

Neural net modeling of equilibria in NSTX-U

J.T. Wai^a, M.D. Boyer^b, E. Kolemen^{a,b,*}

^a*Princeton University, Princeton, New Jersey, USA*

^b*Princeton Plasma Physics Laboratory, Princeton, New Jersey, USA*

Abstract

Neural networks (NNs) offer a path towards synthesizing and interpreting data on faster timescales than traditional physics-informed computational models. In this work we develop two neural networks relevant to equilibrium and shape control modeling, which are part of a suite of tools being developed for the National Spherical Torus Experiment-Upgrade (NSTX-U) for fast prediction, optimization, and visualization of plasma scenarios. The networks include *Eqnet*, a free-boundary equilibrium solver trained on the EFIT01 reconstruction algorithm, and *Pertnet*, which is trained on the Gspert code and predicts the non-rigid plasma response, a nonlinear term that arises in shape control modeling. The NNs are trained with different combinations of inputs and outputs in order to offer flexibility in use cases. In particular, *Eqnet* can use magnetic diagnostics as inputs and act as an EFIT-like reconstruction algorithm, or, by using pressure and current profile information the NN can act as a forward Grad-Shafranov equilibrium solver. This forward-mode version is envisioned to be implemented in the suite of tools for simulation of plasma scenarios. The reconstruction-mode version gives some performance improvements compared to the online reconstruction code real-time EFIT (RTEFIT), especially when vessel eddy currents are significant. We report strong performance for all NNs indicating that the models could reliably be used within closed-loop simulations or other applications. Some limitations are discussed.

Keywords: neural net, equilibrium, shape control

1. Introduction

In tokamaks, a critical aspect of achieving high performance plasma operation and maximizing scientific exploration is the development of fast and reliable model predictions of the plasma behavior. Fast model predictions are often a prerequisite to implementing real-time control algorithms. Additionally, if these are not available in real-time, but on a slightly slower timescale, they can be used by physics operators to inform decisions between shots. Machine learning is a tool that is gaining increasing usage in the nuclear fusion community and can be useful in this type of situation where it is desirable to quickly synthesize data into interpretable outputs. Some examples include its usage in profile predictions [1–3], turbulent transport [4], identifying scalings laws [5, 6], and predicting instabilities and disruptions [7–13].

In this work, we present two neural nets (NNs) *Eqnet* and *Pertnet* developed for the NSTX-U tokamak which are relevant to equilibrium and shape control modeling. Along with recent modeling of neutral beam heating [14] and electron profiles [3], these neural nets are part of a suite of tools being developed for fast prediction, optimization, and visualization of plasma scenarios. Closed-loop simulations of a plasma discharge are generally too slow for

usage between shots due to time-intensive steps of repeatedly solving the free-boundary Grad-Shafranov equation and linearizing around the given equilibrium. The neural nets here were developed to perform these two time-intensive tasks as a step towards fast simulation and eventually, numerical optimization of scenarios. Additionally, we explore and report on additional modes of operation of the networks that have alternative applications.

Section 2 introduces *Eqnet*, a NN-based architecture trained to predict the free-boundary plasma equilibrium based on various sets of inputs. In this work, we distinguish between several modes of operation (forward, forward-control, reconstruction, reconstruction-control) as diagrammed in fig. 1. The *forward* mode predicts the equilibrium flux surfaces using the coil currents and plasma profiles as inputs. In contrast the *forward-control* mode directly predicts control-relevant shaping parameters that are normally derived from the flux surfaces, including quantities such as the x-point locations, elongation, and outer gap distance. Similarly the *reconstruction* and *reconstruction-control* modes predict the same outputs, but use diagnostic information such as magnetic probes as inputs to the NN. Previous neural net equilibrium solvers [15–22] were developed in the spirit of reconstruction-control mode—mapping diagnostics to shaping parameters—with some exceptions such as [23] which reconstructs flux surfaces and [24] which is a fixed-boundary solver.

*Corresponding author

Email addresses: jwai@princeton.edu (J.T. Wai),
ekolemen@princeton.edu (E. Kolemen)

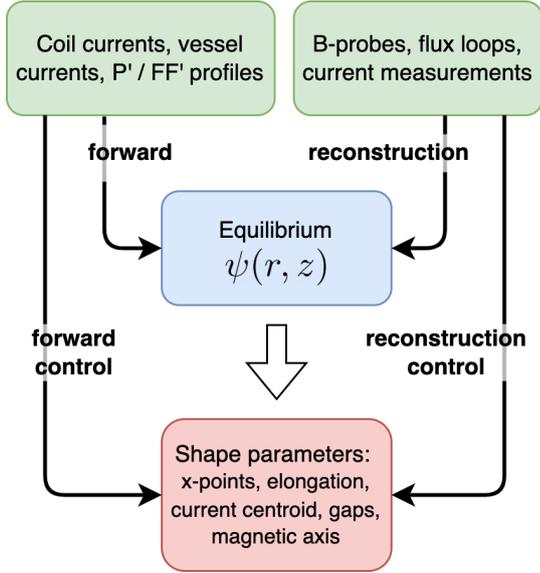


Figure 1: Diagram of the data inputs and outputs that were used to train different versions of Eqnet. The forward modes solve eq. (1) directly whereas the reconstruction modes are a diagnostic best-fit solution. The control modes predict extracted shaping parameters instead of the flux surfaces.

For control simulations, another important calculation is ‘linearizing the equilibrium’ which includes obtaining the non-rigid plasma response, how current in the plasma shifts and redistributes in response to applied fields from the poloidal field (PF) coils. It is also common to model the plasma response to the normalized internal inductance l_i and plasma poloidal beta β_p . This effect is nonlinear, equilibrium-dependent, and is addressed by codes such as Gspert/Gsupdate [25, 26] and Create-L/NL [27, 28]. In section 3 we introduce Pertnet, a NN-based architecture for identifying the plasma response that was trained on the Gspert code.

Eqnet and Pertnet were developed for the simulation goals stated above, although they have been trained to be flexible in inputs and outputs in order to demonstrate potential relevance to a variety of tasks. (To be clear, it is not a single model where inputs and outputs are toggled on and off, but multiple versions of the model trained with different inputs and outputs.) More specifically, some of the motivating applications are:

- A fast plasma simulator for between shot scenario design (Eqnet and Pertnet): Presently available plasma “flight simulators” are not fast enough to run between shots, not to mention performing numerical optimization over the simulations. The key tasks in a flight simulator are solving for equilibria, and evolving currents and profiles according to the shape control model and plasma transport. Future work will focus on integrating Eqnet and Pertnet with NN-based profile evolution [3] in order to develop this simulator.
- Increased availability and reliability of during-shot

equilibrium reconstruction (Eqnet): NSTX-U relies on real-time EFIT (RTEFIT) for shape control measurements, but RTEFIT is not available until roughly 100-150ms into the shot when the plasma current is sufficiently high. A work-around method uses an additional gap controller [29] during this early time, but this adds complexity to the system in tuning the additional controller and dealing with controller transfer. A simpler method would be to use Eqnet as a backup for when RTEFIT is unavailable, since the reconstruction mode of Eqnet can use the same inputs and outputs as RTEFIT. Additionally, RTEFIT suffers inaccuracies when vessel currents are large or dynamic plasma changes lead to insufficient convergence time, which can be improved by using Eqnet (see section 2.5.4).

- Vertical growth rate monitoring (Pertnet): One of the outputs that can be obtained from Pertnet is the growth rate of the plasma vertical instability. Due to hardware actuator constraints, any individual tokamak has a limit on the growth rate that can be controlled before the plasma is lost to a wall collision, which was a difficulty for NSTX-U since it is a spherical tokamak with highly elongated plasmas. One possibility is to use Pertnet to monitor the growth rate in real-time so that corrective actions can be taken before an instability limit is reached.
- Online updates to the shape controller (Pertnet): In most current tokamak shape control algorithms, feedback gains are calculated offline based on a single reference linearization, and these parameters are used for the entire discharge and across discharges. Using Pertnet, the equilibrium-dependent response predictions could be served to the controller rapidly (<1ms) during the shot to optimize controller performance.

Development of these tasks will be the subject of future work. The rest of this paper is divided as follows. Section 2 introduces Eqnet including the data structure, model architecture, prediction results, and a simple application problem of using Eqnet to design a novel equilibrium. Section 3 describes Pertnet, including the nonrigid plasma response model and relevant background, NN architecture, and performance results.

2. Eqnet

2.1. Problem statement

The goal of this neural net is to identify the plasma equilibrium $\psi(R, Z)$ which is the solution to the free-boundary Grad-Shafranov equation [30, 31]:

$$\begin{aligned}
\Delta^* \psi &= -\mu_0 R J_\phi \\
J_\phi &= J_\phi^{pla} + J_\phi^{ext} \\
J_\phi^{pla} &= R P'(\psi) + \frac{F F'(\psi)}{\mu_0 R}
\end{aligned} \tag{1}$$

Here ψ is the poloidal flux, J_ϕ is the toroidal current density, $P(\psi)$ is the pressure profile and $F(\psi) := R B_\phi$ is related to the poloidal current. J_ϕ^{ext} is the current density in the shaping coils and surrounding vacuum vessel structures, and serves as a boundary condition for the free-boundary problem. Analytical solutions are known for only a limited number of cases and in general eq. (1) must be solved numerically in an iterative process.

2.2. Data inputs and outputs

The ground truth training data is composed of equilibria generated from the EFIT01 reconstruction algorithm [32, 33] during the 2015-2016 NSTX-U campaign. EFIT01 is a magnetics-only equilibrium reconstruction code which is run automatically after each shot and thus offers a sizeable database of samples (~ 25000 equilibria, spaced 5-10ms apart, from 220 shots). One limitation of magnetics-only reconstructions is that the pressure and current profiles are not well-constrained. On NSTX-U there does exist a small number of discharges for which kinetic equilibria exist (pressure and current constrained by Thomson and Motional-Stark-Effect diagnostic data). However, EFIT01 is sufficiently accurate for boundary and shape control purposes and its usage is beneficial due to the higher number of samples. It is known that EFIT01 is sufficiently accurate because the NSTX-U shape controller internally uses RTEFIT, and EFIT01 is more accurate than RTEFIT. These two codes are mostly identical but without real-time constraints EFIT01 can run for more iterations until convergence.

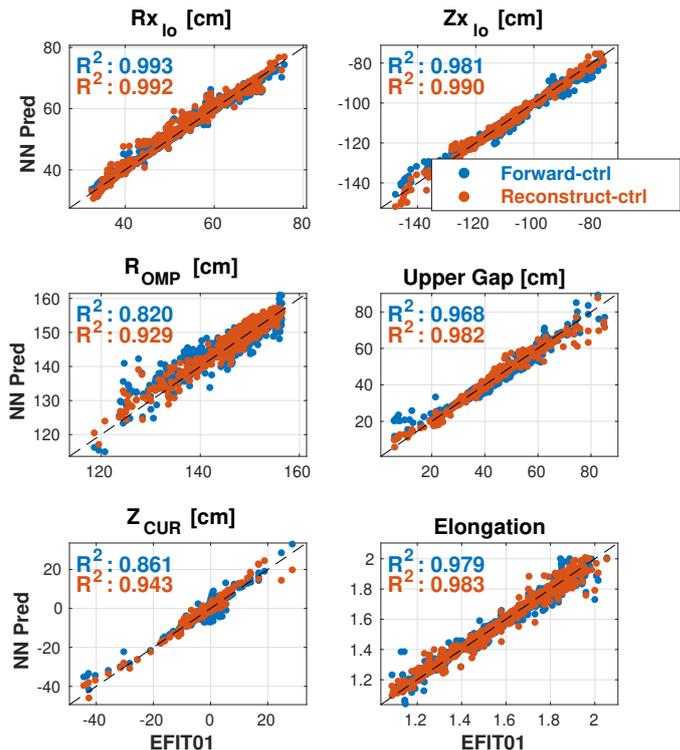


Figure 2: Regression plots of the Eqnet predictions vs EFIT01-measured parameters including: lower x-point position (Rx_{lo} , Zx_{lo}), plasma outer midplane radius R_{OMP} , upper control segment gap (see fig. 7), vertical position of the current centroid Z_{curr} , and plasma elongation. The reconstruction mode performs slightly better than forward mode as determined by R^2 coefficient.

As mentioned, the data inputs and outputs vary depending on the mode of operation. In the forward mode, the NN maps the coil currents, vacuum vessel currents, P' and FF' profiles to the equilibrium flux map. This mode is useful for simulations and optimizing coil current trajectories, since most of these data inputs are readily available from the shape control model (described in more detail in section 3.1). The shape control model integrates a state vector \mathbf{I} consisting of the plasma, coil, and vessel currents and thus all the current sources are available as inputs at each time step. Often, additional states β_p and l_i are included to track information about the pressure and current (l_i corresponds with the peaked-ness of the plasma current profile). These 2 scalars are almost information-equivalent to knowing the EFIT01 P' and FF' profiles, which by design consist of 1 and 3 basis functions respectively. (We train with the P' and FF' profiles instead of β_p, l_i for flexibility in use cases, since eq. (1) is the standard formulation).

The output of the model is the equilibrium flux values at all grid locations ψ_g , although for better data compression, the NN actually outputs only the plasma flux ψ_g^{pla} . The total flux is the sum of the plasma flux and flux from the external sources.

$$\begin{aligned}\psi_g &= \psi_g^{pla} + \psi_g^{ext} \\ \psi_g^{ext} &= M_{gc}I_c + M_{gv}I_v\end{aligned}\quad (2)$$

Here, M_{gc} and M_{gv} are the mutual inductance matrices between each grid point and the coil and vessel elements respectively, and I_c and I_v are the currents in the coil and vessel elements. The advantage of predicting the plasma flux ψ_g^{pla} is that it has less spatial variation than the total flux leading to better principal component analysis (PCA) data compression and higher accuracy in the final solution. Since the quantities that constitute ψ_g^{ext} are known, there is no disadvantage to only predicting ψ_g^{pla} .

A bit of care must be taken when considering Eqnet’s reconstruction version, since in most reconstruction applications the algorithm would only have access to measurements of the currents I_c and I_v . This can have a moderate impact: for example, on NSTX the relative error difference between measured and fitted currents was 0.56% and 4.0% for coil currents and vessel currents respectively [33]. When training the reconstruction version, we use the measured values of I_c and I_v in eq. (2). In order to still obtain the correct total flux prediction, the NN is trained to predict the plasma flux plus any flux error due to the measurement process. This model is trained using 75 B-probes, 55 flux loops, and 35 vessel current signals and we report strong accuracy in the reconstruction, slightly outperforming the forward mode (fig. 2) and RTEFIT (section 2.5.4).

The control-oriented modes do not predict the flux but instead predict parameters that are normally derived from the equilibrium. These include: the plasma minor and major radius, elongation, gap sizes of several control segments, positions of the upper and lower x-points, and position of the current centroid. The advantage of this mode is that it bypasses computational steps such locating the touch-point/x-point and tracing the boundary. A shape controller can provide feedback based directly on the NN outputs. A potential downside to this mode is that the outputs are pre-determined, offering less flexibility. If, for example, the position of the shape control segments are adjusted, the NN would have to be retrained with the new control segments.

2.3. Data processing

The 220-shot dataset is first organized chronologically by shot number, with an average ~ 100 samples per shot. The chronologically latest 10% of shots are set aside for testing. All of the shot numbers and examples referenced in this manuscript are from this test data set. The remaining 90% of shots are later divided and used for training and validation which is described in more detail in section 2.4. This data split procedure is used to mimic the standard experimental application in which a model would be trained on all available data and then need to perform adequately in a future experiment. Random splitting would likely

over-estimate the NN capabilities since samples are highly correlated within a shot, and to a lesser-degree, among nearby shots.

Each of the target and predictor variables is dimensionally reduced by performing principal component analysis (PCA) on the training set. The number of components kept is selected to achieve 99.5% captured variance, and is summarized for each variable in fig. 4. Another rule-of-thumb sometimes used to select the number of components is the scree test shown in fig. 3 for some representative signals. This test plots the singular values of the data and selects a minimum number of components based on the location of the “elbow” where the values begin to level off. We verify that the number of components selected is always greater than this point of maximum curvature.

It should be noted that we do not perform any selection process for including samples as has been done in some previous NN solvers, such as restricting to the flattop portion of the discharge. We include all available samples ranging from $t=30$ ms near the end of breakdown phase and concluding at shot termination (end of rampdown, disruption, and/or loss of control). Indeed, a sizeable fraction of shots terminate off-normal ($> 30\%$ terminate within 0.6s, while the target pulse length is generally 1 – 2s) which is a significant contributor to the range and variance in the data, and important for learning.

In order to adequately learn non-flattop equilibria, one step that proved to be necessary for good performance was performing a PCA “merging” procedure on the target variable ψ_g^{pla} . This was done by separating the dataset into rampup, flattop, and rampdown/pre-disruption phases of the discharge, performing PCA on each of these phases separately, and combining the PCA components together to form a single basis. The motivation is that the flattop phase tends to be over-represented because I_p is constant and the flux surfaces and plasma position are mostly constant. In contrast, much of the important variation occurs during rampup and rampdown/pre-disruption. Naively using PCA on the combined data did not represent the rampup or rampdown equilibria well, as determined by visual inspection of the projected flux surfaces. This problem is not fixed by oversampling the rampup and rampdown equilibria, likely because these samples have lower I_p and therefore lower magnitude of ψ_g^{pla} so they still do not sufficiently contribute to the variance in the data matrix.

The PCA-merging step is performed as follows: suppose we have a data matrix X of dimensions $n_{features} \times n_{samples}$ with sample mean μ , and that we have performed PCA on X and have a set of principal components U and coefficients W . We recall that the data matrix can be projected from the components and coefficients as

$$X = UW + \mu \quad (3)$$

Thus for two (truncated) PCA projections $X_1 = U_1W_1 + \mu_1$ and $X_2 = U_2W_2 + \mu_2$, the combined projection $[X_1 \ X_2]$ is within the span of the combined components and means

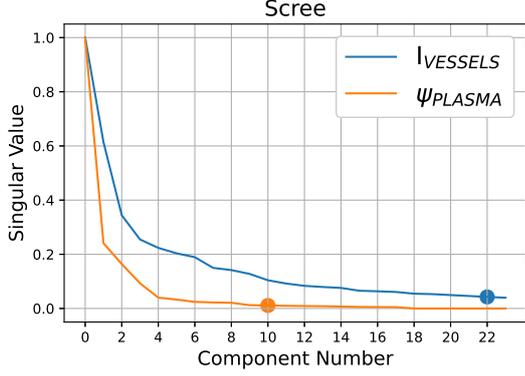


Figure 3: A scree test for 2 representative variables. The number of components is verified to be larger than the “elbow” location.

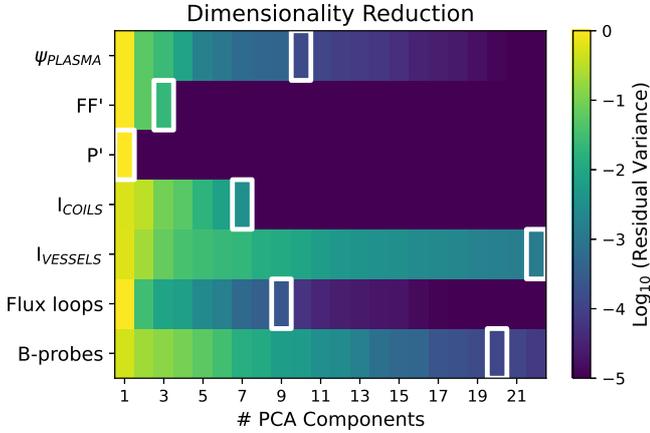


Figure 4: Dimensionality reduction of the Eqnet dataset using Principal Component Analysis (PCA). The outlined boxes indicate 99.5% captured variance and the number of components used for each variable.

$\bar{U} := [U_1 \ U_2 \ \mu_1 \ \mu_2]$. Thus we can write

$$[X_1 \ X_2] = \underbrace{[U_1 \ U_2 \ \mu_1 \ \mu_2]}_{\bar{U}} \bar{W} \quad (4)$$

for some set of weights \bar{W} , or a version with nonzero mean which is in the same PCA-like form as eq. (3).

$$[X_1 \ X_2] = \underbrace{[U_1 \ U_2 \ \mu_1 \ \mu_2]}_{\bar{U}} \bar{W} + \mu^* \quad (5)$$

$$\mu^* = (\mu_1 + \mu_2)/2$$

Equation (5) is not a valid PCA projection because the columns of \bar{U} are not orthogonal, but we can obtain a set of orthogonal components U^* by performing, for example, a singular value decomposition of \bar{U} . The merged coefficients are then $W^* = U^{*T}([X_1 \ X_2] - \mu^*)$. This simple preprocessing step proved to be sufficient for compressing data across a wide range of samples. The final result compresses the ψ_g^{pla} grid of size $[65 \times 65]$ to just 10 PCA

coefficients.

Figure 4 summarizes data compression results for the various inputs and outputs. For the flux values, significant dimensionality reduction is achieved. As mentioned, the P’ and FF’ profiles consist of 1 and 3 basis functions respectively, which is captured by this mode decomposition. The highest dimensional inputs are related to the vessel currents and B-probe diagnostics. After dimensionality reduction, we normalize the inputs to have zero mean and unit variance, to arrive at the neural network inputs. The standardization and PCA decompositions are then inverted to map neural network outputs to physical outputs.

2.4. Model architecture and hyperparameter tuning

The model architecture including data processing steps is shown in fig. 5. The core of the model uses fully-connected artificial neural nets (ANNs) which are created and trained in PyTorch with the Adam optimizer. We use a small (5-member) ensemble of NNs and average their predictions, before inverse-transforming the results back to the original data-space. Recall that 10% of shots are reserved for testing with the remaining 90% shots going to training and validation. Each NN ensemble member is trained on a random subset (90%) of shots from within this training/validation set (i.e. an individual NN will see 90% of the 90%). This ensemble method is used for two reasons—first, to increase the robustness of the final prediction, since the various applications tend to need stable predictions, and second, as a method for determining statistical properties of the predictions. (This ensemble method is essentially 5-fold cross-validation, with each ensemble member representing one iteration of the cross-validation procedure.) As will be shown in section 2.5.3, the agreement between ensemble members can be a useful tool for gauging uncertainty in the prediction, and actually does correlate with underlying uncertainty in the equilibrium reconstruction.

The network hyperparameters are chosen by performing scans of the number of hidden layers, hidden layer size, dropout fraction, optimizer learn rate, batch size, and non-linearity activation function. Results of this procedure are shown in fig. 6. The hyperparameter landscape is fairly convex, making it straightforward to identify a suitable parameter combination. We note that the NN tends to perform better with larger network sizes (depth and breadth), zero dropout, and small batch sizes. We also note a jump in performance when using the ELU activation function as opposed to the standard ReLU.

2.5. Performance

2.5.1. Model timing

A key motivation for this work is the computational speed of neural networks. The time to train an individual NN for 300 epochs is roughly 8-10 hours on PPPL’s portal cluster (single core AMD Opteron 6136 2.4GHz CPU). More important is the inference time which was measured

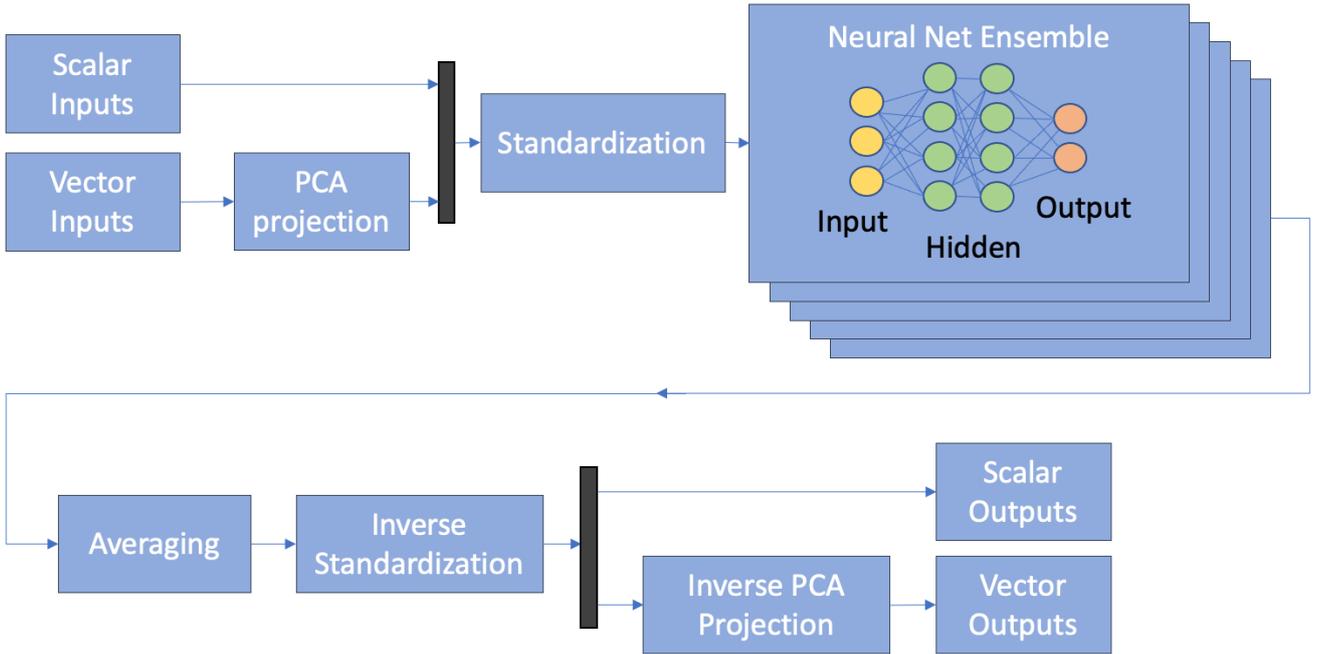


Figure 5: Model architecture for Eqnet and Pertnet including data processing steps. The core of the model is an ensemble of fully-connected networks.

on a Macbook Pro 3.5GHz Dual-Core Intel i7 CPU and ranges from $810\mu\text{s}$ to $940\mu\text{s}$ depending on the particular NN version. This was computed in Python and represents the inference time for a single NN, since it would be natural to parallelize the inference of each NN ensemble member. This is already fast enough to support real-time operation (approximately 1ms, which is also approximately the time needed by RTEFIT to complete a cycle), although additional speed increases are expected for when the model is translated into a compiled programming language. Therefore we expect any real-time applications could be supported without issue.

2.5.2. General performance

Figure 7 shows some representative results for NN flux surface predictions (forward-mode) taken from the test dataset. Also depicted are the x-points, magnetic axis, and control segments used in computing the upper and lower gap. The 4 samples shown represent different categories of equilibrium shapes such as limited plasma produced during rampup, lower single-null (LSN) during flattop, upper single-null (USN) during flattop, and a LSN from a disrupting plasma. In the first three cases the flux surfaces are mostly indistinguishable by eye at this scale. One does observe errors in the disruption sample. Since this type of event is off-normal it is less well-represented in the training data. This suggests that for less-represented samples or shapes that have not yet been explored on NSTX-U (e.g. negative triangularity), there is room for future improvement by obtaining more experimental data or sup-

plementing with synthetic data.

A regression test of shaping parameters is shown in fig. 2 for the control-oriented NNs. We observe that the reconstruction mode outperforms the forward mode slightly in most variables as measured by the R^2 coefficient. ($R^2 = 1$ is perfect prediction) A possible explanation is that the reconstruction mode has significantly more NN inputs to learn from, including highly localized information from the diagnostics. The hardest variables to predict are the outer midplane radius and current centroid position with R^2 values 0.820/0.929 and 0.861/0.943 respectively. As another quantitative measure of accuracy we report absolute prediction error values at the 50th (median) and 90th percentile levels in table 1. For most geometric features the median errors are generally only a few mm, and the 90th percentile errors are about 1cm.

Prediction errors

Model	Percentile	$R_{x,LO}$ [mm]	$Z_{x,LO}$ [mm]	R_{CUR} [mm]	Z_{CUR} [mm]	R_{OMP} [mm]	Gap _{,UP} [mm]	Gap _{,LO} [mm]
Forward	50	1.7	4.6	5.9	2.4	4.4	4.2	3.2
Reconstruction	50	1.6	2.9	5.7	2.4	3.5	2.5	3.5
Dead B-probe	50	5.8	8.9	10.4	8.3	8.3	10.4	12.8
Forward	90	6.0	11.1	15.9	10.3	15.7	12.2	15.2
Reconstruction	90	5.3	9.6	13.8	6.7	12.1	7.8	10.0
Dead B-probe	90	11.6	17.2	22.4	17.9	27.0	21.2	29.8

Table 1: Prediction errors for several models computed at the 50th and 90th percentiles. Most errors are less than 1cm. Zeroing out a magnetic probe results in an approximate doubling of errors.

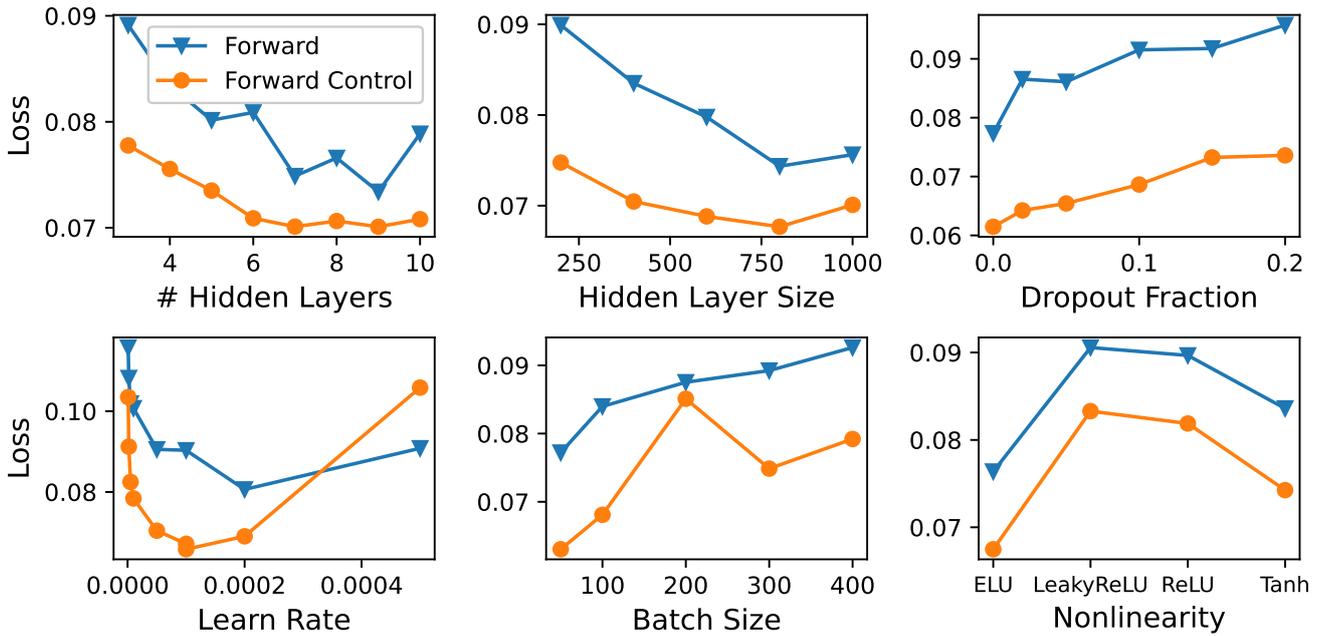


Figure 6: Hyperparameter grid scans used to identify the best performing Eqnet model. The loss value is computed using the validation dataset. The landscape is fairly convex making it straightforward to identify good parameter combinations.

While the large majority of samples are predicted accurately, it is true that some samples have errors of multiple cm which is high for shape control. The time trace predictions shown in fig. 8 give some qualifying insight. From the R_{OMP} and Z_{CUR} plots we observe multiple oscillations in the signals, which tend to occur very early in the shot (rapidly changing conditions during current rampup) or near the end of a shot (from disruption, loss of control, or rapidly varying conditions during rampdown). In fact, for errors ≥ 3 cm, 68% of bad NN predictions occur during the first or last 100ms of the shot, and 100% within the first or last 200ms. These dynamic features are representative of the NSTX-U campaign and is in part due to the control difficulties of associated with a short-pulse, highly-elongated, unstable plasma. The worst errors always occur during these fast transient events, and represents real physical uncertainty since on this fast time scale induced vessel currents do not have time to decay, and have a shielding effect on the diagnostics. In other words the underlying EFIT01 equilibria also have more uncertainty at these times.

The timetraces in fig. 8 includes predictions from the forward version, where shaping parameters have been derived from the flux surfaces, and from the forward-control version (direct prediction of the parameters). Both sets of NN predictions match each other and EFIT01 closely. In these figures, the control version has slightly less error overall but this is not universal. Neither version is substantially superior to the other so the main tradeoff is in terms of the desired output type, with flux surface prediction being the more flexible option.

2.5.3. Ensemble agreement as an indicator of prediction uncertainty

We advocate that the disagreement between ensemble members can be used as an indicator for prediction uncertainty. The logic is that since the NNs were trained with different shot data and initialization weights, they will tend to give different predictions for samples that are hard to predict. This is not a direct conversion; often the difference between NNs is smaller than the difference of the NN vs EFIT01. The shaded red areas in fig. 8 is a measure of agreement between the underlying NN ensemble members in Eqnet, plotted at ± 3 standard deviations. We observe the disagreement between NNs is not constant but varies throughout the shot, for example, increasing substantially for all signals in the last 300ms of the shot. This is also the time period where prediction errors are larger, and where it is known that there is more uncertainty in the reconstruction (due to the fast dynamics at beginning or end of shot). This information could be used to help with safe deployment of NN-based applications, for example, by relying on Eqnet only if the disagreement is below a certain threshold.

2.5.4. Comparison vs RTEFIT

As discussed in the introduction, one of the possible applications for Eqnet is to perform real-time equilibrium reconstruction during shots. On NSTX-U this task is performed by the RTEFIT algorithm. The reconstruction version of Eqnet can functionally act like RTEFIT since it is trained to predict equilibria using the same magnetic signals as inputs. Additionally, Eqnet is trained on the

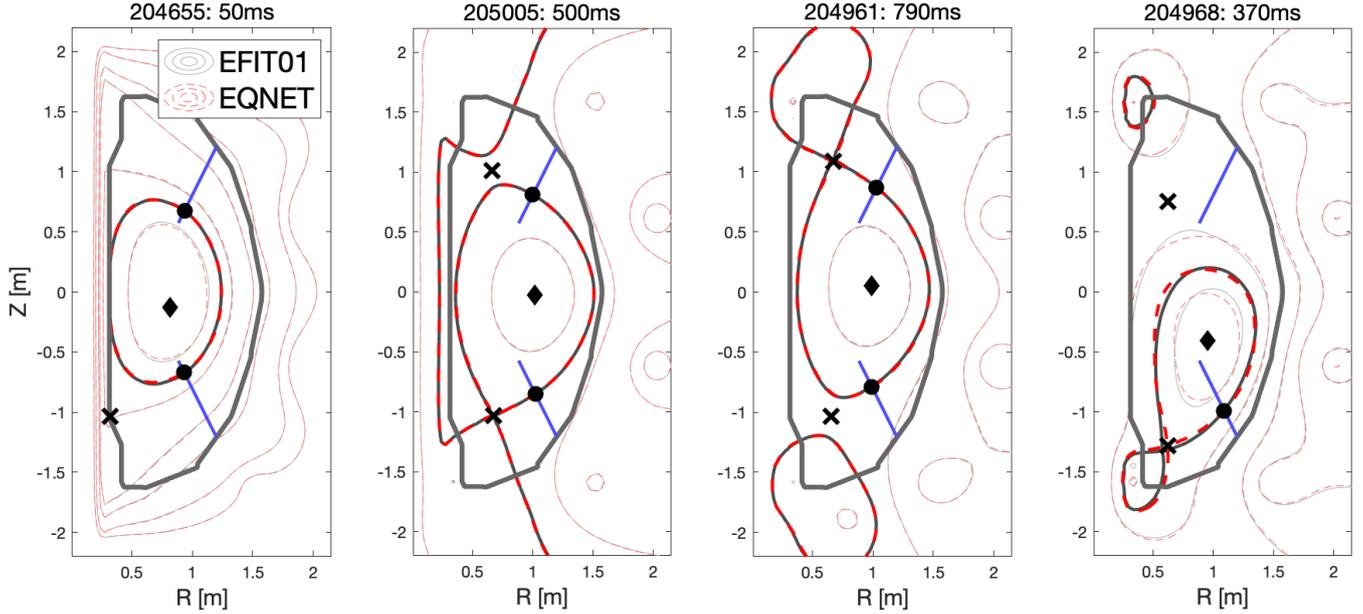


Figure 7: Equilibrium flux surfaces for several test shots as determined by EFIT01 and the Eqnet neural network. The boundary contour is shown in bold. Also shown are control segments (blue), magnetic axis (black diamond), and x-points.

EFIT01 outputs which is a more accurate reconstruction than RTEFIT, so there is a possibility for Eqnet to outperform RTEFIT by some standards.

To clarify the relationship, EFIT01 is the offline magnetics-only reconstruction and RTEFIT is the online real-time version. There are a few key differences between these codes. Most importantly, EFIT01 performs more iterations and achieves better convergence of the Grad-Shafranov solution, whereas RTEFIT only has time to compute one iteration per equilibrium. Second, EFIT01 can perform non-causal filtering of the diagnostics since it has access to future measurements for applying, e.g., drift correction. Lastly, RTEFIT is only available starting at 100-150ms into the shot whereas EFIT01 is run offline for times larger than ~ 30 ms.

One source of error for RTEFIT stems from occasional failure to reach convergence. This failure is more likely to occur during large changes to the plasma, such as strong oscillations, or if the boundary-defining point is switching between various locations. Another source of error is related to the vessel eddy currents. As noted in the original NSTX implementation, the vessel current measurements are computed via Ohm’s law $I = V_{meas}/R$ neglecting dynamical effects [33, 34], which is accurate enough to prevent “an unphysical solution, yet is large enough to permit the solution to deviate from the computed value during periods of fast current transients.” [33] Vessel currents can also have the effect of temporarily shielding magnetic sensors from changes that are happening to the plasma, negatively impacting the reconstruction.

Figure 9 compares several reconstructed signals from EFIT01, Eqnet, and RTEFIT for a few shots from the

test dataset, and highlights that Eqnet can outperform RTEFIT especially in relation to the two error sources (convergence and induced vessel currents).

The first panel 9a depicts R_{OMP} the radius of the plasma boundary at the outer midplane. Eqnet matches EFIT01 closely, while RTEFIT deviates from EFIT01 at $t = 0.10 - 0.13$ and again at $t = 0.20 - 0.29$ s. The former period is due to RTEFIT being switched on and requiring several iterations to converge. The latter error is due to the vessel current effect. During this time $0.20 - 0.29$ s the plasma was changing from limited to diverted and moving radially outward. This radial plasma movement induces strong vessel currents, and RTEFIT does not coincide with EFIT01 until this motion stops and eddy currents die down.

In 9b we observe the measurements of the upper gap during a shot that was experiencing undesired vertical oscillations. The RTEFIT signal is phase-shifted behind EFIT01 by 5ms indicating the length of time the vessel dynamics are exerting strong influence. By contrast Eqnet is in-phase with EFIT01 demonstrating that the neural net has learned how to compensate for these dynamic effects.

Convergence-related issues are shown in 9c and 9d. During the time period shown in 9c, the plasma was oscillating significantly, and the boundary point was switching rapidly between an inner-wall touch point, upper x-point, and outer-wall touch point. This led to insufficient convergence for RTEFIT, and so Z_{maxis} , the vertical position of the magnetic axis swings wildly by > 50 cm. (Note that RTEFIT will aggressively move the magnetic axis, but the boundary measurements are not changing by quite this magnitude). One of these equilibria at $t = 1.23$ s is plot-

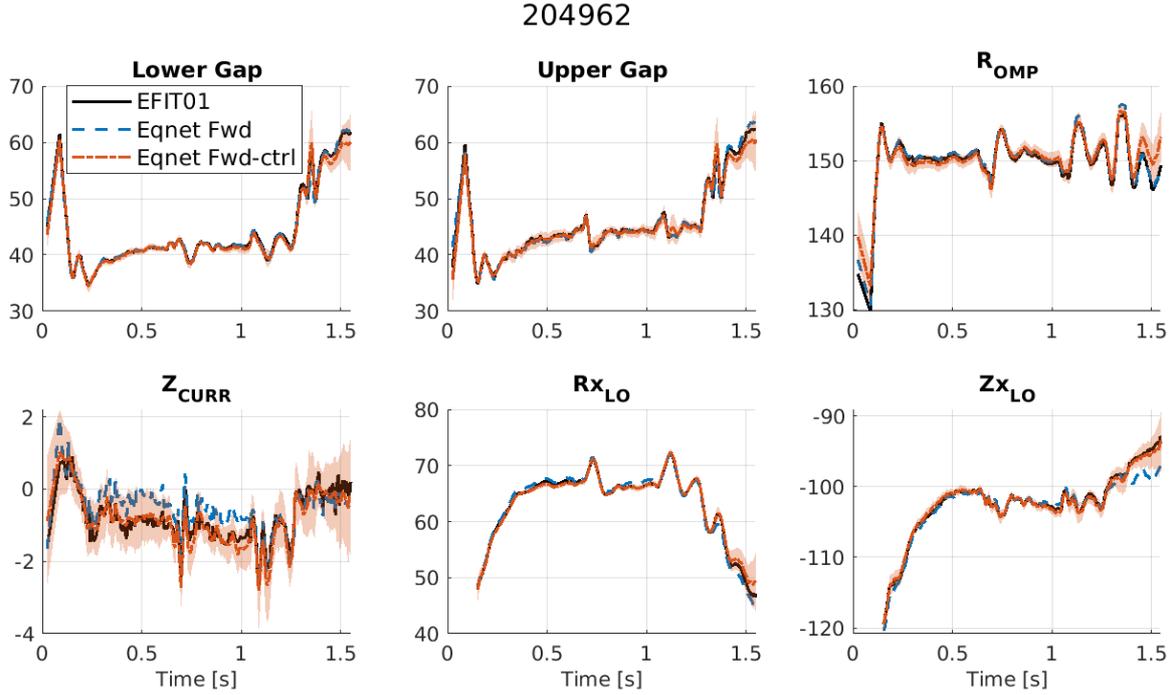


Figure 8: Time traces of several geometric quantities for test dataset shot 204125, as measured by EFIT01 and predicted by Eqnet. The lower and upper gaps are the intersection of the plasma boundary with the control segments as shown in fig. 7. The shaded red region represents ± 3 standard deviations between the predictions of NN ensemble members.

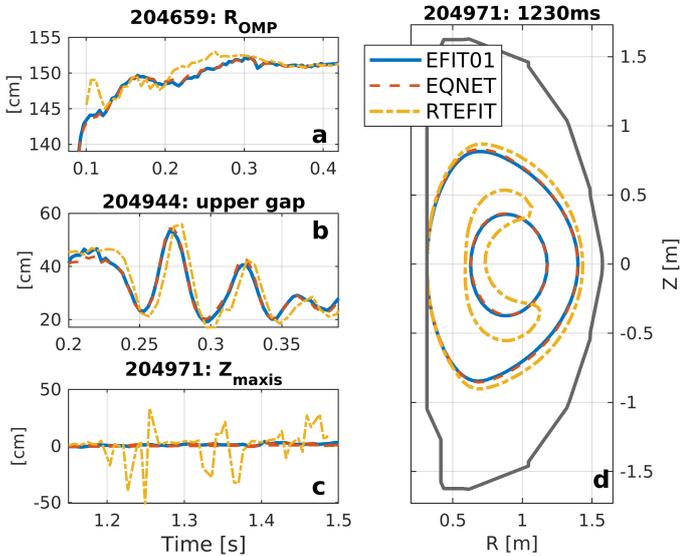


Figure 9: Comparison of Eqnet and RTEFIT. The RTEFIT error and 5ms phase shift in (a) and (b) is due to dynamic events inducing strong vessel eddy currents. Eqnet has learned to compensate for this. The RTEFIT mismatch in (c) and (d) is due to insufficient convergence.

ted in 9d. While the boundary is relatively close to the EFIT01 boundary, it is clear that the internal flux contours are very different.

These benefits of a NN-based approach must be weighed against disadvantages such as a weaker ability to generalize, and the need for high-quality training data that may

not exist for, e.g. a newly built machine. While a NN may not be immediately usable during new machine commissioning, it could increasingly offer benefits as a reconstruction backup as data is collected over time.

2.5.5. Robustness to diagnostic failure

Another aspect to consider for deployment is robustness of the NN to sensor failure. For example, if one of the magnetic probes dies, it would be beneficial if experiments could continue as normal with only a slight decrease in performance. To investigate this issue, we trained the reconstruction-version with all sensor inputs, and then zeroed out one of the magnetic probes and measured the predictions. This was repeated for all 75 magnetic probes, and the average of these results is recorded in table 1. We note that most errors have roughly doubled compared to the original reconstruction-mode predictions. For many of the signals, the 90th percentile errors are now in the range of 2-3cm. This would likely be workable, but with reduced operational capability. For example, the R_{OMP} target would likely have to decrease to give margin for wall clearance. These results were obtained using light (1%) dropout during training to randomly zero out some magnetic sensor inputs. It is likely possible to improve these results further by optimizing the dropout rate, to strike a balance between peak performance and robustness.

2.6. Design Application

As a simple application task we attempt to solve an inverse equilibrium design problem using the Eqnet forward-

version. The design problem is: given a target shape, what are the coil currents that will transform an equilibrium into that shape. Notationally, we have target shape \mathbf{y} and the neural net is a map from inputs \mathbf{x} to an equilibrium $f(\mathbf{x})$. (We restrict to only modifying coil currents, not all of the network inputs.) The design procedure is carried out by using the Matlab function `fminsearch`, which is a standard optimization algorithm based on the Nelder-Mead method. At each iteration, the search makes a call to the NN (which evaluates $f(\mathbf{x})$) and adjusts the inputs to minimize the cost function $J := (f(\mathbf{x}) - \mathbf{y})^2$.

The target shape is obtained by beginning with a LSN diverted equilibrium and increasing the elongation from 1.65 to 1.95. The target boundary also has a smaller minor radius 54cm vs 58cm and larger inner and outer gaps. The initial, target, and final shapes are shown in fig. 11, and coil currents in fig. 10. Note that the boundary “final” boundary is not the NN-prediction (which exactly matches the target), but the result of a standard free-boundary calculation using the NN-designed coil currents. The equilibrium produced by the NN is indeed much closer to the target, with some small differences in the boundary at the lower and upper. Given that this is a novel equilibrium for the NN, which has not seen this exact equilibrium before and was not generated via EFIT, we expect this gives a reasonable estimate of generalization error. This level is also acceptable for feedforward planning applications, since feedback control can correct for small errors.

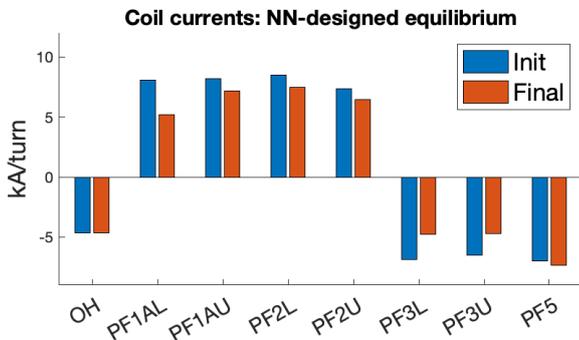


Figure 10: Coil currents designed by with `fminsearch` + Eqnet in order to construct an elongated equilibrium.

3. Pertnet

3.1. Plasma response model

“Pertnet” is a neural network we have designed to predict the nonrigid plasma response, which is a term related to the Grad-Shafranov equation that arises in the development of the shape control model. This section includes some background context on the shape control model and plasma response. For further detail on shape control, we refer readers to the review paper by Ambrosino and Albanese [35]. The nonrigid plasma response describes how the plasma current shifts and redistributes in response to

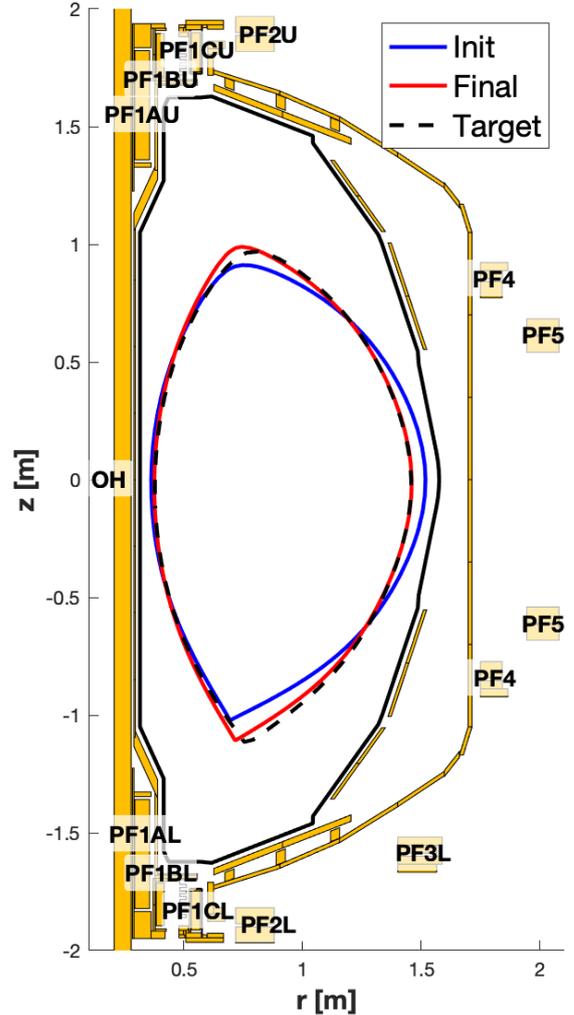


Figure 11: Equilibrium designed with `fminsearch` + Eqnet, by adjusting the coil currents in order to match the target boundary. The target shape is an elongated $\kappa = 1.95$ LSN equilibrium.

changes in the currents of surrounding conductors. The response can be written in units of current or units of flux which are related by the grid mutual inductance:

$$\frac{\partial \psi_g^{pla}}{\partial I_i} = M_{gg} \frac{\partial I_g^{pla}}{\partial I_i} \quad (6)$$

Here, ψ_g^{pla} is the grid flux due to plasma currents, I_g^{pla} is the plasma current distribution, and M_{gg} is the mutual inductance matrix for the grid. We use the “pla” superscript to distinguish flux that is due to plasma sources, since the total flux at any point is due to both plasma and external sources. The physical interpretation is that when the current in an external conductor is perturbed, this changes the background magnetic field. Since the plasma is carrying current it experiences a Lorentz force

that causes it to shift back into equilibrium. We rely on the Gspert code [25] to obtain the plasma response, which is based on a linearization to the Grad-Shafranov equation. In addition to obtaining the response due to current changes, it is often useful to obtain the response due to changes in global parameters such as β_p and l_i . These are included in the Pertnet predictions, although not explicitly discussed here.

The plasma response arises in several places in the shape control model. For example, the model often begins by writing a circuit equation for all the conducting elements within the tokamak:

$$\begin{aligned} v_i &= R_i I_i + \dot{\psi}_i \\ &= R_i I_i + M_{ij} \dot{I}_j + \frac{\partial \psi_i^{pla}}{\partial I_j} \dot{I}_j \\ &= R_i I_i + M_{ij} \dot{I}_j + M_{ig} \frac{\partial I_g^{pla}}{\partial I_j} \dot{I}_j \end{aligned} \quad (7)$$

where v_i is the applied voltage from external power supplies, R_i is the resistance of each element. We see that the plasma response contributes to the flux change at conductors and must be accounted for when modeling current evolution. Additionally, this term is needed for modelling how various feedback parameters (x-points, current centroid) change with respect to the inputs. Consider some inputs x that define the equilibrium (x includes the coil currents), and some shaping parameters y that are derived from the equilibrium flux. This relation is

$$y = f(\psi(x)) \quad (8)$$

$f(\cdot)$ is the mapping from equilibrium to parameters, and $\psi(\cdot)$ is the map from inputs to equilibrium. By the chain rule, in order to see how the inputs affect the shape parameters we need access to the plasma response.

$$\frac{\partial y}{\partial x} = \frac{\partial f}{\partial \psi_g} \left(\frac{\partial \psi_g^{ext}}{\partial x} + \frac{\partial \psi_g^{pla}}{\partial x} \right) \quad (9)$$

We train 2 versions of the Pertnet model – a standard flux mode that calculates the flux response $\frac{\partial \psi_g^{pla}}{\partial x}$ and a control-oriented mode that directly estimates the parameter responses $\frac{\partial y}{\partial x}$. To illustrate this mapping from flux response to parameter response via eq. (9), we show the derivation for 3 parameters (current centroid response, x-point response, and vertical growth rate) which are the main parameters predicted by the Pertnet control mode:

3.1.1. Current centroid response

The current centroid position is the current-weighted average of each coordinate over the plasma domain Ω :

$$Z_{cur} = \frac{1}{I_p} \sum_{\Omega} Z_g I_g \quad (10)$$

Applying eq. (9) and eq. (6) gives

$$\frac{\partial Z_{cur}}{\partial x} = \frac{1}{I_p} \sum_{\Omega} Z_g \left(M_{gg}^{-1} \frac{\partial \psi_{g,pla}}{\partial x} \right) \quad (11)$$

3.1.2. X-point response

The x-point is the location at which the flux gradient is equal to zero, so finding the x-point is equivalent to finding the root of the equation $F(\mathbf{s}) = [\psi_r(\mathbf{s}) \ \psi_z(\mathbf{s})]^T$. Using Newton's method with the Jacobian of F defined J_F , one could find the root by successively applying updates as

$$\Delta \mathbf{s} = -J_F(\mathbf{s})^{-1} F(\mathbf{s}) \quad (12)$$

If we have already converged to an x-point, but then apply a perturbation to the flux, the response of the x-point is

$$\frac{\partial (r_x, z_x)}{\partial x} = \frac{\partial \Delta \mathbf{s}}{\partial x} = -J_F(\mathbf{s})^{-1} \begin{bmatrix} \frac{\partial \psi_r}{\partial x} \\ \frac{\partial \psi_z}{\partial x} \end{bmatrix} \quad (13)$$

which depends explicitly on the plasma response via the last terms since

$$\frac{\partial \psi_r}{\partial x} = \frac{\partial}{\partial r} \left(\frac{\partial \psi^{ext}}{\partial x} + \frac{\partial \psi^{pla}}{\partial x} \right) \quad (14)$$

3.1.3. Vertical growth rate

The vertical growth rate is an instability that can be related to the evolution of currents in conducting structures. Defining the response term in eq. (7) as $\frac{\partial \psi_i^{pla}}{\partial I_j} := X_{ij}$, this equation can be rewritten in matrix form as

$$\mathbf{v} = \mathbf{R}\mathbf{I} + (\mathbf{M} + \mathbf{X})\dot{\mathbf{I}} \quad (15)$$

or equivalently

$$\begin{aligned} \dot{\mathbf{I}} &= \mathbf{A}\mathbf{I} + \mathbf{B}\mathbf{v} \\ \mathbf{A} &:= -(\mathbf{M} + \mathbf{X})^{-1}\mathbf{R} \\ \mathbf{B} &:= (\mathbf{M} + \mathbf{X})^{-1} \end{aligned} \quad (16)$$

which is the standard form of the shape control state space model. The vertical growth rate is taken as the largest unstable eigenvalue of the \mathbf{A} matrix.

$$\gamma = \max(\text{real}(\text{eig}(\mathbf{A}))) \quad (17)$$

3.2. Data Inputs and Outputs

Similar to Eqnet, we train 2 versions of the Pertnet model – a “standard” flux mode that calculates the flux response $\frac{\partial \psi_g^{pla}}{\partial x}$ and a “control” mode that predicts the parameter responses $\frac{\partial y}{\partial x}$. The standard mode is more general in that all aspects of the shape control model can be derived from it, albeit with more computational effort. On the other hand, control mode skips some computational

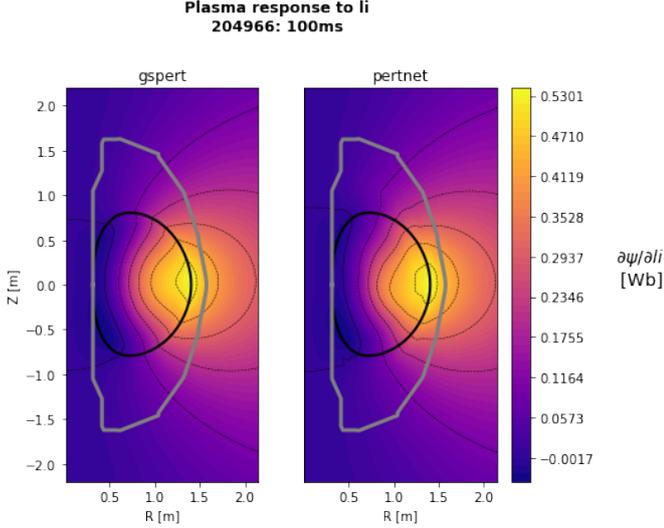


Figure 12: Comparison of the plasma response to an l_i perturbation as calculated by the Gspert code and pertnet neural network approximator. The plasma boundary is overlaid in black.

steps in finding the shaping parameter response directly. This would make it useful for experiments in providing fast estimates to the shape controller. A potential downside is that if the desired outputs are modified then the NN requires retraining. Additionally, the control mode does not predict all aspects of the shape control model, such as what the full \mathbf{X} matrix is.

For both modes, the inputs are a complete description of the equilibrium. We assume more inputs here than we did for Eqnet because we are linearizing around the current equilibrium and can thus assume access to it, either via a NN prediction or reconstruction algorithm. The inputs include: $\psi(r, z)$ as calculated by EFIT01, the EFIT01 pressure $P'(\psi)$ and $FF'(\psi)$ profiles, current distribution J_ϕ , coil and vessel currents, I_p , coordinates of the boundary contour, ψ at the magnetic axis and boundary, and the (R, Z) position of the magnetic axis, x-points, and current centroid.

For the standard mode, the output is the flux response on the grid. Note that from eq. (6) we could predict in units of flux or units of currents. For the NN, it is advantageous to predict in flux units. This is because the current response can contain highly localized information about the current variation in nearby grid cells. By contrast the flux response has less spatial variation and has better data compression. In other words, the grid mutual inductance M_{gg} acts like a smoothing or averaging filter when transforming from current to flux.

3.3. Architecture and hyperparameter tuning

We use the same set of ~ 25000 equilibria samples and chronological train-validation-test data splitting method used in Eqnet (section 2). For each sample, we run Gspert to find the plasma response to the 8 active PF coils, 40 vessel circuits, I_p , β_p , and l_i . For the standard flux mode,

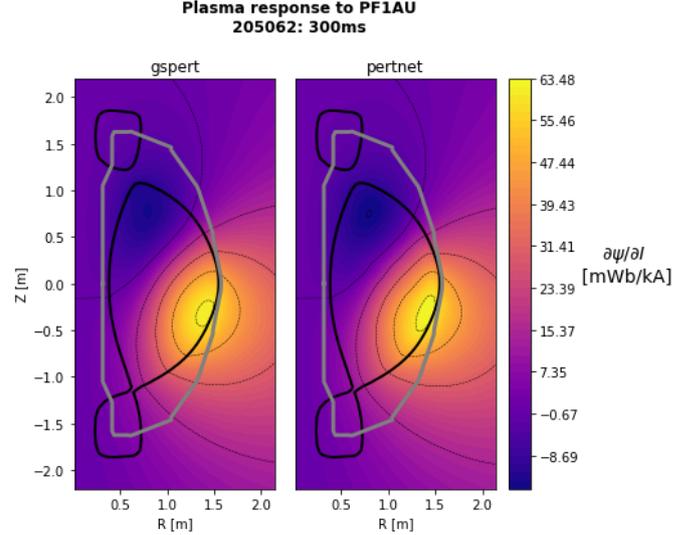


Figure 13: Comparison of the plasma response to a PF1AU current perturbation as calculated by the Gspert code and pertnet neural network approximator. The plasma boundary is overlaid in black.

we train an individual NN to predict the response of each quantity. This performs better than predicting the flux response to all inputs simultaneously, which does not dimensionally reduce very well, since each response quantity can vary significantly in magnitude. For the control mode, which predicts responses of individual parameters that are of much lower dimension than the flux response, it was found that a single NN was sufficient.

Principal component analysis is used to compress all the inputs and outputs, choosing the number of components to capture 99.5% variance, with results of this procedure shown in fig. 17. Only some of the response quantities are plotted. Note that the equilibrium currents J_ϕ has a much larger representation than the flux ψ , again indicating that the flux is generally smoother than the current. All response variables have the similar order of magnitude (~ 10 components). Some variables, such as the ohmic (OH) coil have expectedly higher variation. The OH coil spans the entire vertical length of NSTX-U, and thus can interact with the plasma from multiple angles depending on plasma position. After PCA, the coefficients are then normalized to zero mean and unit variance. The full architecture of the model is identical to Eqnet and shown in fig. 5.

Training is performed in Pytorch with the Adam optimizer. Hyperparameters are selected by performing grid scans over the various model parameters such as size and depth, with results shown in fig. 18. Early on, it was observed visually that using the L1 loss (absolute error) function performed better than the standard L2 loss (sum of squares). This is theorized to be a result of noise and outliers in the dataset. In particular, a small fraction of samples were observed to have spurious results due to running Gspert automatically on all 25k+ equilibria. These

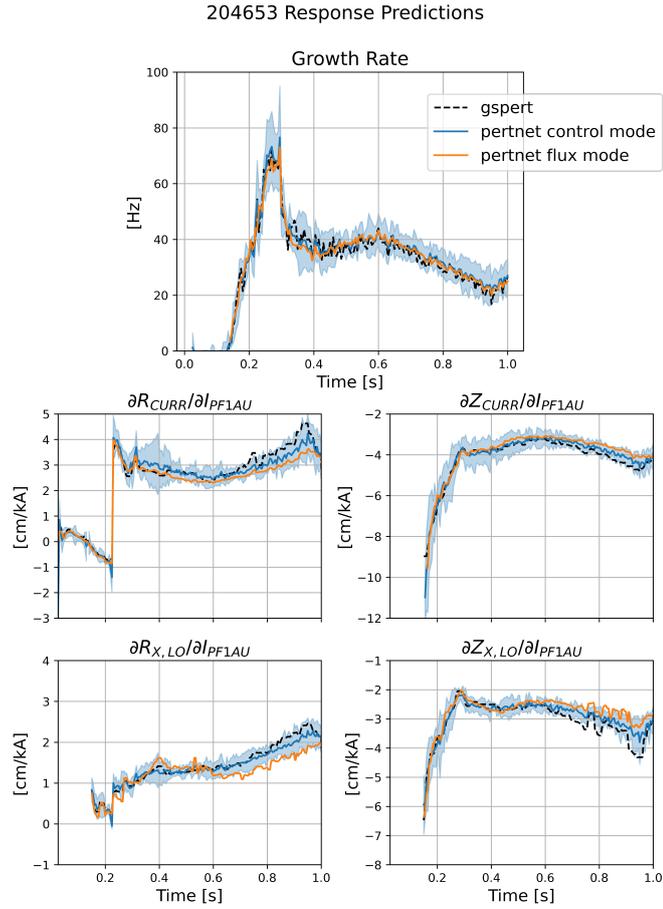


Figure 14: Predictions of the plasma response for several parameters (growth rate, current centroid response to PF1AU, and x-point position response to PF1AU) as calculated by Gspert and versions of Pertnet. The Pertnet control mode directly predicts these parameters, while the pertnet flux mode predicts the flux response which maps to the parameter responses. The shaded blue region corresponds to ± 4 standard deviations of the disagreement between NN ensemble members.

include errors such as not converging fully to features of interest, or not identifying the boundary correctly in a situation with multiple x-points and touch points. An attempt was made to perform automated filtering of the outputs, but manual inspection of all results was not performed. In addition, the samples are theorized to have a lot of noise due to the inherent noise associated with taking derivatives. Since L2 penalizes squared errors, the network was probably learning too much from predictions with large errors that should have been identified as outliers. Another possible indicator of stochasticity is that the performance improves by including 5 – 10% dropout on the network weights during training. The final network is smaller than Eqnet with the standard/control modes having 4/5 layers with hidden layer size 400/800. Note that the hyperparameters differ significantly between the modes because they are actually problems with relatively different structure, since the standard mode has a NN for each flux response,

Pertnet-Std Regression

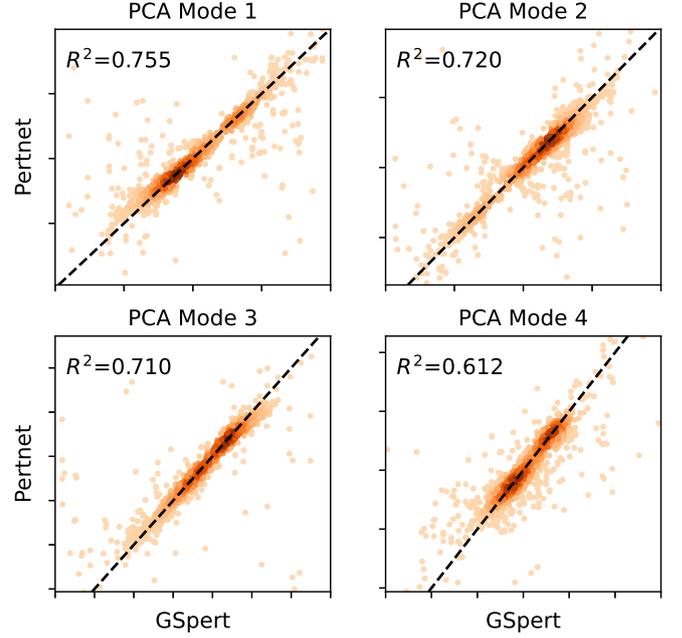


Figure 15: Regression of the Pertnet flux prediction vs calculated values from Gspert. The data shown is the test dataset for the response to PF5. Each plot shows the coefficient of the corresponding PCA mode of the flux response.

and the control mode is a single NN that predicts all parameter responses.

3.4. Performance

3.4.1. Model timing

Using the same setup as before (Macbook 3.5GHz Dual-Core Intel i7 CPU), the inference times for individual NNs are measured at $230\mu\text{s}$ (standard flux mode) and $710\mu\text{s}$ (control mode). This indicates these NNs could be adapted to real-time applications without issue. Since the flux mode is composed of many NNs (one for each response variable) this may require a parallelized implementation in order to keep computational times low.

3.4.2. General prediction performance

Performance of the flux response prediction is shown in fig. 15 where we have plotted the NN-predicted PCA coefficient vs the Gspert values, for the first 4 PCA modes. These predictions are for the flux response to PF5 the midplane outboard shaping, for the entire test dataset. The densest clusters lie close to the true values, though there is scatter for individual sample points. As expected, the prediction is most accurate on the first PCA mode ($R^2 = 0.755$) and decreases slightly with each increasing mode number.

Figure 12 and fig. 13 show individual flux responses to two different perturbation sources, l_i and the upper inboard coil PF1AU. The predicted vs actual flux responses

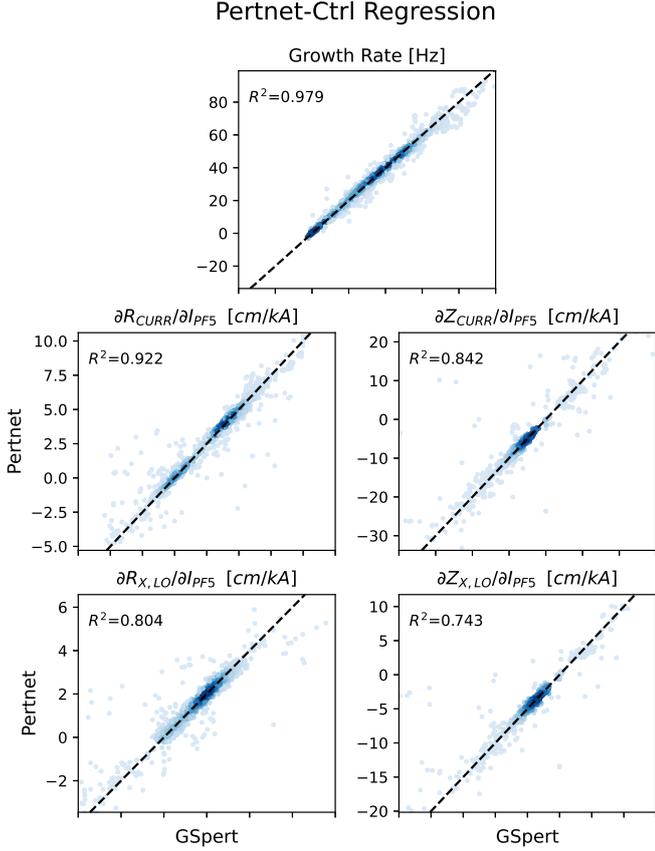


Figure 16: Regression of Pertnet parameter responses vs the Gspert values, based on predictions from Pertnet-control mode.

are difficult to distinguish by eye, although there are some slight differences. Predictions of the timetraces for parameter responses are plotted in fig. 14 for as single shot 204653. These include the current centroid response, x-point response, and growth rate as obtained by eqs. (11), (13) and (17). Even though the flux response is not trained to predict these parameters, the extracted parameter response is close to the Gspert values. However, if only interested in the parameter response, the control mode outperforms slightly. For example, in the second half of the shot, Gspert estimates that the radial position of the x-point ($R_{X,LO}$) will grow increasingly more sensitive to the PF1AU coil. The control-mode captures this trend, though the standard flux mode does not trend as strongly. The growth rate prediction for both models matches Gspert very accurately, which is a nice confirmation of prediction of all the NN models. Note that the growth rate prediction is unlike the other parameters, in that it does not depend on the response to just one input coil, but all the coils simultaneously (via eqs. (16) and (17), since the \mathbf{X} matrix includes effects from all coils). Even if individual flux responses can each have inaccuracies, this is an indicator that the dominant effects upon the coil current dynamics is being captured.

Lastly, fig. 16 shows the performance of Pertnet-control

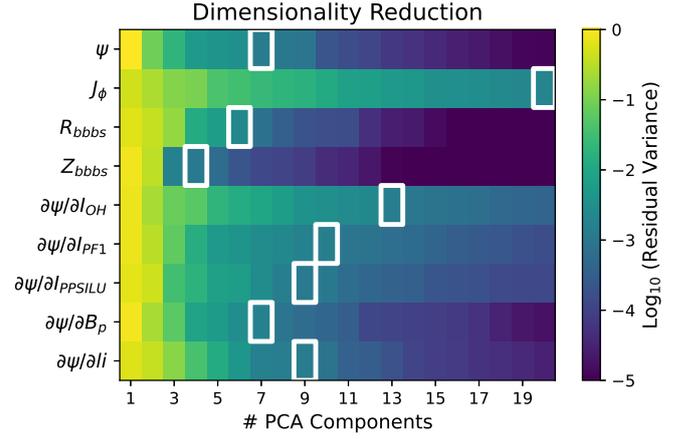


Figure 17: Dimensionality reduction of the Pertnet dataset using Principal Component Analysis (PCA). The outlined boxes indicate 99.5% captured variance and the number of components used for each variable.

mode for predicting individual parameter responses, in this case to PF5. We note that the individual responses generally have slightly higher correlation coefficients than prediction of the flux modes. While it is not a linear transformation between flux and parameters, this is consistent with the observation on fig. 14 that direct parameter prediction performed slightly better.

The most accurate prediction is the growth rate ($R^2 = 0.979$) while the the least accurate is the Z-derivative of the lower x-point ($R^2 = 0.743$). We record that 50th and 90th percentile errors for these respective parameters are 1.2/4.6Hz and 0.16/1.00[cm/kA]. This is a fairly large difference between the 50th and 90th percentiles. Again, this is a result of the largest responses and prediction errors being driven by a small number of samples as a result of the dynamic periods during rampup and loss-of-control (e.g. for the worst 5% of predictions, 88% occur within the first or last 200ms of the shot). This tends to manifest itself worse in the vertical direction than in the radial direction, consistent with the vertical instability effect. Note the scale difference for R-coordinate response vs the Z-coordinate responses, and that the R^2 coefficient for radial responses is higher indicating better predictions for this direction.

4. Conclusion

In this work we have introduced two neural nets for the NSTX-U tokamak, as part of a suite of prediction and simulation tools being developed for optimizing plasma scenarios. The neural nets include Eqnet, a reconstruction or forward free-boundary equilibrium solver, and Pertnet which calculates the nonrigid plasma response. As a simple application, we use Eqnet to design the coil currents for a desired target equilibria, indicating good performance and stability. Both neural nets are trained and tested with different combinations of input and output data. We

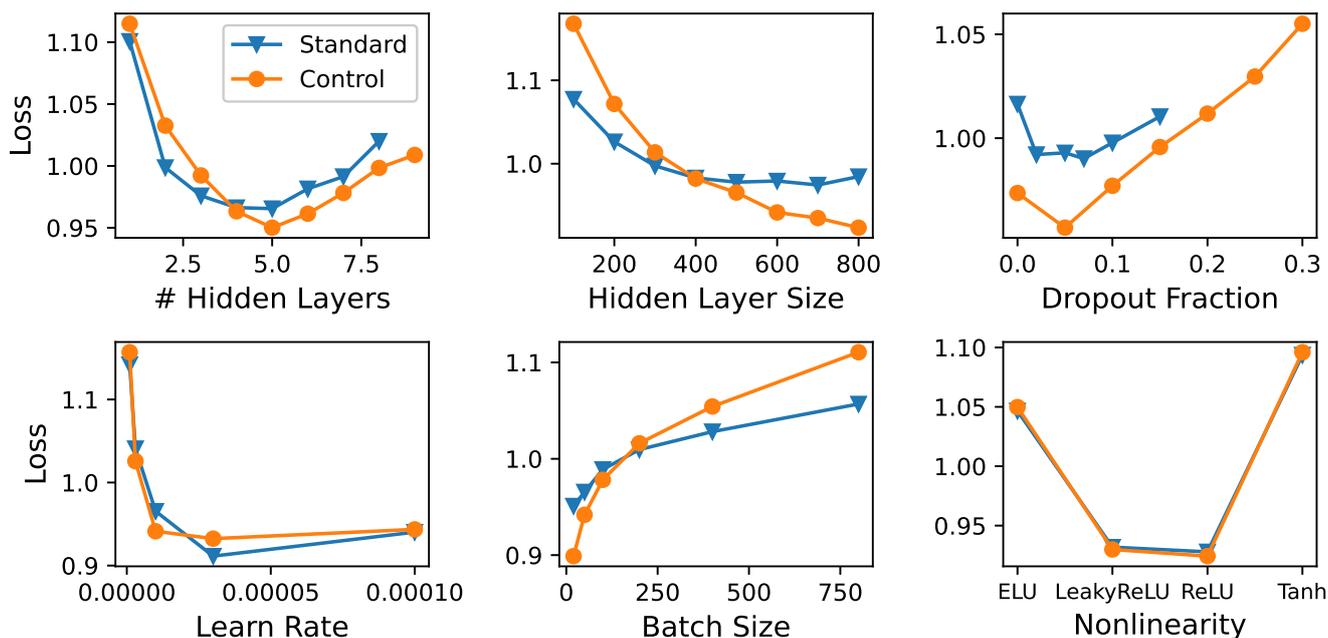


Figure 18: Hyperparameter grid scans used to identify the best performing Pertnet model. The optimal model depends on which Pertnet mode is under consideration, with the control mode favoring larger networks than the standard mode.

find that using Eqnet to reconstruct equilibria from magnetic diagnostics gives some performance improvements over RTEFIT especially when vessel eddy currents are significant. We report strong performance for Pertnet in predicting the plasma response. If only interested in the response of various parameters, it is simpler (fewer NNs to train) and marginally more accurate to predict these directly. However, predicting the full flux response is needed in some situations and can give good estimates of the parameters with additional computational steps.

Acknowledgements

This work was supported by the US Department of Energy Grant under contract numbers DE-AC02-09CH11466, DE-SC0015480, and DE-SC0015878.

References

- [1] A. Jalalvand, et al., Real-time and adaptive reservoir computing with application to profile prediction in fusion plasma, *IEEE Transactions on Neural Networks and Learning Systems* (2021) 1–12. doi:[10.1109/tnnls.2021.3085504](https://doi.org/10.1109/tnnls.2021.3085504).
- [2] J. Abbate, et al., Data-driven profile prediction for DIII-d, *Nuclear Fusion* 61 (2021) 046027. doi:[10.1088/1741-4326/abe08d](https://doi.org/10.1088/1741-4326/abe08d).
- [3] M. Boyer, J. Chadwick, Prediction of electron density and pressure profile shapes on NSTX-u using neural networks, *Nuclear Fusion* 61 (2021) 046024. doi:[10.1088/1741-4326/abe08b](https://doi.org/10.1088/1741-4326/abe08b).
- [4] H. LI, et al., Machine learning of turbulent transport in fusion plasmas with neural network, *Plasma Sci. Technol.* 23 (2021) 115102. doi:[10.1088/2058-6272/ac15ec](https://doi.org/10.1088/2058-6272/ac15ec).
- [5] P. Gaudio, et al., An alternative approach to the determination of scaling law expressions for the l–h transition in tokamaks utilizing classification tools instead of regression, *Plasma Phys.*

- Control. *Fusion* 56 (2014) 114002. doi:[10.1088/0741-3335/56/11/114002](https://doi.org/10.1088/0741-3335/56/11/114002).
- [6] A. Murari, et al., Machine learning for the identification of scaling laws and dynamical systems directly from data in fusion, *Nuclear Instruments and Methods in Physics Research* 623 (2010) 850–854. doi:[10.1016/j.nima.2010.02.080](https://doi.org/10.1016/j.nima.2010.02.080).
- [7] C. Rea, et al., A real-time machine learning-based disruption predictor in DIII-d, *Nucl. Fusion* 59 (2019) 096016. doi:[10.1088/1741-4326/ab28bf](https://doi.org/10.1088/1741-4326/ab28bf).
- [8] M. Lungaroni, et al., On the potential of ruled-based machine learning for disruption prediction on JET, *Fusion Engineering and Design* 130 (2018) 62–68. doi:[10.1016/j.fusengdes.2018.02.087](https://doi.org/10.1016/j.fusengdes.2018.02.087).
- [9] K. Montes, et al., Machine learning for disruption warnings on alcator c-mod, DIII-d, and EAST, *Nucl. Fusion* 59 (2019) 096015. doi:[10.1088/1741-4326/ab1df4](https://doi.org/10.1088/1741-4326/ab1df4).
- [10] Y. Fu, et al., Machine learning control for disruption and tearing mode avoidance, *Physics of Plasmas* 27 (2020) 022501. doi:[10.1063/1.5125581](https://doi.org/10.1063/1.5125581).
- [11] A. Pau, et al., A machine learning approach based on generative topographic mapping for disruption prevention and avoidance at JET, *Nucl. Fusion* 59 (2019) 106017. doi:[10.1088/1741-4326/ab2ea9](https://doi.org/10.1088/1741-4326/ab2ea9).
- [12] J. Kates-Harbeck, et al., Predicting disruptive instabilities in controlled fusion plasmas through deep learning, *Nature* 568 (2019) 526–531. doi:[10.1038/s41586-019-1116-4](https://doi.org/10.1038/s41586-019-1116-4).
- [13] A. Piccione, et al., Physics-guided machine learning approaches to predict the ideal stability properties of fusion plasmas, *Nucl. Fusion* 60 (2020) 046033. doi:[10.1088/1741-4326/ab7597](https://doi.org/10.1088/1741-4326/ab7597).
- [14] M. Boyer, et al., Real-time capable modeling of neutral beam injection on NSTX-u using neural networks, *Nuclear Fusion* 59 (2019) 056008. doi:[10.1088/1741-4326/ab0762](https://doi.org/10.1088/1741-4326/ab0762).
- [15] E. Coccia, et al., Identification of noncircular plasma equilibria using a neural network approach, *Nuclear Fusion* 34 (1994) 1349–1363. doi:[10.1088/0029-5515/34/10/i05](https://doi.org/10.1088/0029-5515/34/10/i05).
- [16] J. Lister, H. Schnurrenberger, Fast non-linear extraction of plasma equilibrium parameters using a neural network mapping, *Nuclear Fusion* 31 (1991) 1291–1300. doi:[10.1088/0029-5515/31/7/005](https://doi.org/10.1088/0029-5515/31/7/005).

- [17] R. Albanese, et al., Identification of plasma equilibria in ITER from magnetic measurements via functional parameterization and neural networks, *Fusion Technology* 30 (1996) 219–236. doi:[10.13182/fst96-a30752](https://doi.org/10.13182/fst96-a30752).
- [18] C. M. Bishop, et al., Real-time control of a tokamak plasma using neural networks, *Neural Computation* 7 (1995) 206–217. doi:[10.1162/neco.1995.7.1.206](https://doi.org/10.1162/neco.1995.7.1.206).
- [19] L. Lagin, et al., Application of neural networks for real-time calculations of plasma equilibrium parameters for PBX-M, Elsevier, 1993, pp. 1057–1061. doi:[10.1016/b978-0-444-89995-8.50205-8](https://doi.org/10.1016/b978-0-444-89995-8.50205-8).
- [20] Z.-J. Zhu, et al., Estimation of plasma equilibrium parameters via a neural network approach, *Chinese Physics B* 28 (2019) 125204. doi:[10.1088/1674-1056/ab55d1](https://doi.org/10.1088/1674-1056/ab55d1).
- [21] A. A. Prokhorov, et al., The plasma shape control system in the tokamak with the artificial neural network as a plasma equilibrium reconstruction algorithm, *IFAC-PapersOnLine* 53 (2020) 857–862. doi:[10.1016/j.ifacol.2020.12.843](https://doi.org/10.1016/j.ifacol.2020.12.843).
- [22] B. Wang, et al., Artificial neural networks for data analysis of magnetic measurements on east, *Journal of Fusion Energy* 35 (2015) 390–400. doi:[10.1007/s10894-015-0044-z](https://doi.org/10.1007/s10894-015-0044-z).
- [23] S. Joung, et al., Deep neural network grad-shafranov solver constrained with measured magnetic signals, *Nuclear Fusion* 60 (2019) 016034. doi:[10.1088/1741-4326/ab555f](https://doi.org/10.1088/1741-4326/ab555f).
- [24] B. P. van Milligen, et al., Neural network differential equation and plasma equilibrium solver, *Physical Review Letters* 75 (1995) 3594–3597. doi:[10.1103/physrevlett.75.3594](https://doi.org/10.1103/physrevlett.75.3594).
- [25] A. S. Welander, et al., Nonrigid, linear plasma response model based on perturbed equilibria for axisymmetric tokamak control design, *Fusion Science and Technology* 47 (2005) 763–767. doi:[10.13182/fst05-a778](https://doi.org/10.13182/fst05-a778).
- [26] A. Welander, et al., Closed-loop simulation with grad-shafranov equilibrium evolution for plasma control system development, *Fusion Engineering and Design* 146 (2019) 2361–2365. doi:[10.1016/j.fusengdes.2019.03.191](https://doi.org/10.1016/j.fusengdes.2019.03.191).
- [27] R. Albanese, F. Villone, The linearized CREATE-I plasma response model for the control of current, position and shape in tokamaks, *Nucl. Fusion* 38 (1998) 723–738. doi:[10.1088/0029-5515/38/5/307](https://doi.org/10.1088/0029-5515/38/5/307).
- [28] R. Albanese, et al., CREATE-NL: A robust control-oriented free boundary dynamic plasma equilibrium solver, *Fusion Engineering and Design* 96-97 (2015) 664–667. doi:[10.1016/j.fusengdes.2015.06.162](https://doi.org/10.1016/j.fusengdes.2015.06.162).
- [29] M. Boyer, et al., Plasma boundary shape control and real-time equilibrium reconstruction on NSTX-u, *Nuclear Fusion* 58 (2018) 036016. doi:[10.1088/1741-4326/aaa4d0](https://doi.org/10.1088/1741-4326/aaa4d0).
- [30] H. Grad, H. Rubin, Hydromagnetic equilibria and force-free fields, *Proceedings of the 2nd UN Conf. on the Peaceful Uses of Atomic Energy* 31 (1958) 190.
- [31] V. Shafranov, Plasma equilibrium in a magnetic field, *Reviews of Plasma Physics* (1966) 103.
- [32] L. Lao, et al., Reconstruction of current profile parameters and plasma shapes in tokamaks, *Nuclear Fusion* 25 (1985) 1611–1622. doi:[10.1088/0029-5515/25/11/007](https://doi.org/10.1088/0029-5515/25/11/007).
- [33] S. Sabbagh, et al., Equilibrium properties of spherical torus plasmas in NSTX, *Nuclear Fusion* 41 (2001) 1601–1611. doi:[10.1088/0029-5515/41/11/309](https://doi.org/10.1088/0029-5515/41/11/309).
- [34] D. A. Gates, et al., Vessel eddy current measurement for the national spherical torus experiment, *Review of Scientific Instruments* 75 (2004) 5090–5093. doi:[10.1063/1.1813091](https://doi.org/10.1063/1.1813091).
- [35] G. Ambrosino, R. Albanese, Magnetic control of plasma current, position, and shape in tokamaks: a survey or modeling and control approaches, *IEEE Control Systems Magazine* 25 (2005) 76–92. doi:[10.1109/MCS.2005.1512797](https://doi.org/10.1109/MCS.2005.1512797).