training through positivity and consistency constraints and provide realistic noise statistics representative of a large set of discharges.

D. Network training

To train the network shown in Fig. 1, the dataset was partitioned into 90% training and 10% testing subsets. The partitioning was performed at the pulse level to prevent data leakage, ensuring that samples from the same discharge do not appear in both sets. The network architecture was constrained by the PCS memory requirement (vi); hidden layer sizes of $(96,128,64)$ and $N_{\text{low}} = 16$ provided a good trade-off between accuracy and model size.

The loss function parameters (Eq. 6) were empirically chosen, $\alpha = 10^{-5}$, $\beta = 0.1$, and $\gamma = 10^{-4}$ to balance enforcement of physical constraints with predictive accuracy. Training was performed in Python using the JAX framework[20] and the Optax optimization library, employing the AdamW optimizer[21] with a learning rate of 10^{-3}. Batch sizes were limited by available GPU memory to 2048×1000 synthetic samples and 2.9×10^5 experimental samples per epoch. JAX provided automatic differentiation, just-in-time (JIT) compilation, and efficient GPU acceleration.

Training was run for 3000 epochs until the validation loss plateaued (Fig. 3). The training loss continued to decrease, indicating that further training would result in overfitting. With an epoch time of ~ 4.2 s on an NVIDIA RTX A5000 Laptop GPU, the total training time was under 4 hours. Single-threaded inference of the PCNN on a 13th Gen Intel Core i7-13700H processor requires approximately $10\,\mu$s, well below the $50\,\mu$s control cycle latency of the DIII-D PCS, making the model suitable for real-time deployment.

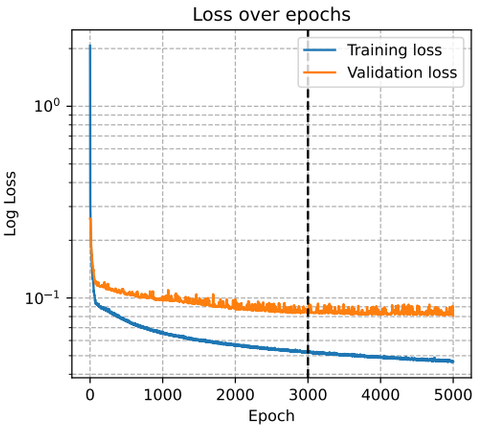


FIG. 3. Training and validation loss during neural network optimization using the AdamW algorithm. Validation curve after $\sim$ 3000 epochs plateaued.

IV. RECONSTRUCTION OF RADIATED POWER

The accuracy of the PCNN is illustrated in Fig. 5 by histograms comparing the inferred power from noise-free synthetic data with the corresponding ground-truth power. The coefficient of determination (R^2) and symmetric mean absolute percentage error (sMAPE) are used to quantify the reconstruction accuracy in each region. For inferred and reference datasets x and $\tilde{x}$, respectively, sMAPE is defined as

$$sMAPE = \frac{2\sum_i |x_i - \tilde{x}_i|}{\sum_i (|x_i| + |\tilde{x}_i|)}. \tag{9}$$

The sMAPE is as low as 2% for the total radiated power and 5% for the core radiated power. In contrast, the largest errors are observed in the divertor regions, with sMAPE values of 16% for the upper divertor, 40% for the inner strike point, and 24% for the outer strike point.

For the experimental data, where ground truth is not available, performance is assessed via the positivity constraint (Tab. I). High positivity fractions are obtained for large regions, while lower values are observed for the upper divertor and P_{axis}. This is expected, as the radiated power in these regions is often small and close to the noise level.

TABLE I. Positivity fraction on experimental dataset

Quantity	$P>0$	Quantity	$P>0$
P_{tot}	99.9%	P_{ldivi}	95.1%
P_{core}	99.2%	P_{ldivo}	94.4%
P_{axis}	86.7%	P_{udiv}	84.8%
P_{sol}	85.1%	P_{udivi}	75.0%
P_{ldiv}	93.1%	P_{udivo}	85.1%

10 November 2026 16:03:57

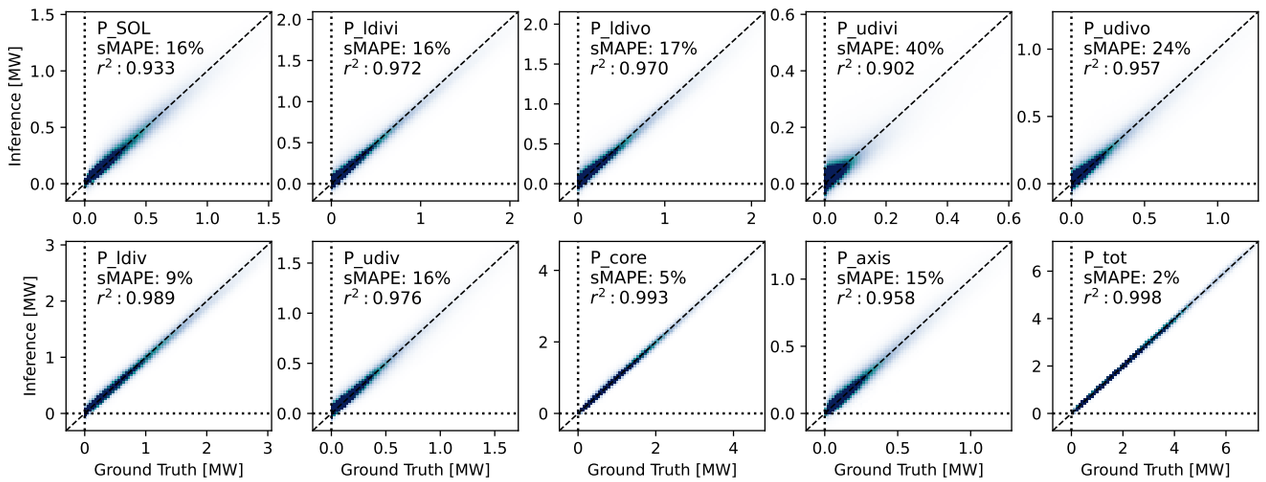


FIG. 4. Histograms of radiated power inferred by the PCNN from noise free synthetics data are compared with the ground true values.

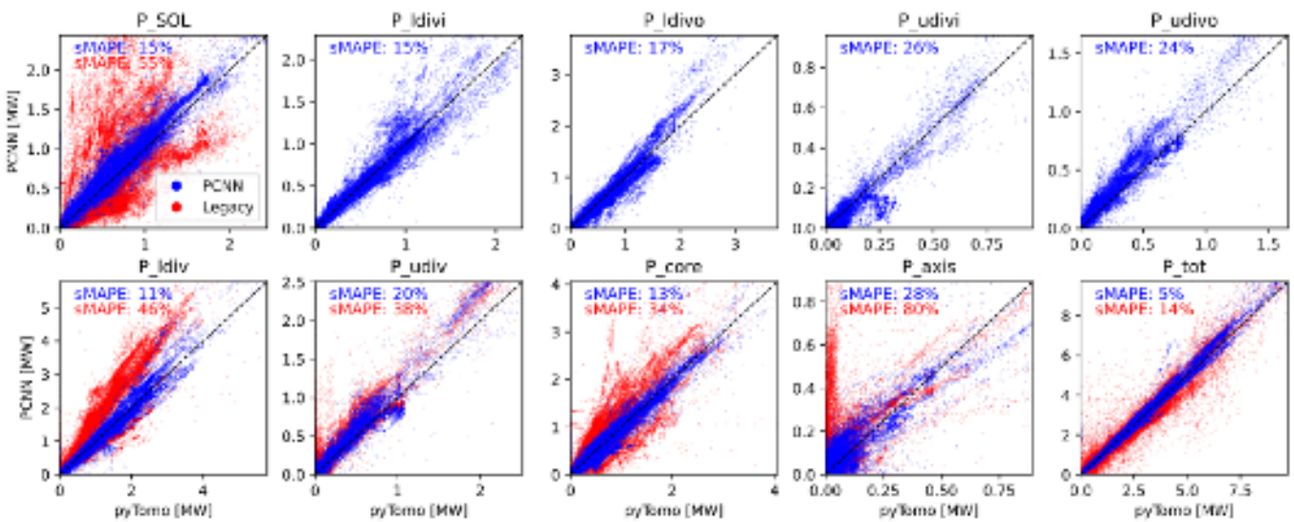


FIG. 5. Subplots show radiated power inferred by the PCNN (blue) and legacy method (red), compared with power obtained by integrating 2D *pyTomo* tomograms over the radiative regions. Individual points represent measured experimental brightness from the original dataset in Sec. III B. Error is estimated using sMAPE defined by Eq. (9)

A. Validation on Experimental Data

Model performance on experimental data is evaluated in comparison with the conventional tomographic reconstruction from the *pyTomo* code, used to construct the augmented synthetic dataset. Importantly, the original experimental measurements and corresponding *pyTomo* reconstructions were never directly used during training, as the synthetic data were generated by re-projecting already modified radiation profiles using the projection matrix $\mathbb{T}$, discarding original measurements. Note that the trained PCNN should be considered an independent method for power reconstruction, sharing the same forward model (via the geometry matrix $\mathbb{T}$ used to create synthetic measurements), rather than a surrogate for the *pyTomo* code.

Furthermore, the PCNN inference is based on real-time bolometer signals[22,23], which must rely on simple causal signal processing for deconvolution of bolometer signals, introducing ~ 40 ms lag, whereas offline tomography employs more robust non-causal filtering. The *pyTomo* reconstruction is therefore not treated as the ground truth, but rather as a reference, since it is also affected by uncertainties inherent to the ill-posed inversion problem.

For comparison, a "legacy" method is plotted. This approach uses fixed weights for P_{tot}, P_{ldiv}, and P_{udiv}, assuming certain distance between detector and radiative center, and a 1D spline-based tomographic inversion for P_{core}, P_{axis}, and P_{SOL}[5]. However, only the weight-based signals were available in the PCS for plasma control[22].

The results are shown in Fig. 5, where PCNN and legacy predictions are compared against *pyTomo*. Both methods accurately reproduce total radiated power, with the PCNN exhibiting significantly lower scatter. This improvement is attributed to the learned channel weighting, which suppresses noise contributions, in contrast to the geometry-based

ACCEPTED MANUSCRIPT
Review of Scientific Instruments
AIP Publishing
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.
PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0348668
10 November 2026 16:03:57

ACCEPTED MANUSCRIPT
Review of Scientific Instruments
AIP Publishing
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.
PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0348668

10 November 2026 16:03:57

weights, unaware of experimental uncertainty, used in the legacy approach. A significant systematic discrepancy was found in the legacy method calculation of a lower divertor power, where the power is overestimated by about 50%, while the PCNN provides very good agreement. The reduced accuracy in upper divertor power estimation is partly attributable to a low P_{udiv} in lower single-null discharges. Discharges in USN shape with $P_{udiv} > 1$ MW exhibit good agreement with the reference. P_{core} is slightly underestimated relative to *pyTomo*, particularly at low powers. Finally, P_{axis} shows increased scatter, reflecting both the inherent difficulty of the measurement and the low power magnitude. Nevertheless, the model maintain good agreement at higher radiation values, which are the most critical for early detection of impurity accumulation.

B. Detection of near-axis W accumulation

The primary motivation for training the PCNN is to provide a reliable signal for early detection of impurity accumulation. When identified at an early stage, impurity buildup can be efficiently suppressed by electron cyclotron (ECH) heating, provided such on-axis heating is possible. The performance of the neural network was tested on an ELM-free, wide-pedestal QH-mode discharge with slow tungsten accumulation. Tungsten slowly migrated from the wall, where it was deposited in previous experiments, into the plasma core. Accumulation started shortly after the LH transition at 1.0s and continued till the end of the discharge despite near-axis ECH heating. Core radiated power from PCNN closely follows the power from *pyTomo* code. Small difference in the near-axis power is within the uncertainty of the tomographic reconstruction.

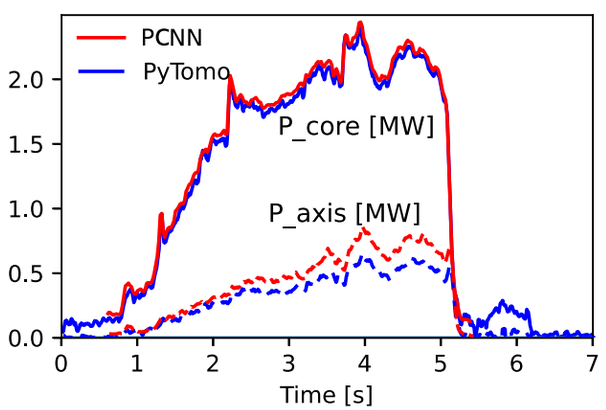


FIG. 6. Time trace of core radiated power ($\Psi_N < 0.9$) and near axis radiation inside of $\Psi_N < 0.2$ during slow tungsten accumulation in discharge 190549. Power inferred by PCNN is in red and from conventional tomography by *PyTomo* code in blue

V. DISCUSSION

The actual accuracy of the predicted power is difficult to quantify because the ground truth is unavailable for experimental data. The trained PCNN shows good agreement with both the synthetic dataset and the conventional tomographic algorithm *pyTomo*. However, systematic errors in the LOS geometry, etendue, and absolute bolometer calibration are shared by both methods because they use the same forward model, represented by the geometry matrix $\mathbb{T}$, and the same calibrated measurements. Moreover, PCNN predictions depend on the procedure used to generate the synthetic training dataset, which augments radiation profiles obtained from *pyTomo*. Since tomographic inversion based on only two views and 48 channels on DIII-D is an ill-posed problem, the reconstructed radiation profiles are not unique and depend on the choice of prior.

Physics-based phantoms, for example radiation profiles calculated using plasma-edge models such as SOLPS[24], could reduce biases introduced by the priors used in conventional tomography. However, the scarcity of detailed edge simulations currently makes them unsuitable for constructing a sufficiently large and diverse training database.

For practical applications, the sMAPE values obtained from the synthetic validation in Fig. 5 provide an estimate of the reconstruction error associated with the PCNN, but should not be interpreted as a complete measure of the uncertainty in experimental radiated power.

VI. SUMMARY

In this work, we have developed an physics constrained neural network model for real-time calculation of the radiated power from foil bolometers on the DIII-D tokamak. Based on plasma shape parameters, the model predicts linear weights based on actual plasma shape and applies them to line-integrated brightness to derive radiation in multiple regions of interest. Because the weights are trained on a mixed synthetic and real experimental dataset, they are aware of typical experimental noise in each channel and produce lower noise and higher-accuracy signals than the currently implemented algorithm based on fixed-geometry weights. High signal-to-noise is particularly valuable for reliable measurement of near-axis radiation, required for early detection of impurity accumulation. The algorithm is integrated into the DIII-D PCS system and will be combined with an ECH controller for real-time radiation control.

ACKNOWLEDGMENTS

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award(s) DE-FC02-04ER54698, DE-SC0014264 and DE-AC02-09CH11466. Z.B. acknowledges support from the U.S. Department of Energy, Office of Science, SciDAC/Advanced Scientific Computing Research under Award Number DE-AC0205CH11231. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy User Facility.

Disclaimer: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency

ACCEPTED MANUSCRIPT
Review of Scientific Instruments
AIP Publishing
This is the author's peer reviewed, accepted manuscript. However, the online version of record will be different from this version once it has been copyedited and typeset.
PLEASE CITE THIS ARTICLE AS DOI: 10.1063/5.0348668

thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DATA AVAILABILITY STATEMENT

The data supporting this study's findings and the source code of the algorithm are available from the corresponding author upon reasonable request.

[1]M. Cappelli, A. Cardinali, V. Zotta, G. Pucella, M. Brambilla, S. Gabriellini, R. Gatto, M. Zerbini, L. Garzotti, D. Van Eester, *et al.*, Fusion Engineering and Design **224**, 115555 (2026).
[2]R. Dux, C. Giroud, R. Neu, A. Peeters, J. Stober, K.-D. Zastrow, A. U. Team, *et al.*, Journal of nuclear materials **313**, 1150 (2003).
[3]P. C. Hansen, Inverse problems **8**(6), 849 (1992).
[4]L. Ingesson, B. Alper, B. Peterson, and J.-C. Vallet, Fusion science and technology **53**(2), 528 (2008).
[5]A. Leonard, W. Meyer, B. Geer, D. Behne, and D. Hill, Review of scientific instruments **66**(2), 1201 (1995).
[6]T. Odstrčil, T. Pütterich, M. Odstrčil, A. Gude, V. Igochine, U. Stroth, A. U. Team, *et al.*, Review of Scientific Instruments **87**(12) (2016).
[7]R. Rossi, A. Murari, T. Craciunescu, I. Wyss, D. Mazon, A. Pau, A. Costantini, M. Gelfusa, J. Contributors, and E. T. E. Team, Nuclear Fusion **65**(3), 036030 (2025).
[8]D. R. Ferreira, P. J. Carvalho, and H. Fernandes, IEEE Transactions on Plasma Science **48**(1), 36 (2019).
[9]Z. Wang, Z. Zhang, D. Li, Y. Wei, Z. Yang, R. Yang, C. Wang, and Y. Dong, Journal of Fusion Energy **43**(2), 52 (2024).
[10]C. Mai, L. Hu, L. Xu, Y. Chao, K. Chen, Y. Chen, and E. Team, Plasma Physics and Controlled Fusion **64**(11), 115009 (2022).
[11]D. R. Ferreira, P. J. Carvalho, I. S. Carvalho, C. Stuart, P. J. Lomas, and J. Contributors, Fusion Engineering and Design **164**, 112179 (2021).
[12]C. M. Bishop, *Pattern Recognition and Machine Learning*, Information Science and Statistics (Springer, New York, 2006), ISBN 978-0-387-31073-2.
[13]J. R. Ferron *et al.*, Nuclear Fusion **38**(7), 1055 (1998).
[14]R. Conlin, K. Erickson, J. Abbate, and E. Kolemen, Engineering Applications of Artificial Intelligence **100**, 104182 (2021).
[15]M. Margo, B. Penaflor, H. Shen, J. Ferron, D. Piglowski, P. Nguyen, J. Rauch, M. Clement, A. Battey, and C. Rea, Fusion Engineering and Design **150**, 111368 (2020).
[16]A. Rothstein, H. J. Farre-Kaga, J. Butt, R. Shousha, K. Erickson, T. Wakatsuki, P. Steiner, S. Kim, A. Jalalvand, and E. Kolemen, Nuclear Fusion **66**(7), 076050 (2026).
[17]L. Ingesson *et al.*, *The mathematics of some tomography algorithms used at JET* (JET Joint Undertaking Abingdon, UK, 2000).
[18]M. L. Reinke, I. Hutchinson, J. Rice, N. Howard, A. Bader, S. Wukitch, Y. Lin, D. Pace, A. Hubbard, J. Hughes, *et al.*, Plasma Physics and Controlled Fusion **54**(4), 045004 (2012).
[19]M. Carr, A. Meakins, M. Bernert, P. David, C. Giroud, J. Harrison, S. Henderson, B. Lipschultz, F. Reimold, E. M. Team, *et al.*, Review of Scientific Instruments **89**(8) (2018).
[20]J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, *et al.*, GitHub repository (2018), `https://github.com/jax-ml/jax`.
[21]I. Loshchilov and F. Hutter, arXiv preprint arXiv:1711.05101 **10** (2019).
[22]D. Eldon, L. Casali, I. Bykov, C. Chrystal, K. Erickson, A. Hyatt, A. Leonard, A. Moser, A. Nelson, T. Odstrčil, *et al.*, Plasma Physics and Controlled Fusion **67**(1), 015018 (2025).
[23]D. Eldon, E. Kolemen, D. Humphreys, A. Hyatt, A. Järvinen, A. Leonard, A. McLean, A. Moser, T. Petrie, and M. Walker, Nuclear Materials and Energy **18**, 285 (2019).
[24]S. Makarov, D. Coster, T. Gleiter, D. Brida, M. Muraca, R. Dux, P. David, B. Kurzan, X. Bonnin, M. O'Mullane, *et al.*, Contributions to Plasma Physics **64**(7-8), e202300139 (2024).

10 November 2026 16:03:57