



Enabling adaptive pedestals in predictive transport simulations using neural networks

Downloaded from: <https://research.chalmers.se>, 2026-04-05 02:56 UTC

Citation for the original published paper (version of record):

Gillgren, A., Fransson, E., Yadykin, D. et al (2022). Enabling adaptive pedestals in predictive transport simulations using neural networks. *Nuclear Fusion*, 62(9).
<http://dx.doi.org/10.1088/1741-4326/ac7536>

N.B. When citing this work, cite the original published paper.

PAPER • OPEN ACCESS

Enabling adaptive pedestals in predictive transport simulations using neural networks

To cite this article: A. Gillgren *et al* 2022 *Nucl. Fusion* **62** 096006

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

You may also like

- [Advanced biomaterial strategies to transplant preformed micro-tissue engineered neural networks into the brain](#)
J P Harris, L A Struzyna, P L Murphy et al.
- [Measurement and simulation of unmyelinated nerve electrostimulation: *Lumbricus terrestris* experiment and numerical model](#)
A Šaroli, Z Živkovi and J P Reilly
- [8th Ibero-American Congress on Sensors \(IBERSENSOR 2012\)](#)
Idalia Ramos and Jorge J Santiago-Avilés

Enabling adaptive pedestals in predictive transport simulations using neural networks

A. Gillgren^{1,*}, E. Fransson¹, D. Yadykin¹, L. Frassinetti²,
P. Strand¹ and JET Contributors^a

¹ Department of Space, Earth and Environment, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden

² Division of Fusion Plasma Physics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden

E-mail: andreas.gillgren@chalmers.se

Received 11 March 2022, revised 13 May 2022

Accepted for publication 1 June 2022

Published 18 July 2022



CrossMark

Abstract

We present PEDESTAL Neural Network (PENN) as a machine learning model for tokamak pedestal predictions. Here, the model is trained using the EUROfusion JET pedestal database to predict the electron pedestal temperature and density from a set of global engineering and plasma parameters. Results show that PENN makes accurate predictions on the test set of the database, with $R^2 = 0.93$ for the temperature, and $R^2 = 0.91$ for the density. To demonstrate the applicability of the model, PENN is employed in the European transport simulator (ETS) to provide boundary conditions for the core of the plasma. In a case example in the ETS with varied neutral beam injection (NBI) power, results show that the model is consistent with previous studies regarding NBI power dependency on the pedestal. Additionally, we show how an uncertainty estimation method can be used to interpret the reliability of the predictions. Future work includes further analysis of how pedestal models, such as PENN, or other advanced deep learning models, can be more efficiently implemented in integrating modeling frameworks, and also how similar models may be generalized with respect to other tokamaks and future device scenarios.

Keywords: fusion, pedestal, AI, machine learning, neural networks, integrated modeling

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

1. Introduction


The edge region of tokamak plasmas imposes great challenges for designing and predicting the performance of future devices, such as ITER and DEMO, but also for modeling existing devices, such as JET, ASDEX Upgrade, and TCV. It has been observed that sharp temperature and density gradients can

form near the plasma edge at sufficient heating, which suppresses transport at the edge and gives a pedestal structure in the radial profiles. The characteristics of the pedestal are of high importance to the fusion community since it is associated with a high confinement state referred to as H-mode, which was first discovered in the ASDEX tokamak in 1982 [1].

Accurately predicting the pedestal is generally not a trivial task, and integrated modeling efforts can be inhibited by insufficient predictive capability for scenarios where the edge of a plasma cannot simply be assumed. EPED [2, 3] is a model for predicting the pedestal, which combines a description of the transport in the pedestal through the kinetic ballooning mode (KBM) constraint and an ideal linear magnetohydrodynamics

* Author to whom any correspondence should be addressed.

^a See Joffrin *et al* 2019 (<https://doi.org/10.1088/1741-4326/ab2276>) for the JET Contributors.

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

(MHD) stability code (ELITE) [4, 5]. Specifically, EPED is accurate in regimes where type I edge localized modes (ELMs) are triggered by ideal peeling ballooning instabilities [6, 7]. However for other ELM types, or for cases where the transport of the pedestal is not understood, where the KBM constraint does not describe the observations well, other methods are required [8, 9]. Additionally, first-principle based models for the pedestal have historically been too time-consuming to employ in some of the relevant integrated modeling environments, such as the European Transport Simulator (ETS) [10, 11]. This issue is currently being tackled by development of machine learning surrogate models that are showing promising results by drastically reducing the computational time while still capturing the dynamics of a first-principle based model [12, 13]. However, to our knowledge, there is still no first-principle based model that covers other ELM types or that does not require assumptions about the pedestal density.

Pedestal predictions based on empirical data are a different approach that similarly has been pursued before. For instance, the IPB98(y, 2) thermal confinement scaling [14] has been around for more than two decades, and the Cordey scaling [15] is a power law for the pedestal stored energy that is still in use today for setting boundary conditions in core transport models. Additionally, recently produced power laws [7] for the pedestal density and temperature serve similar purposes as those we intend to assess.

In this paper we present an alternative empirical pedestal model based on neural networks called PEdestal Neural Network (PENN). Our aim is to investigate how well the pedestal temperature and density can be predicted from global engineering parameters and main plasma parameters using this classical machine learning approach. Our other aim is to enable fast pedestal predictions in integrated modeling frameworks using PENN. As opposed to the other mentioned empirical models, we here choose to include main plasma parameters as input parameters since this imposes core-edge coupling, which provides self-consistent solutions in the integrated modeling framework. Core-edge coupling has recently been explored, for instance, in the EPED implementation for the OMFIT framework [16]. This work is in the same spirit, the main difference being that our model is empirical, and that the pedestal density is an output of our model instead of being an input as in EPED. In the scope of this paper, we utilize empirical data from the EUROfusion JET pedestal database [7], and we employ the model in the ETS.

The structure of the paper is as follows. In sections 2 and 3, we discuss the JET pedestal database and how we curate the data for our purposes. In section 4, we outline the training procedure and describe how we choose the input parameters for the model. In this section, we also show how the model performs on unseen data entries and how uncertainty of the predictions are calculated. We proceed by describing the implementation of PENN in ETS in section 5, which includes our procedures for handling ions as the database currently only consists of data for the electrons. This is followed up, in section 6, with results of a case example in ETS. We continue with exploratory work in section 7, and the paper ends, in section 8, with a discussion and conclusions.

2. Data set

The data utilized here comes from the EUROfusion JET pedestal database, including pulses both from JET-C (carbon wall) and JET-ILW (ITER-like wall). A thorough description of the database and its contents can be found in [7]. Here, we provide a general overview of the most important parts for our purpose.

The pedestal parameters in the database are derived from experimental electron temperature and density profiles for H-mode plasmas, where the key diagnostics used to obtain the profiles is the high resolution Thompson scattering (HRTS) [17]. Specifically, the pedestal parameters are determined in the pre-ELM phase (selected as the time interval 70%–99% of the ELM cycle) during a stationary phase that is at least 0.5 s long. To ensure that the phase is stationary, a check is performed for the power, β_N , gas rate, radiated power, and line integrated density. As part of the validation, all experimental profiles and fits have been visually checked and a more quantitative criterion based on the reduced chi-square χ_r^2 has been systematically implemented, which is also further described in [7]. For instance, the HRTS data points of a profile could be too scattered to ensure a good fit. Moreover, each data point in the database comes with an estimated error based on the uncertainty of the diagnostics used as well as the spread of the data in the pre-ELM phase.

To extract pedestal properties, such as the pedestal height, a modified hyperbolic tangent (mtanh) function is fitted to the data points of the temperature and density profiles

$$\text{mtanh} = \frac{h_1 - h_0}{2} \left(\frac{(1 - sx)e^x - e^{-x}}{e^x + e^{-x}} + 1 \right)$$

with

$$x = \frac{p - \psi_N}{w/2} \quad (1)$$

where h_1 is the pedestal height, h_0 is the pedestal offset in the scrape-off layer (SOL), s is proportional to the slope inside the pedestal top (also referred to as core slope), w is the pedestal width, and p is the pedestal position, defined as the radial point between the top and bottom of the pedestal. ψ_N is the radial coordinate in terms of normalized average poloidal flux coordinates, where $\psi_N = 1$ represents the last closed flux surface (LCFS). The width and position are also given in poloidal normalized flux coordinates, and only data between $0.8 < \psi_N < 1.05$ have been used in the fitting.

3. Statistics and filtering of data

Since neural networks are best suited for interpolation, it is expected that the model will be applicable for regimes and conditions that dominate the database. Hence, statistics and filtering of the data are important to highlight.

A total of 1835 validated entries from the JET pedestal database are included in the training of the neural networks. Out of these, 352 correspond to JET-C, and 1483 correspond to JET-ILW. For the main ion species, 1721 entries correspond to deuterium plasmas, and only 44 and 70 correspond to

hydrogen plasmas and mixed H/D plasmas respectively. Here, all of the hydrogen entries are associated with JET-ILW. There are 150 entries that are associated with impurity seeding, such as nitrogen and neon seeding. In addition to this, there is a subset of approximately 350 entries associated with a low impurity level done for technical reasons, such as for charge exchange recombination spectroscopy measurements, although this low rate does not affect the pedestal significantly. For the divertor, we have the following number of entries for each configuration: V/V: 190, V/C: 267, V/H: 658, C/V: 16, C/C: 700, C/H: 4. Here, the acronym ‘V/C’ stands for a divertor configuration with the vertical (V) target as the inner strike point, and the corner (C) target as the outer strike point (‘H’ stands for the horizontal target). As mentioned, the statistics presented here aims to guide for which scenarios PENN is best suited for. For instance, PENN is expected to be more reliable for simulations with deuterium plasmas compared to simulations with hydrogen plasmas due to the large difference in data size, since this induces bias. This does not mean that it is impossible to get accurate results with hydrogen plasmas, however the difference in data size should be considered.

We choose to exclude a small subset of entries where it is suspected that other ELM types are present, such as type III ELMs. For future work, a combined model for different ELM types will be of interest, however, for the first main version of this model, we choose to focus on type I ELMs to demonstrate the concept. Additionally, we exclude entries that either include kicks, pellets or RMPs. These features can affect the pedestal and could also be included in the future, however for now we do not desire to involve such effects in the PENN implementation in the ETS.

4. Neural networks for pedestal predictions

This section aims to provide an overview of the neural networks of PENN. First, we discuss the choice of which plasma/engineering parameters that the model uses. We proceed by describing how the neural network architecture is determined and how the model is trained. We then show how well the model performs on unseen data entries of the different categories described by the statistics in section 3. Finally, we describe a method to estimate how reliable a prediction is.

4.1. Input parameters

In the scope of this paper, we only choose from scalars that are included in the JET pedestal database. We also consider which parameters that are available in the framework where we intend to employ the model, which in this case is the ETS. Out of these, our tests with feature extraction and tests regarding how the prediction error varies shows that the global parameters in table 1 contribute to improving the model. Additionally, several of these parameters are independent with respect to each other and have shown to be useful for other models [3, 7, 15, 18]. Each parameter in this database additionally comes with an estimated error based on the measurement uncertainty for the different diagnostics systems. For the training of neural networks, minor error in the parameters are

usually not of concern. However, if the model is applied for a simulation case where there might be high measurement uncertainty for an input parameter for that particular pulse, it might be useful to perform a sensitivity scan for that input parameter.

As mentioned in the introduction, some pedestal models solely rely on engineering parameters as input parameters [7, 15]. For instance, the total power input is usually included, but the global plasma parameter β_N is not always included since it is calculated from the profiles, and the pedestal plays a part in determining the profiles through boundary conditions. In other words, for some applications, it does not make sense to use plasma parameters as inputs as they can be considered to be a part of the solution. However, to capture effects that comes from the interplay between the core and the pedestal in the ETS, it is necessary to incorporate at least one plasma parameter, such as β_N . That said, it is of course possible to train an alternative PENN model with only engineering parameters, which would represent a static pedestal for particular engineering parameter setups. This is beneficial in the sense of not requiring information about β_N , Z_{eff} , and q_{95} , although we have seen that this comes with a significant cost regarding how well the model performs.

Parameters that are excluded are, for instance, related to peeling–ballooning stability analysis, and dimensionless parameters that are non-global, such as the collisionality at the pedestal top. Parameters that are more or less constant throughout the database are also excluded. For instance, the major tokamak radius is excluded since we only look at JET data for this first version of the model.

It could be argued that we are neglecting some parameters that could improve the model if they were included. For instance, we exclude effective mass (M), tokamak wall type, divertor configuration, fuel rate of main species, and a parameter related to impurity seeding. For the effective mass parameter, the distribution is imbalanced with a large spike at $M \approx 2$ corresponding to the deuterium plasmas, and a small spike at $M \approx 1$ corresponding to the hydrogen plasmas. Generally, the parameters that are included in table 1 cover more even distributions which is more suitable for neural networks, especially when no advanced procedures to counter imbalance are used [19]. However, we have seen that by simply including the hydrogen entries during the training, we are able to make accurate predictions on unseen hydrogen entries as well, even if we exclude M as an explicit input parameter. The same reasoning applies for the other four mentioned parameters. For the fueling and impurity seeding, we have seen that the predictions are not improved by including these as parameters, as long as we train on all data. For instance, the model performs poorly on entries with impurity seeding if we exclude entries with impurity seeding during the training, but drastically improves when impurity seeding entries are included in the training even if there is no parameter related to impurity seeding. We have also seen that the model does not improve if we explicitly specify the wall type or the divertor configuration as input parameters. Generally, these results could be interpreted as if the information from these parameters are represented by the combination of the other input parameters in table 1. For instance, two plasmas with identical engineering

Table 1. The output and input parameters of PENN with their ranges. The values correspond to the subset of the JET pedestal database that is described in sections 2 and 3. Each value in the database comes with a corresponding absolute error estimation based on the measurement uncertainty of the different diagnostics. Here, the mean error of each parameter is shown as well as the standard deviation of the error and the maximum error of the parameters.

	Min value	Max value	Mean error	Std error	Max error
Output parameters					
Pedestal temperature T_e (keV)	0.15	1.89	0.030	0.024	0.185
Pedestal density N_e (10^{19} m^{-3})	1.73	11.33	0.070	0.033	0.317
Input parameters					
Global β_N (MHD)	0.58	3.62	0.051	0.028	0.218
I_P (MA)	0.82	4.48	0.009	0.006	0.041
B_T (T)	0.97	3.68	0.001	0.006	0.023
Minor radius a (m)	0.88	0.97	0.003	0.001	0.013
Elongation κ	1.60	1.81	0.006	0.003	0.028
NBI power (MW)	3.5	32.13	0.711	0.523	3.112
Total power (MW)	3.99	38.22	0.299	0.177	2.650
Upper triangularity δ_{up}	0.06	0.58	0.005	0.004	0.032
Lower triangularity δ_{low}	0.24	0.49	0.005	0.005	0.029
Plasma volume (m^3)	69.86	82.18	0.317	0.144	2.026
q95	2.43	6.09	0.026	0.020	0.166
Z_{eff}	1.01	3.74	0.051	0.096	0.352

parameters but different wall will have different β_N and Z_{eff} [20]. If a parameter is assumed to be important for pedestal predictions, but no improvement can be detected by including it as an input parameter, it is reasonable to believe that the other parameters can represent it through a combined proxy. This is clearly an advantage of neural networks, since they allow for such correlations to be learned within their structure. For instance, we have seen that we can predict the wall type with 99% accuracy from the input parameters in table 1 using neural networks, which suggests that such correlations can form.

Some parameters in table 1 could probably be excluded without reducing the performance of the model significantly. The reason for specifically excluding the parameters mentioned above has to do with less availability of these parameters in the ETS, or their unsuitable distributions in the database.

4.2. Training

We employ a feedforward neural network architecture with layers of artificial nodes as the basis for PENN, where we use the principles of backpropagation to optimize the model. Here, the first layer consists of the input parameters shown in table 1, and the final (output) layer consists of two nodes, which represent the electron pedestal density and temperature respectively. The remaining layers in between are referred to as hidden

layers. During the early stages of this work, it was found that simple neural networks performed better than other classical fitting models, such as log-linear regression, but also slightly better than other non-linear machine learning models, such as random forests and support-vector machines.

To find the optimal hyper-parameters of the neural network, we employ the neural network intelligence (NNI) framework developed by Microsoft, which is an open source toolkit designed to automatically search for the optimal network architecture and hyper-parameters for our problem (<https://github.com/microsoft/nni>). Specifically, the NNI framework simplifies exhaustive grid searches and provides visualization tools to analyze the results. For our problem, results show that the following hyper-parameter combination represents the optimal configuration: optimizer: Adam, loss function: MSE, learning rate: $lr = 0.001$, epochs = 150, batch size = 30, activation function: ReLu, and four hidden layers with 60, 50, 30, and 20 nodes each. We also employ standard scaling to all our input and output parameters, where each distribution is converted to obtain a mean value of 0 and a standard deviation of 1. For the actual training of the neural networks, we employ functionalities from the TensorFlow [21] and Scikit Learn [22] libraries. Additionally, we employ the concept of ensemble learning [23], which here means that we train a combined model that consists of 15 individual neural networks that

perform the same task. The purpose of this is to enhance the generalization of the model by giving each network a unique set of weight and bias parameters before the training, and by excluding a unique subset of 10% of the data for each network during the training. Consequently, each network will end up with unique set of weight and bias parameters when trained (although the same hyperparameters still applies for all networks). In practice, this means that when the input parameters of a pulse are forwarded to the model, it will produce 15 predictions of the pedestal top values with slight variations since the individual models are slightly different. This means that the input parameters are the same for all individual networks while the output is not. Here, we use the mean output value from all individual neural network predictions as the final prediction of PENN.

4.3. Accuracy on training set and test set

We here define accuracy as how well the model is predicting the pedestal top values with respect to the actual pedestal top values in the database. This is quantified with the coefficient of determination R^2 where $R^2 = 1$ represents a model that make perfect predictions on all data points, and where $R^2 = 0$ represents a model with equal predictive capabilities as a model that randomly guesses the pedestal top values. To investigate the accuracy of the model, we make predictions on 80% of the data which is included during the training, but also on the remaining 20% of the data which is excluded during the training, which are referred to as the training set and test set respectively. The purpose of this being that we want to show and compare biased and unbiased results. A side note here is that we only use the training set in the hyper-parameter search described in the previous section (divided into a temporary training set and a validation set). In figure 1, the results of the different categories of the database are shown, both for the training set and the test set.

4.4. Uncertainty estimation of predictions

It is straightforward to estimate the error of a prediction when the true values are available, which is the case for the results in figure 1. However, when PENN is employed in the ETS, another approach is required since we do not desire to rely on database values for predictive simulations. The first step is to make sure that each input parameters in a prediction is within its training range highlighted in table 1. Neural networks are notorious for extrapolating poorly [24], thus we can assume that the error increases as we go beyond the minimum and maximum values of these parameters. Technical and physics constraints from theory, such as MHD, may also guide for which scenarios that represent extrapolation in the multidimensional input parameter domain, even if the individual parameters are within their training range. For instance, typically high power requires high gas rate to mitigate heat loads and minimize the impurity flux into the plasma. In general, this leads to a much smaller input parameter space compared to using the input parameters independently. Additionally, we can expect a similar error of the predictions as the mean absolute error in the test set as seen in figure 1, which is 0.0585 keV for

the electron temperature and 0.412 m^{-3} (10^{19}) for the electron density.

For a more thorough data driven approach to determine if a prediction is reliable, we utilize features from the ensemble learning method described in section 4.2. Specifically, when a prediction is made, the distribution of the individual predictions in the ensemble can be viewed as a proxy for how well the networks agree, the idea here being that a wide distribution of the predictions indicates low confidence in the model as a whole. Thus, in this paper we define uncertainty as the standard deviation of the ensemble of predictions. Moreover, we define normalized uncertainty as the standard deviation of the predictions before they are rescaled to the units of keV and m^{-3} (referring to the standard scaling method described in section 4.2). In figure 2, the normalized uncertainty of the entries in the database is shown, both for the training set and the test set from the accuracy test described in section 4.3. By analyzing the normalized uncertainty instead of the uncertainty, a direct comparison can be made between the temperature (blue) and density (orange) since both parameters are viewed in the normalized space. It can be seen that normalized uncertainty below 0.25 and 0.4 represents a level of confidence comparable to the predictions on the training set and test set respectively, both for the temperature and density.

In practice, this means that when the model is given the input parameters from a pulse, it will produce 15 outputs for the temperature and density each, where the mean value of each output parameter represents the final prediction as explained in section 4.2. Additionally, the model will produce a value for the normalized uncertainty for each output parameter by looking at the standard deviation of the 15 predictions. This value should thus be analyzed when using the model in an application. If the normalized uncertainty associated with a prediction is higher than 0.4, which is approximately the highest value in the test set (see figure 2(b)), the prediction cannot be considered reliable. Even predictions with a normalized uncertainty between 0.25–0.4 should be considered less reliable compared to predictions with a normalized uncertainty below 0.25. High normalized uncertainty could, for instance, be an indication of a collective extrapolation even though the individual input parameters are within the training range. In that sense, this method indicates if the model recognises the pattern of the input parameters.

5. Implementation in the European transport simulator

This section describes an overview of the most relevant features that are involved in the implementation of PENN in the ETS. First, we describe how we calculate the edge in terms of 1D profiles from the pedestal predictions, and secondly, we describe how we estimate ion pedestals from the electron pedestal predictions.

5.1. Constructing the edge from pedestal predictions

In the ETS (version 6), we have access to Interface Data Structures (IDS) from the IMAS framework [25] for different

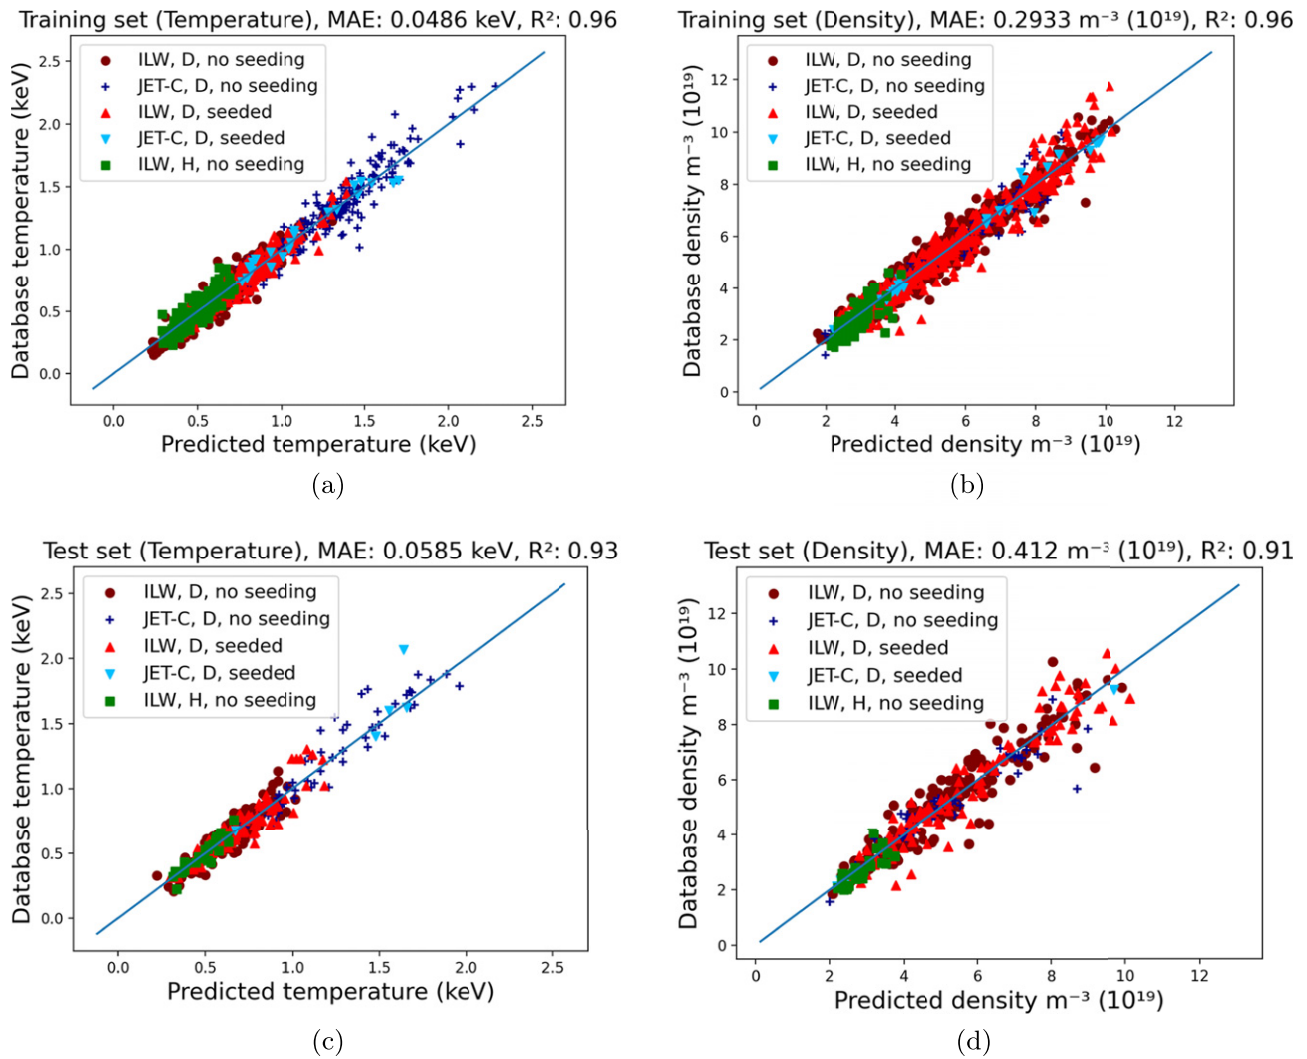


Figure 1. The prediction results of PENN for the training set (a) and (b) and the test set (c) and (d). The prediction corresponds to the x -value of each entry, and the true database value corresponds to the y -value of each entry. The straight lines $y = x$ acts as reference lines for ‘perfect predictions’. The data is categorised by main ion species, wall type, and if the entry contains impurity seeding. The measurement uncertainty for each entry in the database is excluded in the figures to increase the readability. It should however be mentioned that the database uncertainty is negligible for the density, and mostly relevant for higher temperatures, where the maximum temperature error estimation in the database is approximately 0.19 keV. As can be seen, the prediction accuracy is similar for all categories with the exception of the electron temperature for JET-C, which could be because the measurement error is highest in that regime. In general, larger uncertainties in the input and output parameters lead to a lower accuracy of the model when it is evaluated. As expected, the accuracy on the training set is a bit better compared to the test set, but not significantly better which is important for ensuring that we are not overfitting the model.

pulses, which contains the sought input parameters to make neural network predictions. Once the predictions of the scalars representing the pedestal top values are made, additional information is required to construct profiles that represents the entire edge. In this paper we define the edge as the region between the boundary condition for the core equations and the LCFS. Here, we employ the same mtanh function described in section 2 that is used to extract the pedestal parameters from the experimental profiles. On the contrary, we now form the edge profiles from the pedestal parameters. Thus, we need to estimate the pedestal position, width, core slope, and offset. Throughout the database, the width and position of the pedestal corresponds to narrow distributions, and for simplicity we apply the mean value for each parameter: $\psi_N = 0.98$ for the pedestal position and $\psi_N = 0.045$ for the pedestal width,

both for the electron temperature and density. For the offset, 100 eV is set for the electron temperature as the profiles are systematically shifted to have a separatrix temperature of 100 eV in the HRTS diagnostics [7], and $10^{12} m^{-3}$ is set for the electron density too avoid numerical issues that can occur with too low densities. In practice, the offset values are not reached at the LCFS due to the fixed position and width, however, they serve the purpose of maintaining the H-mode pedestal structure even if the profiles plateau further out in the SOL. For applications where SOL physics are of greater importance, a more detailed approach could be considered.

The last parameter, which is the core slope, is an important parameter since the profiles are extrapolated inwards from the pedestal top to the boundary conditions based on this slope. Here, we calculate the linear tangent of the core close to the

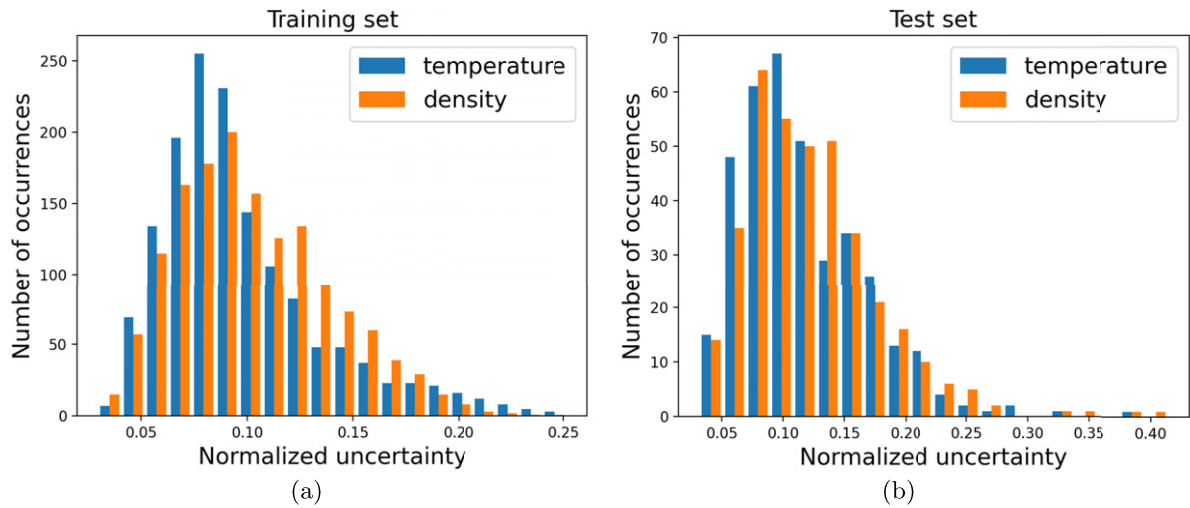


Figure 2. Normalized uncertainty histograms from the ensemble of neural networks for the training set (a) and test set (b) respectively. Here, the blue stacks represent the temperature and the orange stacks represent the density. It becomes clear why it is useful to compare the uncertainty histograms of the different properties in normalized space, as the stacks almost fully overlap (the temperature shows a slightly lower normalized uncertainty).

internal boundary from the initial profiles. In other words, we assume the core slope to be the same as the slope in the initial experimental profiles when we make a new edge prediction. If the slope of the core close to the boundary condition changes during a simulation, for instance due to transport equations, the slope between the boundary condition and the pedestal adapts to match the inner slope.

5.2. Estimating ion pedestals from electron predictions

As previously mentioned, the database currently only consist of pedestal data for the electron temperature and density. For the application in the ETS, additional strategies are required to enable adaptive ion pedestals. For the temperature, we allow the user to set a ratio between the electron pedestal top prediction and the estimated ion pedestal top. The pedestal position, width, and offset are set to be the same for the ion temperature as for the electron temperature. The core slope is obtained in the same way as for the electrons.

The ion density is treated differently since quasi-neutrality needs to be taken into account. If only one ion species is considered, the approach is straightforward since there is only one way quasi-neutrality can be fulfilled at each grid point if impurities are neglected. However, if multiple ion species are involved, several approaches are possible, for instance, by either keeping one or several of the ions species fixed, or keeping the ratio between the ion species constant. In this paper, we choose to look at an example with only one ion species to simplify the features that are not directly related to the machine learning aspects of the model.

6. Application results

6.1. European transport simulator (ETS)

In this section we demonstrate some of the benefits that PENN provides to integrated modeling. We study pulse 84792 from

JET, which is included in a neutral beam injection (NBI) power scan analysis performed by Challis *et al* [26]. The power scan includes pulses with similar values of the engineering parameters except for the NBI power and thus total power. Pulse 84792 is a deuterium pulse that was performed with 13 MW of NBI power, and another pulse in the power scan, 84793, was performed with 4.5 MW of NBI power. These two pulses were performed at the same plasma current and magnetic field (1.4 MA/1.7 T). For this application, we have excluded pulse 84792 and 84793 in the training of the neural networks to avoid bias.

Here, we simulate pulse 84792, both with its original NBI power of 13 MW, but also with 4.5 MW to allow for comparison with the pedestal of pulse 84793. In these simulations, the pedestal and the core are coupled through a boundary condition at $\rho = 0.85$, where PENN is responsible for the pedestal. For the core, we use EDWM [27] for the turbulent transport, NCLASS [28] for the neoclassical transport, and BBNBI [29] for the NBI source terms. EDWM transport has been slightly enhanced in the electron heat channel to better match the original power and pedestal. These settings are the kept for all simulations. We start with profiles from the experiment, and we consider a time window of 1.5 s with a time step of 0.001 s. We evolve the deuterium and electron temperatures, as well as the deuterium density, where we obtain the deuterium density pedestal through predictions of the electron density pedestal together with quasi-neutrality. The core electron density is similarly calculated from quasi-neutrality.

This process is repeated for three scenarios. First, we apply the original NBI power of 13 MW, both for PENN and the core models. Secondly, we apply 13 MW for PENN, but 4.5 MW for the core models, which represents the case not being able to adjust the edge due to the lack of a pedestal model for a new scenario of 4.5 MW. This approach has been common, for instance, in the subject of density peaking in tokamaks, where the effect of changed NBI power is investigated [30–34]. The

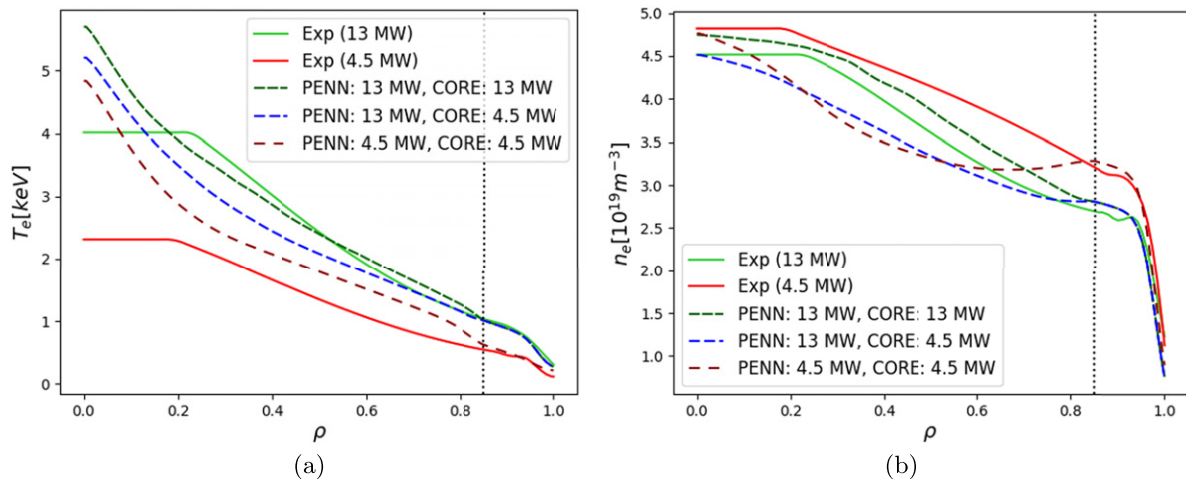


Figure 3. Results of the simulations in the ETS. The electron temperature is shown in (a) and the electron density is shown in (b). The steady-state solutions where PENN has adjusted the pedestal (dashed red) show a similar change in the pedestal as in [26] as well as in the experiment (solid red).

third and final simulation represents a more comprehensive approach where we additionally adjust the edge with PENN according to the new NBI power of 4.5 MW. Hence, to see the full effect of the change in NBI power.

The results in figure 3 shows the steady-state solutions of the three scenarios in the ETS (dashed red, dashed blue, and dashed green), as well as the experimental profiles (solid green and solid red). The black vertical dotted line at $\rho = 0.85$ represents the position of the boundary condition, where PENN provides the pedestal to the right of this line. We see that PENN produces a decrease in pedestal temperature and an increase in pedestal density as the NBI power is decreased (dashed red). This is consistent with the results from the NBI power scan by Challis *et al* [26]. Specifically, by changing the NBI power from 13 MW to 4.5 MW for pulse 84792, we obtain a pedestal that is close to the experimental values of the pedestal for pulse 84793 (solid red). For all simulations, the pedestal predictions show a normalized uncertainty below 0.25 which indicates prediction confidence, as described in section 4.4.

The result for the core, in figure 3, show a good match between the 13 MW experimental profiles (solid green) and the simulation (dashed green). The dashed blue profiles (lower NBI power in the core) indicate a lower temperature and density. Lastly, the simulation where the pedestal and core power were changed are displayed by the dashed red profiles. Noteworthy is the slightly hollow density profile (dashed red) close to the boundary condition which is not seen in the experiment (solid red). This might suggest that the pinch term for the particle transport is too low in the turbulence transport model EDWM. This simulation has the highest collisionality, and as high collisionality suppress the particle pinch we get a flat density profile. The bumps in the temperature profiles are also artifacts from the turbulent transport model EDWM. If the transport model does not give enough flux (particle or heat) compared to the internal sources (NBI or radio-frequency-heating) close to the boundary the gradient will increase. As EDWM, and other turbulent transport models, is gradient

driven this will increase the fluxes until the fluxes from EDWM are as large as the internal sources.

6.2. Standalone NBI dependency

In this section we further investigate NBI power dependency of PENN, and demonstrate how the uncertainty estimation can be used in the analysis. We select four JET pulses from the NBI power scan discussed by Challis *et al* [26]. Here, we gradually increase the NBI power from 0 MW to 30 MW, and make predictions with PENN while keeping the other parameters constant. This is a standalone test, which means that PENN is not coupled to ETS or any other code. Hence, we choose to temporarily drop the plasma parameters β_N , Z_{eff} , and q_{95} as input parameters since we want to avoid arbitrary estimations of these parameters throughout the hypothetical NBI power scan. The four pulses included in this test are: 75738: C-wall low triangularity, 76854: C-wall high triangularity, 84545: ILW high triangularity, and 84792: ILW low triangularity. All of these pulses were performed at the same plasma current and magnetic field (1.4 MA/1.7 T), with the exception of 75738 (1.7 MA/2.0 T). To avoid bias, the pulses are excluded in the training of the neural networks for this particular test.

The results in figures 4(a) and (c) shows the predicted pedestal temperature and density as functions of NBI power for the four different scenarios. The NBI power range discussed in [26] is highlighted in the figure with bold lines. In this regime, the temperature increases and the density decreases with increasing NBI power. Thus, our results show the same dependency as the NBI power scans in Challis *et al* [26]. As seen in figure 4(a), another distinguished feature is that the high triangularity pulses show a higher temperature and density.

As we go to higher values of the NBI power, the temperature plateaux and the density starts to increase past 15 MW. The naive approach would be to trust these results since the JET pedestal database includes entries with NBI power up to 32 MW, meaning that we are within the training range.

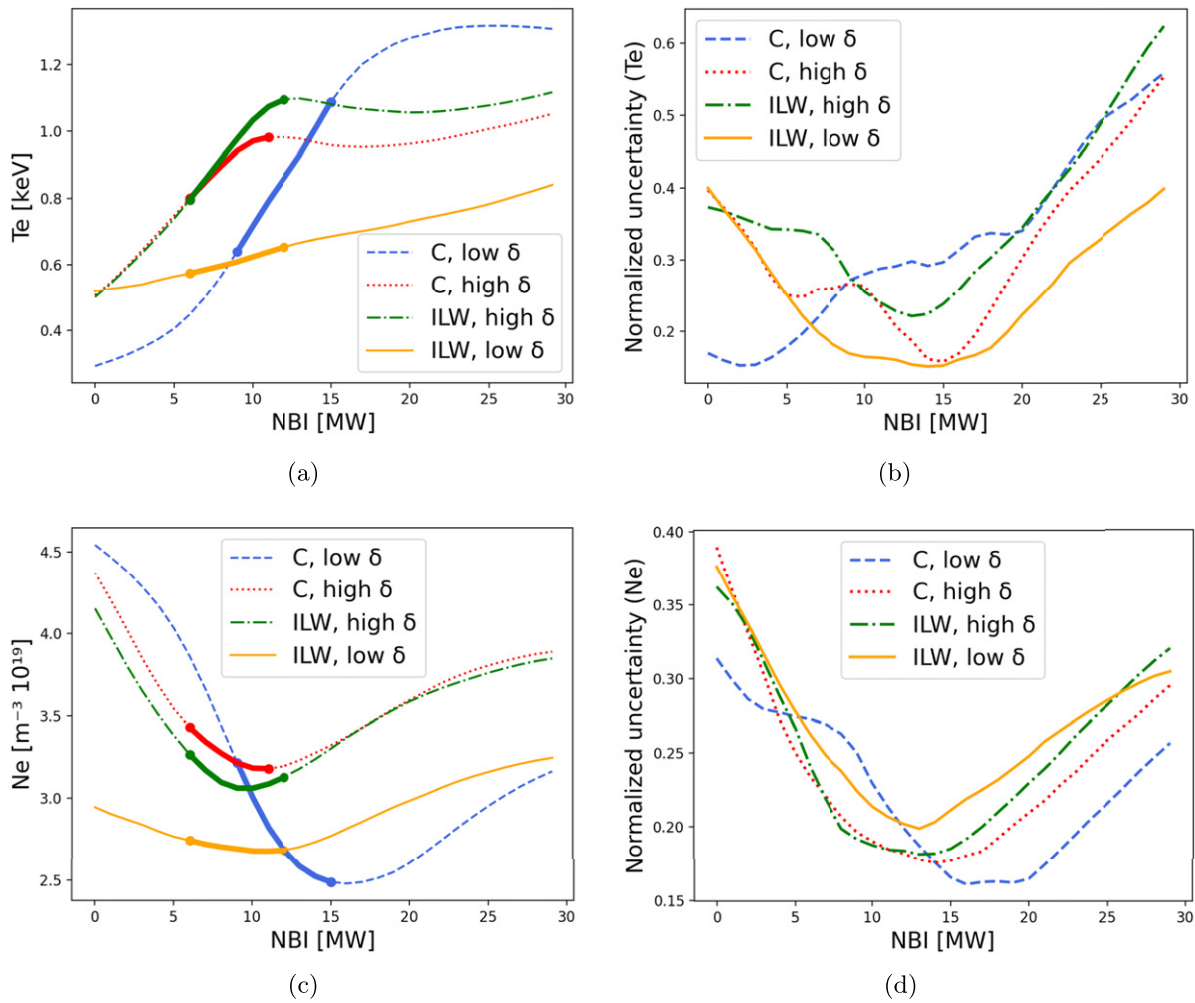


Figure 4. Results from the standalone NBI power scan. These are not real temperatures from JET but predictions from PENN. In (a), the pedestal temperature prediction is seen to increase for the NBI power region discussed in [26] (bold lines). In (b), the associated normalized uncertainty is estimated for the predictions as a function of NBI power. The corresponding result for the density and the normalized uncertainty for the density is seen in (c) and (d). In section 4.4, we saw that a normalized uncertainty above 0.25 represent low confidence. Thus, the model is only reliable in the region between approximately 4–15 MW for pulses with similar engineering parameters (1.4 MA/1.7 T). The results in (b) does not generally discourage PENN predictions with NBI power greater than 15 MW, since higher magnetic field and plasma current is associated with higher NBI power in the JET pedestal database. However, for these pulses with a particular set of engineering parameters, predictions with NBI power greater than 15 MW are discouraged as long as the other parameters are fixed.

However, by employing the uncertainty estimation method described in section 4.4, we see that this assumption cannot be made with confidence. Figures 4(b) and (d) shows the associated normalized uncertainty of the ensemble of neural networks as a function of NBI power for the four pulses. As described in section 4.4, the standard deviation between the predictions of the individual networks in an ensemble network can be associated with uncertainty. Here, it is clear that PENN is more confident in the NBI range that is covered in the NBI power scan for these particular engineering scenarios. This is because pulses in the JET pedestal database with similar engineering parameters have been performed with NBI power in this range. In section 4.4, we stated that a normalized uncertainty above 0.25 indicates low confidence for a prediction. Additionally, even if the normalized uncertainty is below 0.25 for the density for some data points above 15 MW, the rapid increase of normalized uncertainty and the normalized

uncertainty for the temperature should be seen as a warning in this NBI range. Thus, we cannot claim that the temperature will plateau, and that the density will increase for higher NBI power. We can also observe that approximately 4–5 MW represent a lower boundary of confidence for these pulses. This is a typical example of where the model is extrapolating even though the individual parameters are within the training range of the full database.

That said, this analysis only accounts for pulses with similar engineering parameters as in 75738, 76854, 84545 and 84792. The NBI range of validity is shifted based on the other engineering parameters. For instance, most of the entries in the JET pedestal database with NBI power above 20 MW are associated with a magnetic field of 2.0 T–3.5 T, and a plasma current of 1.5 MA–4.5 MA. Thus, in a general sense, predictions with NBI power greater than 15 MW are not discouraged. This

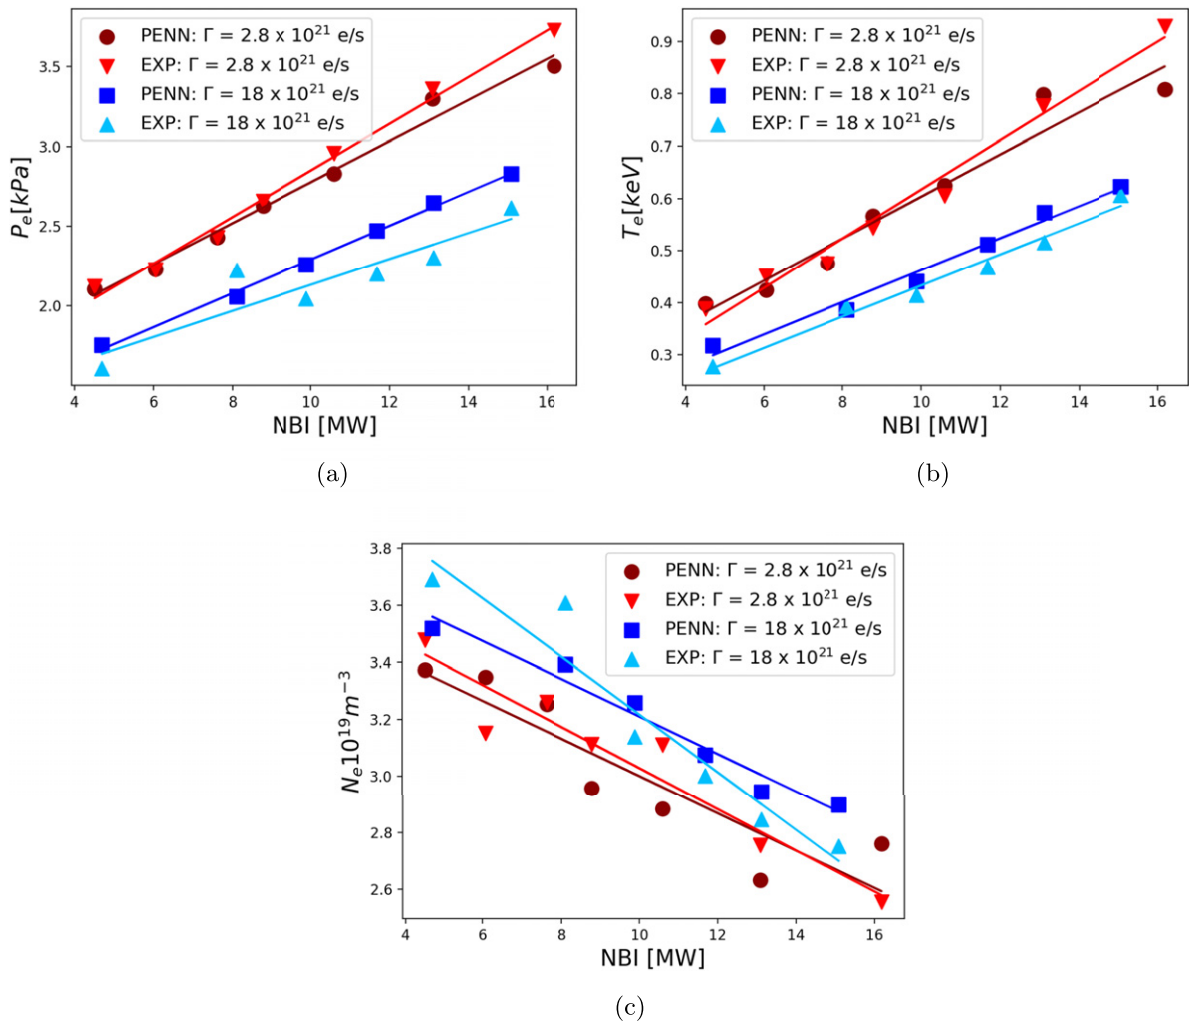


Figure 5. Results of the standalone fueling scan test with PENN. The markers represent the data points and the lines represent linear regressions of the data points to highlight the trends. Here, fueling is not included as an input parameter, however PENN is able to distinguish between the low and high fueling cases anyway due to small differences in the other input parameters. The predictions are compared with experimental data which overall show good agreement. The pedestal temperature (b) shows a stronger dependency with respect to the fueling compared to the pedestal density (c), which leads to an overall decreased pedestal pressure (a) with increased fueling.

uncertainty estimation method rather describes a tool to determine, for instance, which range of NBI power that is appropriate for a particular set of engineering parameters when making predictions with PENN. In a broader sense, this method is not limited to the NBI parameter. The same test can be done for each input parameter of PENN, such as the magnetic field and the elongation.

Of course, there are also physics based reasons for why PENN does not handle high NBI power in this case. The high end of the NBI range for a given magnetic field strength is likely limited by β_N , which triggers MHD instabilities. For obvious reasons, the database does not contain scenarios that are prohibited by physics constraints and engineering limitations. However, the database does not necessarily cover the full multidimensional regime that is allowed from the physics and engineering standpoint. Thus, the machine learning approach serves a purpose of guiding for which scenarios PENN is specifically suited for, which is directly correlated to the scenarios that exist in the

database. In conclusion, both the physics/engineering and the machine learning approach may be used in a complementary way.

6.3. Standalone fueling dependency

As mentioned in section 4.1, we have seen that the accuracy of PENN is not significantly improved by, for instance, providing an explicit input parameter for fueling. However, a reduced pedestal pressure due to increased fueling is regularly observed at JET-ILW [4]. In this section, we show how PENN can capture this effect without providing fueling as an input parameter.

To investigate this topic, we perform predictions on pulses from the 1.4 MA/1.7 T fueling and power scan described in [4]. Specifically, we look at pulses from two power scans between 4.5 MW–16 MW that has been performed at two different fueling rates: $\Gamma = 2.8 \times 10^{21}$ e/s, $\Gamma = 18 \times 10^{21}$ e/s. Since PENN is designed to predict the pedestal temperature and density, we first predict these quantities, and then we estimate the

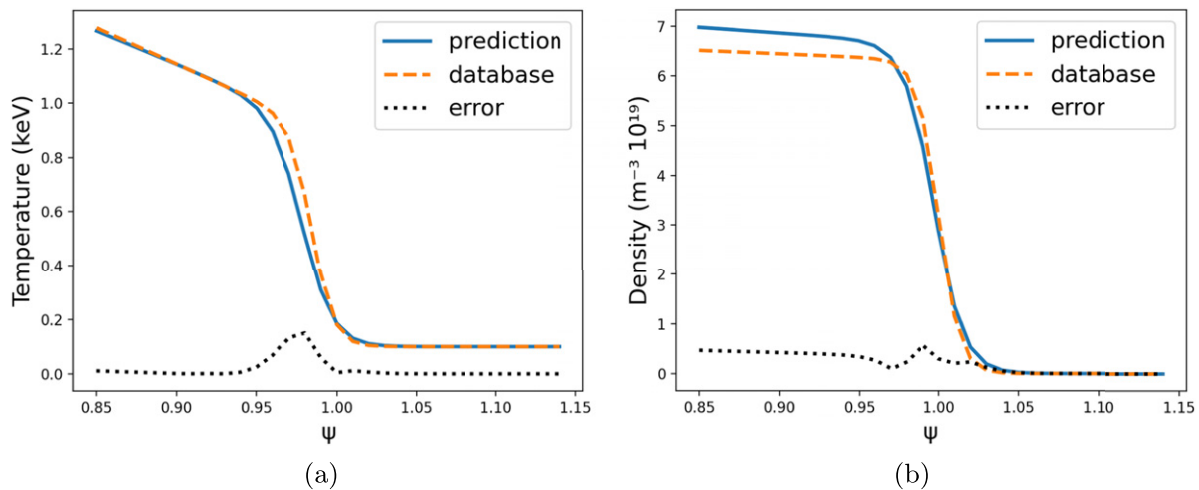


Figure 6. The prediction of the multiple output model is compared with the true values of the database for JET pulse 79757 in the test set. It is an accurate prediction, both for the temperature and density. The error, which is defined as the absolute difference between the prediction and the experiment, is seen to peak at the gradient of the pedestal and then increase slightly toward $\Psi = 0.85$ for the density. The x -axis represents the normalized poloidal flux surface coordinate.

pedestal pressure with $P_e = 1.6 \times N_e T_e$, where 1.6 approximately represents from the Boltzmann constant with T_e in keV, n_e in 10^{19} m^{-3} to have the pressure in kPa.

In [4], it can be seen that β_N shows the same trend as the pedestal pressure for different fueling rates. Thus, we here temporarily drop β_N as an input parameter since there is no challenge in predicting the pedestal pressure with knowledge about β_N for this particular case.

In figure 5, the results of PENN predictions and experimental data are shown for the fueling scan. In this section, we are not necessarily interested in the power dependency which is represented on the x -axis, however, this axis needs to be present to distinguish between the high and low fueling cases since the power varies throughout the scan. It can be seen that PENN is able to capture the effect of higher fueling which leads to a lower pedestal pressure (a) without knowing explicitly about the fueling or β_N . Here, the experimental data is from the actual pedestal pressure in the pedestal database, and the predicted pressure is calculated from the prediction of the pedestal temperature and density. In figures 5(b) and (c), the predictions of the pedestal temperature and density are shown. By comparing the temperature and the density, it is clear that they show opposite trends although the temperature is mostly affected by the variation of fueling rate, while the density shows a weak dependency, which is coherent with the results in [4]. This leads to an overall decrease in pedestal pressure with increased fueling, since the pedestal temperature decreases with increased fueling.

PENN is able to capture this effect due to small changes in the other parameters, such as Z_{eff} , elongation, triangularity, plasma volume and q_{95} . Exactly how these parameter interact as a substitute for the fueling to achieve this result lays hidden in the neural network structure, which requires a significant analysis effort which is out of the scope of this paper. To our knowledge, there is no general downside of using the fueling parameter as an input parameter for pedestal models. In this

paper, we simply exclude it as a parameter since we do not want to rely on explicit fueling information in the ETS where we employ PENN.

7. Further exploration with multiple outputs

In this section, we present an alternative version of PENN. Here, we construct the neural networks with multiple outputs that represent all of the pedestal parameters. By using the mtanh function (1), the output can be converted to 1D edge profiles without additional strategies for parameters, such as the pedestal position and core slope. By forcing the model to predict all of the pedestal parameters simultaneously, there is a slight decrease in the accuracy of the pedestal top values compared to the model that only predicts the pedestal top values. The accuracy of the test set is reduced from $R^2 = 0.93$ to $R^2 = 0.87$ for the temperature and from $R^2 = 0.91$ to $R^2 = 0.89$ for the density. The accuracy of the other parameters are: width (temperature): $R^2 = 0.42$, width (density): $R^2 = 0.51$, position (temperature): $R^2 = 0.62$, position (density): $R^2 = 0.61$, slope (temperature): $R^2 = 0.44$, slope (density): $R^2 = 0.42$. It is interesting that the model is able map the input parameters to all of these pedestal parameters to some extent, even if the accuracy is by far best for the pedestal top values.

An example with JET pulse 79757 is shown in figure 6 where the neural network employs the same input parameters as in the original model, but predicts the pedestal top value, width, position, and core slope simultaneously. This entry is taken from the test set and shows good agreement with the database. The error, which is defined as the absolute difference between the prediction and the experiment is also shown. It can be seen that the error peaks at the steep gradient of the pedestal for both the temperature and density, which is expected since a small difference in pedestal width and position can give large differences in this region. From the pedestal top toward the

core (here the core is defined as the region $\Psi < 0.85$), the error may look different depending on the accuracy of the prediction of the slope and the pedestal top value. If the slope is accurately predicted, but the pedestal top value is not, this will give a constant offset error as the profile expands toward the core. On the other hand, if the pedestal top value is accurately predicted, but the slope is not, the error will gradually increase from the pedestal top position toward the core. The density prediction in figure 6(b) illustrates this to a small extent. Since the accuracy of the pedestal top values are higher than the accuracy of the slopes, it might be useful to move the boundary condition closer toward the pedestal top if this multiple output model is used in a simulation framework that employs boundary conditions, such as the ETS (if this does not affect the simulation negatively in other ways). This is because a predicted slope that is less accurate will have less influence if the distance is shorter as the profile propagates from the pedestal top toward the core.

8. Summary and conclusions

In this paper we have presented PENN, an empirical model for pedestal predictions using neural networks. We have shown the performance of the model on unseen data entries, and how the model can be implemented in simulation frameworks, such as the ETS. The presented method of neural network training and implementation is currently our best approach, however this might be changed in the future as development is made. Moreover, some methodologies are purposely chosen to be less complicated, such as the uncertainty estimation and the procedure to handle ions, since the purpose of this paper is to showcase the principle and its flexibility.

In section 6, we demonstrated how the model can be used in integrated modeling frameworks to enable adaptive boundary conditions for the core of the plasma. Here, results are shown to be consistent with previous studies of how the pedestal is affected by varied NBI power. Additionally, we saw that the method of uncertainty estimation can be useful for interpreting the reliability of the model for different engineering scenarios.

An interesting aspect of PENN is its ability to maintain the same accuracy when removing some of the seemingly important input parameters for pedestal predictions, such as fueling, JET wall type, and a parameter related to impurity seeding. We have also shown that the wall type can be almost perfectly predicted from the other parameters, and that a difference in fueling can be captured through small differences in the other input parameters. In other words, nuances in the other parameters seem to act like a proxy for some of the parameters that are not included explicitly as input parameters. However, analysis of how the small differences in these parameters are coupled to, for instance, the fueling is beyond the scope of this work, but nonetheless opens up an interesting research question for future studies.

The integrated modeling implementation may be improved further in several ways. For instance, since we lack data of the ions in the current JET pedestal database, approaches that

are more physics based may be considered for the estimation of the ion pedestal in relation to the prediction of the electron pedestal. Additionally, we currently assume a constant pedestal position and width. As there are small differences, these effects can be captured with a more sophisticated method. These points are not necessarily related to the machine learning aspect of the model, but nonetheless play an important role for the integrated modeling aspect.

From a perspective of machine learning, future work include investigations related to generalization, transfer learning, and deep learning approaches. For work related to the temporal evolution of the pedestal, recurrent neural networks will be useful. For work related to higher dimensional input parameters for the pedestal, deep neural networks may be employed to truly exploit the strengths that the field of artificial intelligence potentially can provide to fusion energy research.

An appealing aspect of the empirical approach is the lack of required assumptions about the pedestal. Even though the pedestal database from JET contains mostly type I ELM entries, similar models would be able to train on, for instance, type III ELM data, or data from other exotic ELM types. This will be particularly useful for ELM types that do not yet have a theoretical model. Additionally, similar frameworks may be employed to predict other related quantities, such as ELM frequency and energy drop in an ELM sequence, which may be relevant for SOL research and wall load research.

The inter-machine transferability of data driven models like PENN is another important aspect considering, for instance, predictive simulations of future machines like ITER. In this paper, we have only looked at JET data, and generally we can only be confident with predictions for the data domain which we have trained the model on. That said, the results of a transferability investigation will likely depend on if the pedestal physics regime remains similar when employing PENN for a different machine, which is related to different ELM regimes (with an appropriate scaling with respect to the major radius if the size of the machine is different). As mentioned above, when we train the model on one machine, we basically do not require assumptions about the physics regime of the pedestal, however if we apply the model to a different machine, this becomes important. Of course, any form of transferability test would also require assessment of how the different parameters vary between different machines since it would be difficult to go outside the training range of the parameters. In a broader sense, it would be more beneficial to train a model on data from several machines if the aim is to enable transferability. Here, machines with varied size would allow a data driven model to learn the major radius dependency as any other input parameter as well. Theoretically, a neural network model could house several physics regimes within its structure and learn when to apply which regime based on the input. However, even if a data driven pedestal model trained on several existing machines, that can distinguish between different physics regimes is produced, it will still be a challenge to make predictions for ITER and DEMO since there currently are no machines of similar size to train on. Potentially, synthetic ITER/DEMO data from

models, such as EPED, might be used to expand the database for larger machines that are not operational yet.

In conclusion, PENN shows promising results. For future work, it will be useful to explore data sets from other tokamaks, but also for tritium plasmas, and as mentioned before, data sets with a broader range of ELM types for the pedestal. Generally, when more data is available, the model may improve further, both in accuracy and broadness of applications. Nevertheless, the current state of the JET pedestal database has been shown to provide a useful application in PENN for the topics described in this paper.

Acknowledgments

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom Research and Training Programme 2014–2018 and 2019–2020 under Grant Agreement No. 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. This work has also received funding from the Swedish Research Council with the Diary Number 2020-05465, and has been carried out as a part of the Enabling Research project ‘Model for reactor relevant pedestals’ led by Christopher Ham (U4BW Work-order Number: 10189-1008).

ORCID iDs

A. Gillgren  <https://orcid.org/0000-0002-3810-2913>
 E. Fransson  <https://orcid.org/0000-0002-8747-3470>
 L. Frassinetti  <https://orcid.org/0000-0002-9546-4494>
 P. Strand  <https://orcid.org/0000-0002-8899-2598>

References

- [1] Wagner F. *et al* 1982 Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak *Phys. Rev. Lett.* **49** 1408–12
- [2] Snyder P.b., Groebner R.J., Leonard A.W., Osborne T.H. and Wilson H.R. 2009 Development and validation of a predictive model for the pedestal height *Phys. Plasmas* **16** 056118
- [3] Snyder P.b., Groebner R.J., Hughes J.W., Osborne T.H., Beurskens M., Leonard A.W., Wilson H.R. and Xu X.Q. 2011 A first-principles predictive model of the pedestal height and width: development, testing and ITER optimization with the EPED model *Nucl. Fusion* **51** 103016
- [4] Maggi C.F. *et al* 2015 Pedestal confinement and stability in JET-ILW ELMy H-modes *Nucl. Fusion* **55** 113031
- [5] Frassinetti L. *et al* 2017 Global and pedestal confinement and pedestal structure in dimensionless collisionality scans of low-triangularity H-mode plasmas in JET-ILW *Nucl. Fusion* **57** 016012
- [6] Saarelma S., Frassinetti L., Bilkova P., Challis C.D., Chankin A., Fridström R., Garzotti L., Horvath L. and Maggi C.F. 2019 Self-consistent pedestal prediction for JET-ILW in preparation of the DT campaign *Phys. Plasmas* **26** 072501
- [7] Frassinetti L. *et al* 2020 Pedestal structure, stability and scalings in JET-ILW: the EUROfusion JET-ILW pedestal database *Nucl. Fusion* **61** 016001
- [8] Kamada Y. *et al* 2000 Disappearance of giant ELMs and appearance of minute grassy ELMs in JT-60U high-triangularity discharges *Plasma Phys. Control. Fusion* **42** A247
- [9] Ozeki T., Chu M.S., Lao L.L., Taylor T.S., Chance M.S., Kinoshita S., Burrell K.H. and Stambaugh R.D. 1990 Plasma shaping, edge ballooning stability and ELM behaviour in DIII-D *Nucl. Fusion* **30** 1425
- [10] Coster D.P., Basiuk V., Pereverzev G., Kalupin D., Zagorksi R., Stankiewicz R., Huynh P. and Imbeaux F. 2010 The European transport solver *IEEE Trans. Plasma Sci.* **38** 2085
- [11] Kalupin D. *et al* 2013 Numerical analysis of JET discharges with the European transport simulator *Nucl. Fusion* **53** 123007
- [12] Kasim M.F. *et al* 2020 Up to two billion times acceleration of scientific simulations with deep neural architecture search (arXiv:2001.08055)
- [13] Meneghini O. *et al* 2017 Self-consistent core-pedestal transport simulations with neural network accelerated models *Nucl. Fusion* **57** 086034
- [14] ITER Physics Expert Group on Confinement and Transport *et al* 1999 Chapter 2: plasma confinement and transport *Nucl. Fusion* **39** 2175
- [15] Cordey J.G. (the ITPA H-Mode Database Working Group and the ITPA Pedestal Database Working Group) 2003 A two-term model of the confinement in Elmy H-modes using the global confinement and pedestal databases *Nucl. Fusion* **43** 670
- [16] Meneghini O. *et al* 2021 Neural-network accelerated coupled core-pedestal simulations with self-consistent transport of impurities and compatible with ITER IMAS *Nucl. Fusion* **61** 026006
- [17] Pasqualotto R., Nielsen P., Gowers C., Beurskens M., Kempenaars M., Carlstrom T. and Johnson D. 2004 High resolution Thomson scattering for Joint European Torus (JET) *Rev. Sci. Instrum.* **75** 3891
- [18] Kit A. *et al* 2022 Deep or Not Deep: Supervised Learning Approaches to Modeling the Pedestal Density *Plasma Physics and Controlled Fusion* (to be submitted)
- [19] Guillaume L. *et al* 2017 Imbalanced-learn: a python toolbox to tackle the curse of imbalanced datasets in machine learning *J. Mach. Learn. Res.* **18** 559–63
- [20] Beurskens M.N.A. *et al* 2014 Global and pedestal confinement in JET with a Be/W metallic wall *Nucl. Fusion* **54** 043001
- [21] Abadi M. *et al* 2015 TensorFlow: a system for large-scale machine learning software available from (tensorflow.org)
- [22] Pedregosa F. *et al* 2011 Scikit-learn: machine learning in Python *J. Mach. Learn. Res.* **12** 2825–30
- [23] Goodfellow I. *et al* 2016 Bagging and other ensemble methods *Deep Learning* (Cambridge, MA: MIT Press) (<http://deeplearningbook.org>)
- [24] Trask A. *et al* 2018 Neural arithmetic logic units *32nd Conference on Neural Information Processing Systems* (Montréal, Canada 2018) *Advances in Neural Information Processing Systems* pp 8035–44 (<https://neurips.cc/>)
- [25] Imbeaux F. *et al* 2015 Design and first applications of the ITER integrated modelling and analysis suite *Nucl. Fusion* **55** 123006
- [26] Challis C.D. *et al* 2015 Improved confinement in JET high β plasmas with an ITER-like wall *Nucl. Fusion* **55** 053031
- [27] Strand P. *et al* 2004 Comparisons of anomalous and neoclassical contributions to core particle transport in tokamak discharges *31st Conf. EPS Plasma Physics* (London, UK 2004) vol 28 (London: European Physical Society) (<https://doi.org/10.1088/0741-3335/46/12B/e01>)

- [28] Houlberg W.A., Shaing K.C., Hirshman S.P. and Zarnstorff M.C. 1997 Bootstrap current and neoclassical transport in tokamaks of arbitrary collisionality and aspect ratio *Phys. Plasmas* **4** 3230
- [29] Asunta O., Govenius J., Budny R., Gorelenkova M., Tardini G., Kurki-Suonio T., Salmi A. and Sipilä S. 2015 Modelling neutral beams in fusion devices: beamlet-based model for fast particle simulations *Comput. Phys. Commun.* **188** 33–46
- [30] Angioni C., Peeters A.G., Pereverzev G.V., Bottino A., Candy J., Dux R., Fable E., Hein T. and Waltz R.E. 2009 Gyrokinetic simulations of impurity, He ash and α particle transport and consequences on ITER transport modelling *Nucl. Fusion* **49** 055013
- [31] Tala T. *et al* 2019 Density peaking in JET-determined by fuelling or transport? *Nucl. Fusion* **59** 126030
- [32] Fable E. *et al* 2019 The role of the source versus the collisionality in predicting a reactor density profile as observed on ASDEX upgrade discharges *Nucl. Fusion* **59** 076042
- [33] Fransson E., Eriksson F., Oberparleiter M., Held M., Mordijck S., Nordman H., Salmi A., Strand P. and Tala T. 2021 Comparing particle transport in JET and DIII-D plasmas: gyrokinetic and gyrofluid modelling *Nucl. Fusion* **61** 016015
- [34] Eriksson F. *et al* 2019 Interpretative and predictive modelling of JET collisionality scans *Plasma Phys. Control. Fusion* **61** 102487