

Application of Machine Learning and Artificial Intelligence to Extend EFIT Equilibrium Reconstruction

L.L. Lao^{1*}, S. Kruger², C. Akcay¹, P. Balaprakash³, T.A. Bechtel^{1,4}, E. Howell², J. Koo³, J. Leddy², M. Leinshauser⁵, Y.Q. Liu¹, S. Madireddy³, J. McClenaghan¹, D. Orozco¹, A. Pankin⁶, D. Schissel¹, S. Smith¹, X. Sun^{1,4}, and S. Williams⁷

¹General Atomics, San Diego, California, USA

²TechX, Boulder, Colorado, USA

³Argon National Laboratory, Lemont, Illinois, USA

⁴Oak Ridge Associated University, Oak Ridge, Tennessee, USA

⁵University of Delaware, Newark, Delaware, USA

⁶Princeton Plasma Physics Laboratory, Princeton, New Jersey, USA

⁷Lawrence Berkeley National Laboratory, Berkeley, California, USA

*Contact author: Lao@Fusion.gat.com

Abstract — Recent progress in the application of machine learning (ML) / artificial intelligence (AI) algorithms to improve EFIT equilibrium reconstruction for fusion data analysis applications is presented. A device-independent portable core equilibrium solver capable of computing or reconstructing equilibrium for different tokamaks has been created to facilitate adaptation of ML/AI algorithms. A large EFIT database comprising of DIII-D magnetic, Motional-Stark Effect (MSE), and kinetic reconstruction data has been generated for developments of EFIT Model-Order-Reduction (MOR) surrogate models to reconstruct approximate equilibrium solutions. A neural-network (NN) MOR surrogate model has been successfully trained and tested using the magnetically reconstructed datasets with encouraging results. Other progress includes developments of a Gaussian-Process (GP) Bayesian framework that can adapt its many hyperparameters to improve processing of experimental input data and a 3D perturbed equilibrium database from toroidal full magnetohydrodynamic linear response modeling using the MARS-F code for developments of 3D-MOR surrogate models.

Keywords — Tokamak equilibrium reconstruction, machine learning, artificial intelligence, Gaussian Process, Model Order Reduction, neural network, 3D perturbed equilibrium

I. INTRODUCTION AND OVERVIEW

Reconstruction of experimental axisymmetric magnetohydrodynamic (MHD) equilibria is fundamental to tokamak research and operation and is an important part of fusion data analysis and plasma control. Equilibrium reconstruction provides essential magnetic geometry and current and pressure profiles information necessary to support tokamak operation and data analysis, and has contributed to several major discoveries of tokamak physics such as the experimental validation of theoretically predicted β stability limits [1] and the negative central-shear operating regime [2,3].

The equilibrium reconstruction and fitting code EFIT is widely used in many tokamaks to reconstruct experimental MHD equilibria. These include reconstructions at the DIII-D, C-MOD, JET, HIT, START, MAST, KSTAR, JT-60U, NSTX, TORE SUPRA, HT-7, HL-2A, EAST, and QUEST devices [4-26]. EFIT reconstructs equilibria by solving the Grad-Shafranov (GS) equation [27, 28] while approximately conserving the available experimental measurements and the imposed physics constraints

$$\Delta^* \psi(R, Z) = -\mu_0 R J_\varphi(R, \psi), \quad (1)$$

$$J_\varphi(R, \psi) = R P'(\psi) + \frac{\mu_0 F F'(\psi)}{4\pi^2 R}. \quad (2)$$

Here, ψ is the poloidal magnetic flux per radian of the toroidal angle φ enclosed by a magnetic surface and $\Delta^* = R^2 \nabla \cdot (\nabla / R^2)$. J_φ , P , and $F = 2\pi R B_\varphi / \mu_0$ are the toroidal current density, pressure, and poloidal current stream function. B_φ is the toroidal magnetic field at the major radius R .

This is an inverse problem. Information provided by plasma responses outside the plasma such as external magnetic measurements and inside the plasma such as spectroscopic and kinetic profile measurements are used to infer and reconstruct the current source information flowing within the plasma. EFIT efficiently searches for the optimized solution vector that best fits the available measurements and meets the imposed physics and free-boundary MHD equilibrium constraints,

by transforming the non-linear optimization problem into a sequence of linear optimization problems using a Picard iteration scheme [4-6].

Machine learning (ML) and artificial intelligence (AI) techniques have been applied to develop practical and efficient means to compute approximate solutions to various magnetic fusion research problems [29-39]. These include early work to solve the equilibrium GS Eq. (1) based on neural network [29] and more recent work to apply deep neural network to solve the equilibrium GS equation with measured external magnetic signals in the KSTAR tokamak based on the offline EFIT equilibrium reconstruction results [36].

In this paper, recent progress in the application of ML/AI algorithms to further test and extend EFIT equilibrium reconstruction capability for fusion data analysis applications are presented, taking advantage of the large database of DIII-D equilibria reconstructions available. A device-independent portable core equilibrium solver EFIT-AI, capable of computing or reconstructing equilibrium for different tokamaks, has been created to facilitate adaptation of ML/AI algorithms. A large EFIT database comprising of DIII-D magnetic, Motional-Stark Effect (MSE), and kinetic reconstruction data has been generated for developments of EFIT Model-Order-Reduction (MOR) surrogate models to reconstruct approximate or to accelerate the search of equilibrium solutions. A neural-network (NN) MOR surrogate model has been successfully trained and tested using magnetically reconstructed dataset with encouraging results. Pearson correlation coefficients R^2 of 0.98-0.99 have been obtained for the poloidal flux solution function ψ and global plasma parameters such as the normalized toroidal beta β_N and safety factor q_{95} at 95% of normalized ψ .

Other progress includes development of a Gaussian-Process (GP) Bayesian framework that can adapt its many hyperparameters to improve processing of experimental input data and construction of a 3D perturbed equilibrium database from toroidal full MHD linear response modeling using the MARS-F MHD code and development of a 3D-MOR surrogate model from the database using the singular-value decomposition (SVD) approach.

This paper is organized as follows. Creation of a device-independent portable core equilibrium solver EFIT-AI is discussed in Section II. Generation of a DIII-D EFIT equilibrium reconstruction database and development of a MOR representation are discussed in Section III. Initial EFIT-AI magnetic reconstruction results based on a NN MOR surrogate model are discussed in Section IV. Development of a Gaussian-Process (PR) Bayesian framework to improve processing of experimental input data is discussed in Section V. Generation of a 3D perturbed equilibrium database using MARS-F and development of a SVD-based surrogate model are presented in Section VI. Lastly, a summary is given in Section VII.

II. DEVICE-INDEPENDENT PORTABLE CORE EQUILIBRIUM SOLVER

In this Section, we summarize the development of a device-independent portable core equilibrium solver EFIT-AI that is capable of computing or reconstructing equilibrium for different tokamaks to facilitate adaptation of ML/AI algorithms.

EFIT has been used in three specific roles: 1. post-processing analysis of experimental discharges, 2. between-shot analysis of discharges for experimental control and planning [40] and 3. real-time reconstructions for the plasma control system [41]. These roles go from increasing accuracy to decreasing accuracy, and concomitantly, from increasing time-to-solution to decreasing time-to-solution. Roles 2 and 3 especially require tight integration into the data acquisition system for each experimental device; e.g., the data acquisition system for DIII-D is significantly different than that for JET.

Historically, this was handled by effectively forking the EFIT source code without the use of version control, resulting in code fragmentation. The increasing development of EFIT features combined with this code fragmentation made maintenance difficult. To overcome these issues and to facilitate adaptation of ML/AI algorithms into EFIT, we are generalizing the core functionality and leveraging modern software tools to modernize EFIT. As an example, we use Gitlab to act as host for using Git as our version control system, and for its excellent continuous integration (CI) features. Our goal is to create a single new version of EFIT named EFIT-AI that can be used for different tokamak devices and compiled and run under a variety of computer platforms over a wide range of compiler options including a high-level optimization, as well as capable to accommodate different spatial resolutions. This new EFIT-AI version creates a solid foundation for implementing ML/AI algorithms into EFIT.

In the past, when EFIT versions were forked to model a new device, the equilibrium solution methods and majority of the reconstruction code were left untouched. Important parameters related to the number of diagnostic channels, external poloidal-field coils, and other machine specificities were hard coded in order to standardize input variable dimensions. This has been

circumvented in the new EFIT-AI version by introducing a new namelist for the machine-dependent parameters that is read in before the normal inputs. This is similar to the approach used by EFIT++ [11]. Another important modernization has been the use of the CMake build system [42]. The use of a scriptable build system allows for more easily enabling or disabling code features based on detection of available software. For example, if the MDS+ library [43] for fetching experimental data is not found, then code that interfaces to the MDS+ system can be disabled easily. Another feature is out-of-place builds whereby the build files are kept separate from the source files and can be placed anywhere on a file system. This enables multiple builds against a single code version allowing for more rapid testing, especially comparison against different compilers and compile options. With these improvements, the new EFIT-AI version has been aggressively improved to allow a single version of EFIT source code that can be built for different tokamak devices with a wide range of compilation and optimization options.

As a result, the equilibrium solver has been made entirely machine independent. This is demonstrated in Fig. 1 by comparison of DIII-D and ITER equilibria computed using separate previous DIII-D and ITER versions of EFIT against results from the new single machine-independent EFIT-AI version. Equilibrium magnetic surfaces from the new EFIT-AI closely follow those from the two previous EFIT DIII-D and ITER versions.

A CTest-based regression test system (CTest being CMake’s testing system) has also been implemented. We are currently testing three different machines (DIII-D, NSTX and ITER) each time the test system is invoked; i.e., the new version of EFIT is machine independent. New machines can be conveniently added. Our CI system currently tests against two different compiler sets (GCC and Clang) and we plan to add other compiler sets to the CI system to reduce the hand-testing.

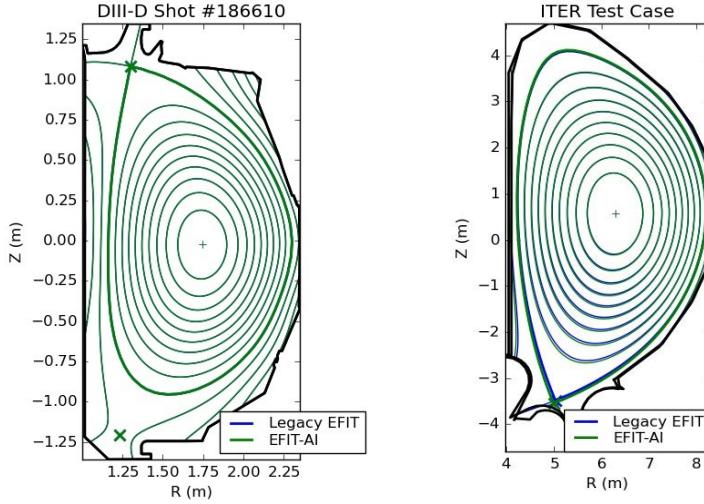


Fig.1. Comparison of DIII-D (left) and ITER (right) equilibrium magnetic surfaces (Blue) computed using two previous DIII-D and ITER versions of EFIT against results (Green) from the single new machine-independent EFIT-AI version. The green contours almost entirely mask the blue lines.

To further demonstrate the improvements to EFIT that these new developments have brought, we consider the performance improvements that have resulted from code improvements which have allowed for more aggressive compiler optimizations. These improvements include explicit variable definitions, well-defined code blocks, removal of constructs that are not universally supported, and use of newer compiler versions. We will compare the wall clock time required by a previous EFIT version used at DIII-D that has limited optimization with the new version being incorporated in EFIT-AI to quantify the improvements made.

For this comparison we have performed a high resolution, in both space and time, reconstruction over a full plasma discharge from a DIII-D negative triangularity shot #180533 that was part of the recent 2019 run campaign. The left side of Fig. 2 shows a comparison of the evolution of β_N , internal inductance ℓ_i , and q_{95} throughout the shot from the new EFIT-AI against those from the previous DIII-D version. The magnetic surfaces reconstructed at the time marked by the vertical red line using these two versions are compared on the right side. These plots have two sets of curves overlaid from the different versions of the code. The blue lines showing the

solution from the previous version of EFIT are completely covered by the green lines from the new EFIT-AI version.

We have also tested the EFIT-AI version of the code on 3 different computer platforms, with combinations of 4 different compilers that were native to the systems. The first is the Iris machine at General Atomics that is composed of Intel Haswell E5-2630 v3 CPUs. This system supports GNU, Intel, and PGI compilers and is where the previous EFIT version was compiled with PGI and limited optimization and sees regular use. While this system has many nodes, it restricts users to only 20 cores per job. The second machine is the Haswell partition of the NERSC machine Cori which has E5-2698 v3 CPUs. Here the GNU, Intel and Cray compilers are supported and the only limit to the number of cores that can be used for a job is the size of the system (76,416 total). The final machine considered is a Dell Precision 3551 laptop with i7-10850H CPU containing 6 cores. GNU is the only compiler tested here.

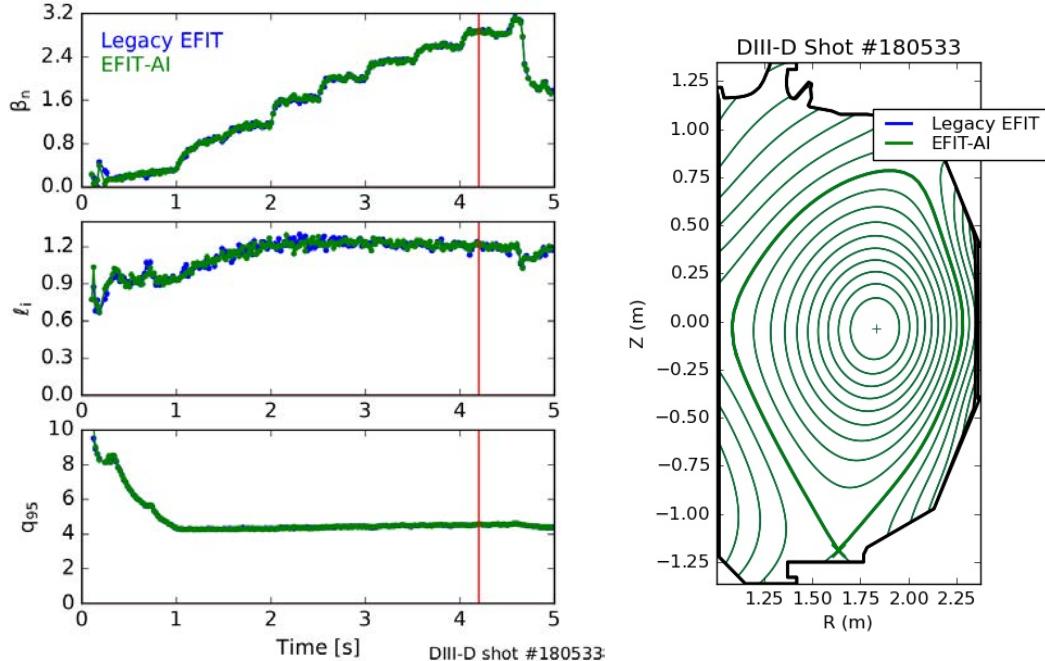


Fig. 2. Left: Time traces of plasma parameters computed by EFIT are shown for DIII-D shot #180533. At 4.2s (marked by the red line) the flux surfaces describing the plasma are shown in the plot on the right. In both of these plots results computed by the previous legacy version of EFIT (Blue) are covered by the results from the new EFIT-AI version (Green).

In this case study, 240 roughly equally spaced time slices from the length of the discharge were reconstructed on a 129×129 spatial grid with a relative convergence tolerance of 10^{-4} or maximum of 100 iterations. The separate times were solved simultaneously using Message Passing Interface (MPI) parallelism across all available cores on each machine (up to 240). This was repeated 3 times for each instance of the code and the runtimes were averaged together to reduce the effects of shot noise. The run time results shown in Fig. 3 demonstrate the significant speed up EFIT-AI offers relative to the previous version. The improved performance is responsible for a majority of the difference in run times between the first column and the others. Just comparing timing from the Iris machine, EFIT-AI is roughly 3 \times faster than the previous version.

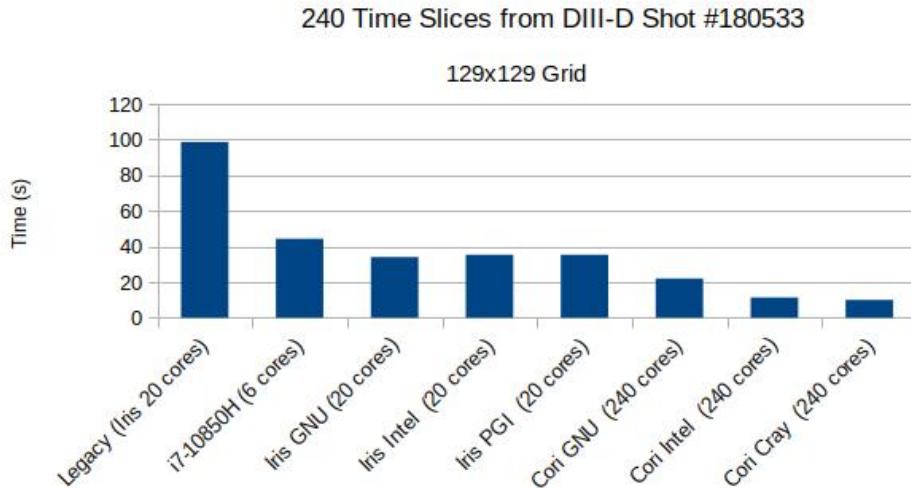


Fig. 3. The run times for 8 different EFIT builds leveraging the maximum available MPI parallelism, are shown for reconstructing 240 time slices on a 129×129 spatial grid. The bar on the left is the previous legacy version of the code with limited optimization while all others are the EFIT-AI version. The EFIT-AI builds encompass optimized installations on 3 different computing systems using 4 different compilers.

The remaining variation seen in Fig. 3 can be attributed to the different CPU clock speeds and the number of cores run in parallel. The i7 laptop has the highest CPU speed at 5.1 GHz, while Cori is at 3.6 GHz, and Iris can run at 3.2 GHz. This is why the time to run on the i7 is similar to the times on Iris, despite having less than 1/3 as many cores available. The significant increase in the number of cores available on Cori allows for an even larger speedup. When the EFIT-AI code

is run on Cori and exploits the maximum available parallelism for this case, we see greater than 7 \times speed up over the previous version run on Iris. The times are all within the shot variation for every compiler available on a chosen system with the exception of GNU on Cori. This build requires 10s longer which we believe to be overhead, since it is also present on higher resolution cases with longer run times, accounting for a smaller fraction of the time taken. This could be a consequence of the GNU version being older on Cori (8.3) than the other systems (9.3).

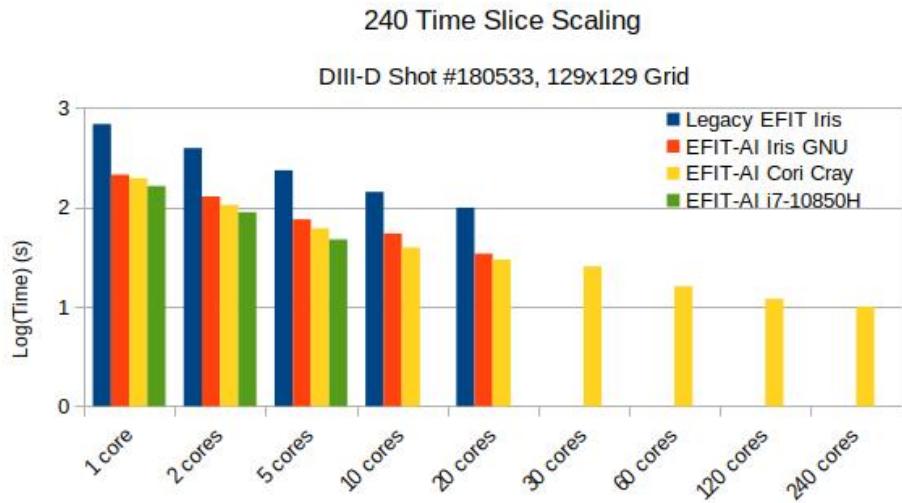


Fig. 4. The run times for 4 different EFIT builds on 3 different compute systems are compared when reconstructing 240 time slices on a 129×129 grid. For each case, the number of cores being used in parallel with the MPI interface is varied up to the available limits of the system. Run times are plotted on a log scale.

While the speed up with increased parallelism is significant, it does not match a strong scaling with the number of cores. Figure 4 shows the results of varying the number of cores used for one build on each system. The run time scales weakly with the number of processors used at similar rates on all systems. With ideal scaling one would expect the EFIT-AI code to take an average of 10s with 20 cores and 1s with 240 cores which is 3 \times faster than observed at 20 cores and 10 \times faster at 240 cores. Part of the reason that the ideal scaling is not observed, is that some time slices require longer to solve than others, with the longest requiring 6s. Therefore, the observed scaling

is actually within 4s ($<2\times$) of the strong scaling at 240 cores and much closer with fewer cores. This discrepancy is likely caused by overhead with the parallel decomposition. Similar performance improvement and scaling of the EFIT-AI code has been observed for higher resolution (257 \times 257) runs and when MSE and CER constraints are included as well.

Lastly, the capability to read HDF5 files with the OMAS equilibrium data format [44] under the OMFIT integrated modeling framework [45] has also been added to EFIT-AI. This allows users easily adjust and re-run cases that were stored as part of the new databases described in the next section. Furthermore, it allows for an easier integration with AI-generated initial conditions.

III. DATABASE AND MOR REPRESENTATION

One of the important applications of machine learning (ML) in the physical models is to serve as a MOR model for the physical model, also known as a surrogate model. The development of a ML-based MOR model for the GS Eq. (1) has been explored previously [29,36,37], but our goal is to extend the prior work in multiple ways. The first is to take advantage of the large database of DIII-D EFIT equilibria reconstructions available to further test and extend the application of ML/AI algorithms to enhance EFIT equilibrium reconstruction capability for fusion data analysis applications. The second is to explore the accuracy of the neural network by exploiting some of the same physical insights that makes EFIT successful.

Machine learning has enjoyed a large resurgence in recent years as a result of the increase in data and resources required to train the neural network models. Simultaneous to the improvements in the machine learning, there has been improvements in the development of tools to generate the database, notably the OMFIT framework [45] and the OMAS/IMAS data standards [44]. Thus, EFIT-AI is providing a bridge between the developments of the fusion community and the ML community to produce an advanced MOR representation of EFIT.

A. EFIT-AI Database

A large database of EFIT equilibria have been generated and gathered using various levels of reconstruction fidelity: including magnetics, Motional Stark Effect (MSE) spectroscopy, and kinetic EFITs. All EFIT reconstructions contain data from external magnetic measurements including flux loops and magnetic probes, plasma Rogowski loops for the plasma current, poloidal-field and Ohmic coils, vacuum toroidal magnetic field strength, and limiter and vacuum vessel positions. At the most basic level of reconstruction is the magnetics EFIT equilibrium database that primarily contains constraints from external magnetic measurements. This database contains EFIT equilibrium reconstructions for the entire DIII-D 2019 experimental campaign, and consists of 438,986 equilibria over 2331 discharges. Plotted in panel (a) of Fig. 5 are the reconstructed β_N versus ℓ_i in this database. Next is the MSE database, which has the additional constraint of the

magnetic field line pitch angle in the plasma core as measured by the MSE diagnostic, and consists of 299,991 equilibria over 1652 discharges. The reconstructed β_N versus ℓ_i are shown in panel (b) of Fig. 5. Finally, the database of kinetic EFITs has been gathered from existing OMFIT kinetic EFITs generated by users on DIII-D. This database has additional constraints of the pressure and current profiles from fitting of plasma temperatures and densities and edge bootstrap current constraints from physics models. As this database was manually generated, it is much smaller than the magnetic and MSE databases with 8963 equilibria over 222 discharges.

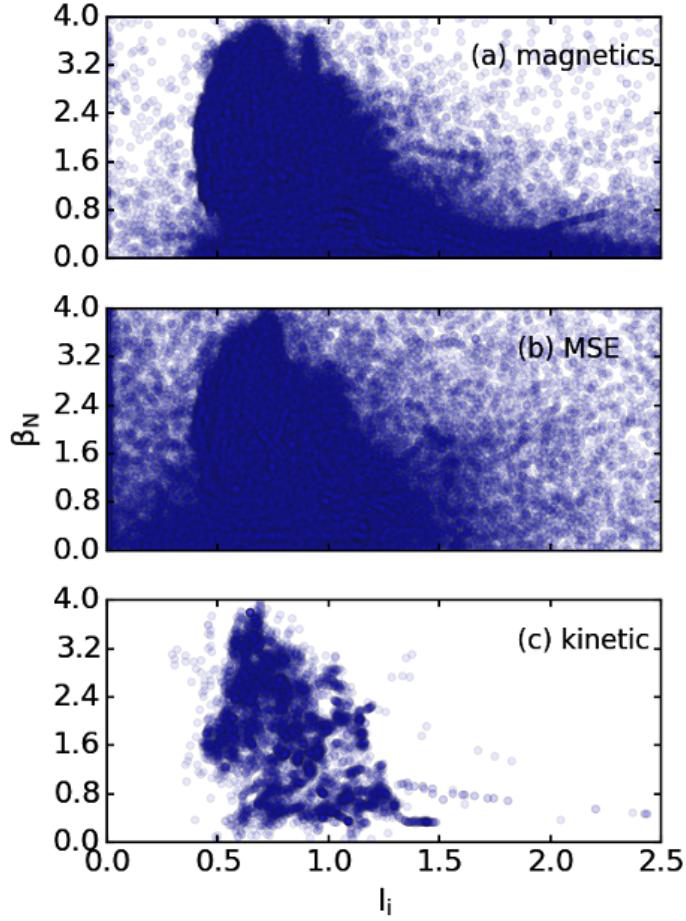


Fig. 5. β_N is shown versus ℓ_i for the magnetics (a), MSE (b), and kinetic (c) EFIT databases.

B. EFIT-MOR

EFIT-MOR offers the promise of moving to dynamic and real-time full kinetic equilibrium reconstructions currently limited to off-line analysis. A form of it has already been applied to the

magnetic reconstructions of the KSTAR tokamak [36]. In that work a fully-connected (or multi-layer perceptron, MLP) NN is trained on the magnetic measurements used in EFIT to reconstruct the poloidal flux functions ψ . The NN training is further constrained by including the equilibrium constraint in the loss function. Another more recent example of equilibrium reconstruction with NNs involves *Eqnet*, a NN trained to predict the free boundary-plasma equilibria in NSTX-U with several modes of operation based on various set of inputs [39]. Our present MOR approach parallels the KSTAR and NSTX-U efforts in some respects and differs from these cited works in other significant ways, which are highlighted below. In particular, to develop EFIT-MORs, two MLP NN implementations have been trained and tested. The first implementation reconstructs $\psi(R, Z)$ on the EFIT rectangular (R, Z) grid given the external magnetic signals and the equilibrium GS constraint. This is similar to both the KSTAR approach and reconstruction mode of *Eqnet* for NSTX-U [39]. The second implementation uses a separate MLP-NN to directly fit the plasma boundary and various global plasma parameters of interest such as β_N , ℓ_i , and q_{95} . This is motivated by the fact that external magnetic measurements alone only yield information on plasma boundary, external magnetic geometry, and global plasma parameters such as poloidal β_p (or β_N) and ℓ_i [46], and that the plasma boundary and these global plasma parameters are essential for plasma transport and stability calculations but are computationally expensive to compute. The estimation of these parameters with a NN, given the magnetic inputs, resembles Ref. [39]’s *Eqnet* implementation under the *reconstruction control mode*.

The input vector for the training of the NNs consists primarily of the external magnetic signals from approximately 123,000 time slices from the 2019 database as described in the previous Section IIIA. This includes tangential component of the magnetic field measured by 76 magnetic probes, the poloidal magnetic flux measured by 44 flux loops, the toroidal plasma current I_p , and 18 measurements of the poloidal-field coil currents; resulting in 139 total *features* for each time slice. The location of the magnetic probes and flux loops are shown as red lines and blue squares in Fig. 6. The F-coils appear as rectangular boxes labeled ‘external coils’ in the same figure. All 139 input space features are normalized by various combinations of the background time-varying

vacuum magnetic field B_0 , the major radius $R_0 = 1.67$ m, and average minor radius $r_0 = 0.6$ m. For example, the magnetic-probe measurements are scaled by $(r_0/R_0)B_0$, which is approximately the poloidal equivalent of the vacuum (toroidal) magnetic field. These normalizations place most of the magnetic measurements on the $[-1, 1]$ range, providing approximately equal initial weights for all input features. The feature weighting is then determined at the end, intrinsically by the learning algorithm. However, a further systematic study that varies the initial weighting could be performed to investigate any bias in the NN architecture. The temporal trace of a few of these dimensionless magnetic signals are plotted in the right panel of Fig. 6 for DIII-D shot #180087. Note that unlike in the KSTAR NN training, the EFIT grid is not included in the input vector space.

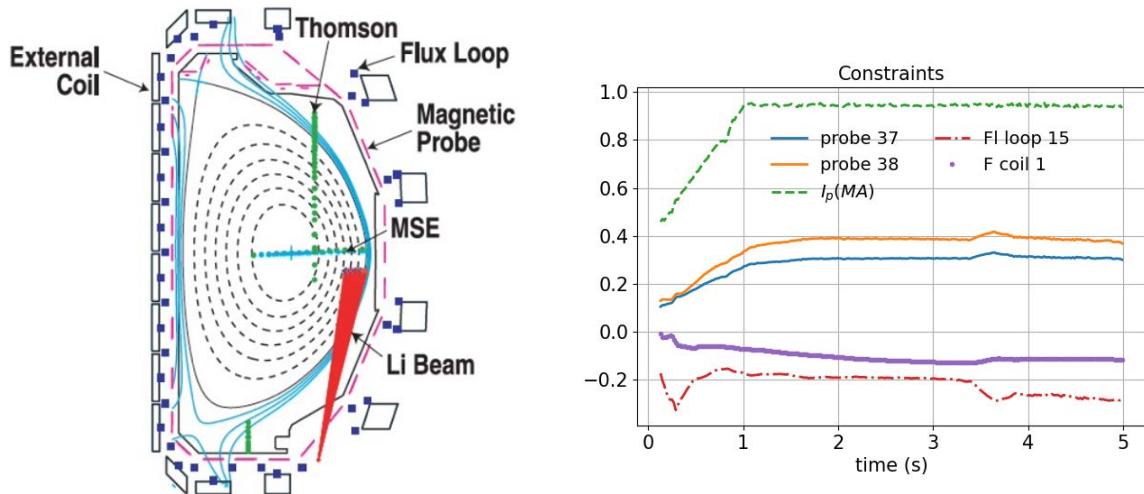


Fig. 6. Left: The DIII-D diagnostics that are used in the EFIT equilibrium reconstructions. The magnetic diagnostics used in this work to teach the neural networks are the magnetic probes, flux loops, and currents in the external coils. A temporal trace of the dimensionless form of several of these signals from DIII-D discharge 180087 are plotted in the right panel.

In the first implementation, the targets (or outputs) of the NN are the poloidal flux ψ on the 129×129 rectangular (R, Z) grid of EFIT. To speed up the training, the output ‘image’ resolution space has been halved, yielding 65×65 pixels for the equilibrium flux surfaces. Thus, the total number of output features for each time slice is $65^2 = 4225$. The poloidal flux ψ is normalized to

lie over the range [0,1] within the plasma, and can assume any value over the approximate range [-10,10] in the vacuum region. The quality of the NN reconstructions is evaluated with the Pearson correlation coefficient R^2 and the peak-signal-to-noise ratio, which is an-image relevant metric defined as follows: $\text{PSNR} = 10\text{Log}_{10}(\text{MAXi}^2) / E_S$, where MAXi is the maximum possible pixel value of the image (ψ) and E_S is the mean-square error.

A fully-connected MLP is used to learn ψ . The present network architecture consists of two hidden layers containing 437 and 1360 nodes, respectively. The network expands to accommodate the much larger size of the output vector compared to the input vector. The number of nodes at each layer is chosen to be the geometric mean of the number of nodes in the neighboring layers, although the NN accuracy is fairly insensitive to the number of nodes in the layers so long as a sufficient number of them is employed. A hyperbolic tangent *Tanh* activation function is used to propagate the information from the input space to the first hidden layer, and from the first hidden layer to the second hidden layer. The activations at the 2nd hidden layer are mapped linearly to the output layer. A linear mapping is necessary to accommodate the extended range of values that ψ assumes in the vacuum region. 80% of the 123,000 samples are used in the training of the NN, while the remaining 16% and 4% are reserved for validation and testing; yielding approximately 5000 time slices for the test set. The loss function of the NN is defined as a mean square error: $L = 1/N_t (|\psi - \psi_{\text{NN}}|^2)$, where N_t is the total number of training or validation samples. Adam (adaptive moment estimation) stochastic gradient descent method [47] is used in the back-propagation step to find the parameters (weights) of the NN that minimize the loss function.

The results are shown below in Figs. 7(a), (c), (e), and (g), which overlay the NN prediction of ψ in dashed red contours with the true ψ solid black contours from the database reconstructed with EFIT. Four randomly selected samples out of the 5000 test slices are shown. Note that none of the four display samples are necessarily from the same discharge. Fig. 7(b), (d), (f), and (h) show the normalized residual $[\psi(R,Z) - \psi_{\text{NN}}(R,Z)] / \max(|\psi(R,Z)|)$ for the same four samples. The residuals for the time slices from the flat-top stage of the plasma are all less than 1% for all

locations on the (R,Z) grid. The residual for the time slice from the ramp up, 7(b), gets as large as a few percent. This is likely a consequence of the high variability of the plasma behavior during ramp up, which might require increasing training sample size and/or employing deeper NNs such as Convolution NN to capture.

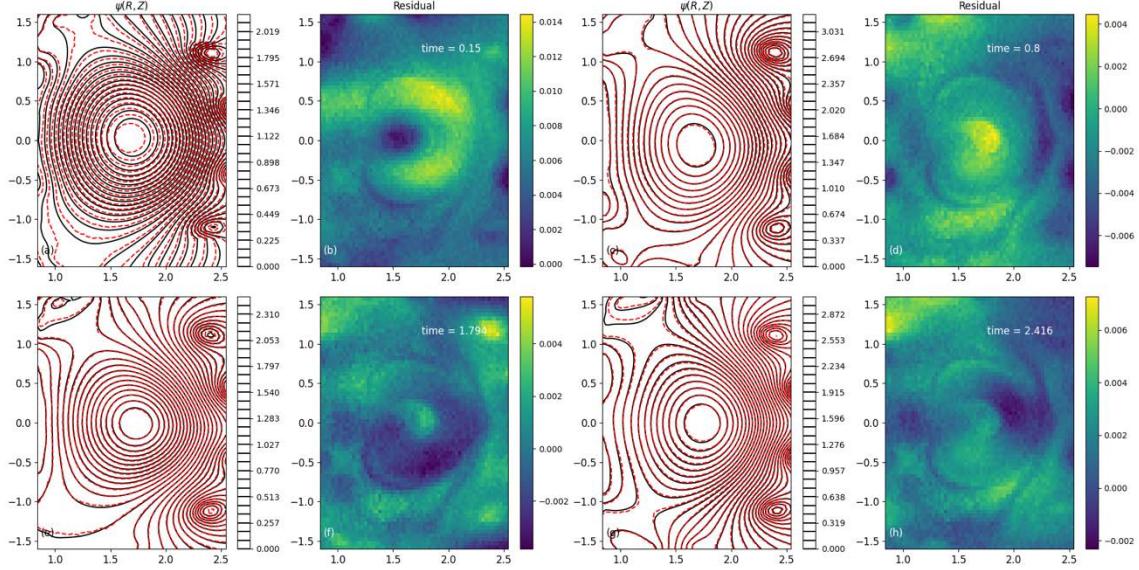


Fig. 7. Panels (a), (c), (e), (g) overlay the poloidal flux surfaces (solid black) from the magnetic database against those (dashed red) predicted using a MLP NN for four randomly-selected times slices out of the 5000 test samples to test the NN reconstruction accuracy. Panels (b), (d), (f), (h) show the normalized residua $[\psi(R,Z) - \psi_{NN}(R,Z)] / \max(|\psi(R,Z)|)$ for the same four samples.

Figure 8(a) shows the quality of the NN reconstructions as a regression plot that compares the true value of ψ at every (R,Z) grid point to the NN-reconstructed ψ at the same grid points for the test set. The MLP-NN yields an $R^2 = 0.998$ and PSNR = 87 over the test set, signaling a high accuracy for the NN reconstructions. Figure 8(b) shows the evolution of the network's training and validation loss function as a function of the network iterations.

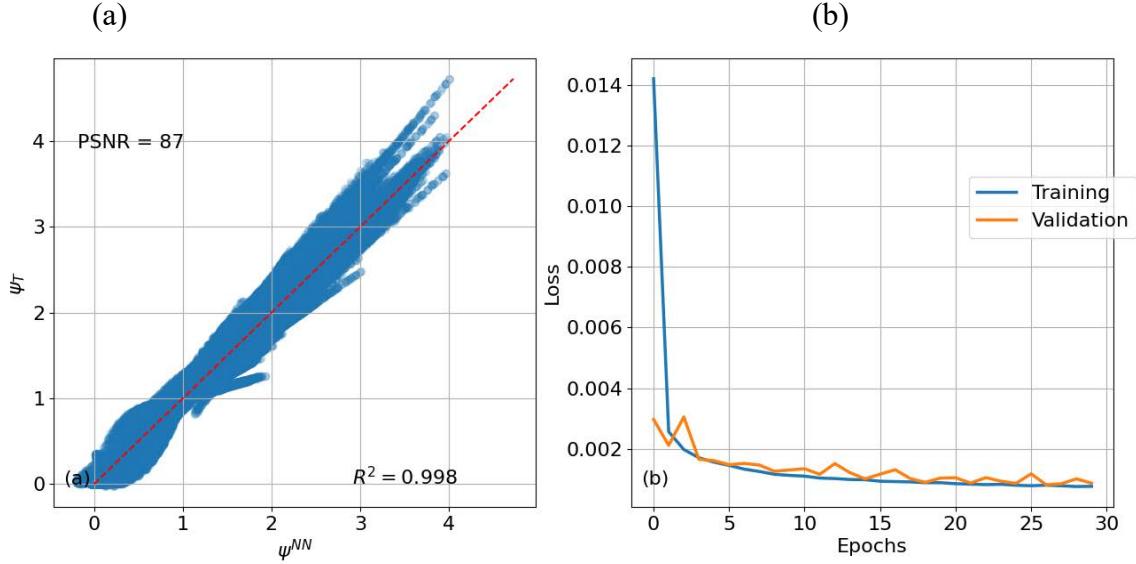


Fig. 8. Training results of MLP NN used to learn $\psi(R,Z)$. (a) True values of all ψ contained within the test data set plotted against ψ predicted by the MLP-NN. (b) Training and validation history showing the loss function versus the NN iterations.

Secondly, a separate neural network is applied to learn β_N , ℓ_i , and q_{95} using the same magnetic inputs as those used in learning ψ . The model-order-reduction with NNs provides a speedy way to calculate these global parameters, which are essential (alongside the plasma boundary) for plasma transport and stability calculations, but are computationally expensive to compute within EFIT. The MLP-NN employed in this task comprises 3 hidden layers, consisting of 53, 20, and 7 nodes. In this case, the architecture of the NN narrows since the output space only contains 3 features for each time slice. A *sigmoid* is used for all four activation functions to propagate the information from the inputs to the outputs. Other activation functions, applied in various combinations, yield similar accuracies. Once again, an E_S loss function is employed with the Adam optimizer to minimize the loss function of the NN. To remove the outliers, a filter is applied where any time

slice with $\ell_i > 2.0$ is discarded. The same 80-16-4 data split into training-validation-test sets is used.

The regression plots for the NN predictions of β_N , ℓ_i , and q_{95} are shown in Fig. 9, which also displays the R^2 of each fit. The results show accurate fitting of β_N and q_{95} by the NN, with $R^2 = 0.98-0.99$, while the prediction of ℓ_i underperforms with a lower $R^2 \sim 0.92$. These fits will likely improve once the NNs are trained over the entire set data.

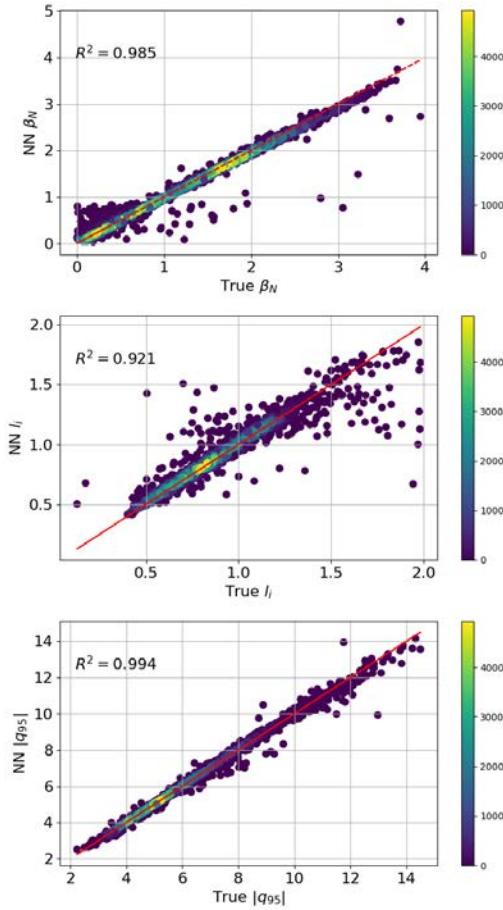


Fig. 9. Regression plots for the three global parameters β_N , ℓ_i , and q_{95} . An MLP NN with three hidden layers is employed to learn these parameters. 80% of the 123,000 equilibria are used in the training of the MLP-NN, while the remaining 16% and 4% are reserved for validation and testing.

In addition, another MLP-NN is implemented in this work to learn the (R, Z) coordinates of the plasma boundary. Each plasma boundary extracted from the magnetic EFIT database is first parametrized in terms of a local polar angle, which is then re-interpolated to lie on a surface comprising 101 points. Thus, the dimension of the output vector changes to 202 in this case, which also changes the network architecture to have 153, 168, and 184 nodes in its three hidden layers. For the training, we map the (R,Z) coordinates to approximately span the range [-1, 1]. The activation functions chosen for the present architecture are *leaky relu* (a modified version of the rectified linear unit function *relu* that admits negative values), *leaky_relu*, *tanh*, and finally *tanh* again to map the activations at the last (3rd) hidden layer to the output layer, resulting in an accuracy of $R^2 = 0.997$ for the plasma boundary reconstruction. This is illustrated in Fig. 10. Other combinations of activation functions could yield similar accuracies. Note that the choice to rescale (R, Z) of the plasma boundary to span [-1, 1], and not [0, 1], naturally necessitates the use of the *tanh* function instead of the *sigmoid* or *relu* to propagate the information from the last hidden layer to the output layer. These techniques can be readily extended to include other global plasma parameters of interest such as those magnetic geometry parameters related to the plasma boundary surface: elongation, triangularity, squareness, boundary gaps, X-point and divertor strike points, as was comprehensively done in Reference 39 for the NSTX-U tokamak.

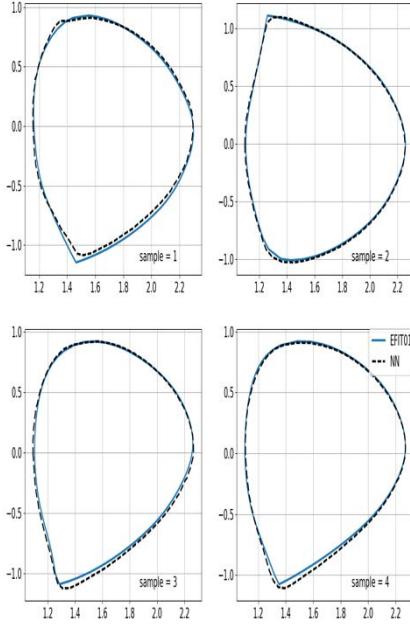


Fig. 10. NN reconstructions of the plasma boundary (dotted-dashed blue) overlaid against the true plasma boundary (solid black) for 4 randomly drawn samples out of the $\sim 5,000$ samples of the test set. An MLP-NN with three hidden layers is employed to learn the radial and axial (R, Z) coordinates of plasma boundary. 80% of the 123,000 EFIT magnetic equilibria are used in the training, while the remaining 16% and 4% are reserved for validation and testing.

Lastly, the learning of the flux surfaces can be further constrained by including in the NN loss function the toroidal current density $J_{\text{tor}}(R, Z)$ in terms of the two stream functions P and F as shown on the RHS of the GS Eq. (1). This provides a way to impose the important equilibrium force balance condition on the NN. This reconstruction exercise focuses only on the volume within the plasma boundary, which is currently taken as known from the EFIT equilibria and applied as a filter in the pre-processing stage. However, as demonstrated above, the NNs are capable of learning the plasma boundary accurately given the magnetic inputs. In our case, training the NN directly on J_{tor} , without the application of the Δ^* operator, produces noise-free J_{tor} and eliminates the requirement of including the grid in the NN training, as was done in the KSTAR work. This enhanced NN still uses the MLP-NN structure to predict $\psi(R, Z)$, then predicts $J_{\text{tor}}(R, Z)$ given the NN-produced $\psi(R, Z)$, using a convolutional neural network (CNN) [47]. CNNs are locally connected networks which are well-known for extracting spatial representations. They can

simultaneously extract key features from images and also reduce the number of trainable parameters combined with pooling operations, which is suitable for reconstructing J_{tor} given ψ .

IV. NEURAL-NET DIII-D MAGNETIC RECONSTRUCTION RESULTS

To test the accuracy of the MLP-NN in predicting the plasma boundary and global plasma parameters, the predictions of the NN MOR model described in Section IIIB are compared against those from the magnetic EFIT database EFIT. We have again chosen to use the DIII-D negative triangularity shot #180533 shown in Fig. 2 for this analysis with the same parameters as described in section II. For a more applicable comparison, we have repeated the solutions using only a single processor. With this, the previous version of EFIT requires on average of 682 seconds and the optimized EFIT-AI version requires 212 seconds.

This shot was excluded from the set of shots that the MOR model was trained on, but similar shots with negative triangularity from the campaign were included in the training. These negative triangularity shots only make a small fraction of approximately 123,000 time slices used in the training and validation of the NNs, as described in Section IIIB. Here, we leverage two trained NNs to predict global plasma parameters of interest as well as the full flux-surface topology for 240 time slices during the progression of shot 180533. The NN-based MOR method is able to produce these results in 0.345 seconds, which is $\sim 600\times$ faster than the full EFIT reconstructions. These are produced with a Python code that was intended to test out the method and has not been optimized for performance. Therefore, these numbers represent a lower bound for the possible efficiency of these methods.

The normalized β_N , internal inductance ℓ_i , and edge safety factor q_{95} are shown as a function of discharge time in Fig. 11. The minor differences between the true and inferred values exhibit little sensitivity to both network architecture and the choice of the activation functions. They also vary among discharges. The flux surfaces predicted by the surrogate model are compared with the direct EFIT solutions for four representative times in Fig. 12. The chosen times are shown as dotted vertical lines in Fig. 11. For the particular example shown, the NN accurately predicts the flux surfaces inside the plasma, while it does not always reproduce the magnetic topology near the

field nulls, as evidenced by Fig. 12 (a) and (c). These deviations are likely to get smaller as the training set is expanded to the full 2019 EFIT database.

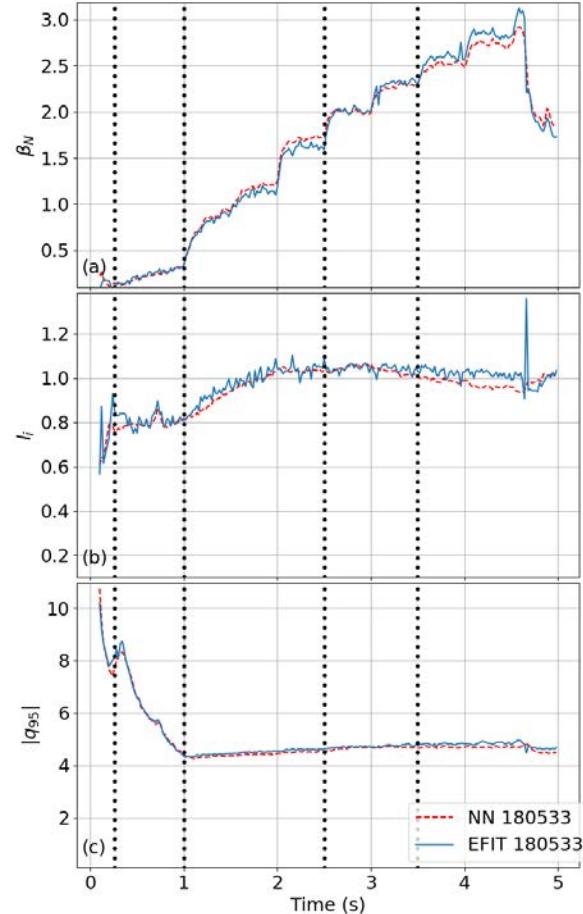


Fig. 11. The NN inference of the global plasma parameters (dashed red) against direct EFIT reconstructions of the same quantities (solid blue) as a function of time for a negative triangularity DIII-D discharge (180533). Shown are (a) the normalized plasma beta β_N , (b) the internal plasma inductance l_i , and (c) the edge safety factor q_{95} as a function of the discharge time.

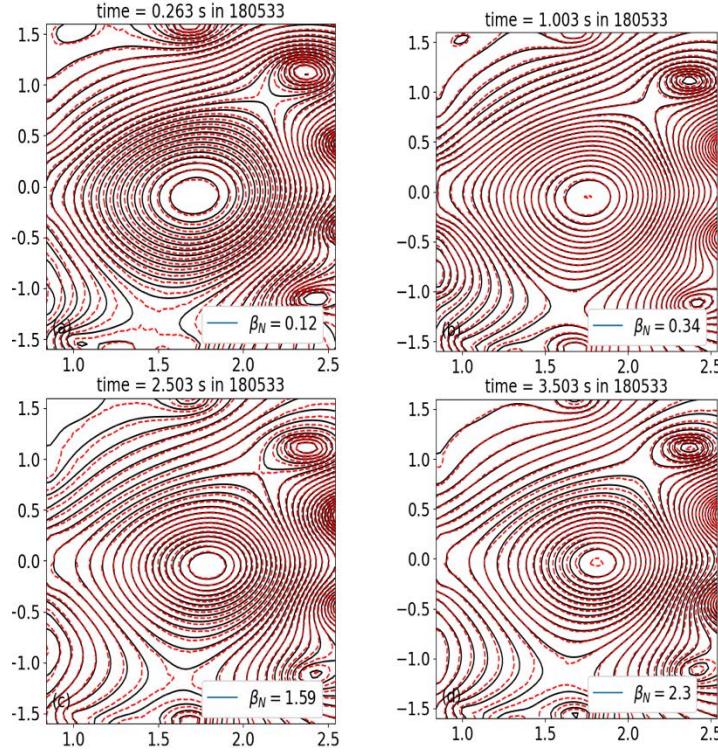


Fig. 12. The NN prediction of the negative-triangularity flux surfaces (dashed red) for DIII-D discharge 180533 overlaid against the flux surfaces reconstructed by EFIT (solid black) for the same discharge. Shown are flux contours from four different times during the discharge that correspond to 263 (early ramp-up), 1003 (end of ramp up), 2503, 3503 (flat top) ms during the discharge. These instances are marked as the dotted vertical lines in Fig. 11.

V. GAUSSIAN PROCESS TO IMPROVE PROCESSING OF EXPERIMENTAL DATA

Equilibrium reconstruction utilizes input from diagnostics to refine and inform about the experimental conditions of the plasma. These input data are in a variety of forms, depending on the diagnostics from which they are obtained. Two key profiles obtained through the Thomson scattering (TS) diagnostic are the electron temperature and density profiles, from which a pressure profile is constructed. These data can be noisy and can contain outliers, so a reliable and insightful fitting technique is required to make full use of the data. Currently, this can be done by identifying the regime of the plasma (H-mode, L-mode, etc.) and then using parameterized functions to do least-squares fitting of the data. This has several drawbacks. Firstly, outliers negatively affect these sorts of fitting techniques by pulling the fit away from the true profile. Secondly, an extra step is required in first identifying the regime, and the plasma may be in some alternative state that does not fit nicely into H-mode or L-mode, like containing an ITB. Lastly, these techniques do not intrinsically involve any sort of error propagation to help estimate the error on the fit based on the error on the data.

In this Section, we summarize our development of Gaussian Process Regression (GPR) method that solves these problems in a robust and reliable way. GPR has gained wide use since its introduction by Svensson for soft X-ray tomography in 2011 [48], It has since grown to include other diagnostics [49-53] and has broadened to be used for verification and validation [54,55] and in equilibrium reconstruction as well [56-58]. The latter is our ultimate goal, and here we are discussing the issues related to L-mode and H-mode parameter regimes.

GPR is a fitting technique that establishes a probability distribution over all possible fits to a set of data points. With such a distribution, it is simple to take the mean and variance to account for the best fit and the error on that fit. The underlying assumption of GPR is that this distribution is a multivariate normal distribution, so drawing samples from this distribution results in fits for the data. The trick is then to construct the distribution and optimize these hyperparameters so the fits that are obtained are high quality. This is done by choosing an appropriate kernel function and

then by either optimization of the hyperparameters that parameterize this kernel or through Markov Chain Monte-Carlo iteration to explore the hyperparameter space. Either way, the fits themselves are completely non-parametric, and so minimal knowledge of the data set is required a priori.

For the TS data, we have chosen to use a linear combination of two Matern kernels, each with their own length-scale hyperparameter that is constant in space. However, the kernels switch on and off at hyper-parameterized locations along the radial axis. Additionally, to handle outliers we have implemented a student-T likelihood distribution that is used during the hyperparameter optimization. This allows the mean fit to remain accurate as the student-T distribution has the potential for heavy tails that can automatically account for the large error associated with outliers. With GPR, it is not required to choose the hyperparameters directly, but instead to allow the data to inform and optimize the values of the hyperparameters. Additionally, the variance obtained from the fit is directly related to the input error on the data points, automatically including error propagation. We have chosen to utilize GPFlow library to implement the GPR for the TS data as because it is performant (based on TensorFlow) and contains a variety of methods that are well maintained and validated.

The results of the GPR methods we implemented are shown in Fig. 13 for both L-mode (Left) and H-mode (Right). The dotted line is the underlying true profile while the dots are synthetic data that includes gaussian noise and a number of outliers. The red line indicates the mean fit and the red shaded area is the 95% error bounds. The Student-T likelihood function accounts for these outliers very well, as seen in Fig. 13, as the mean is not moved away from the underlying true profile. The variance on these outlier data points is not meaningful, but that is acceptable because the noise on these points is also not-Gaussian but instead due to some instrument malfunction or

spectral fitting error.

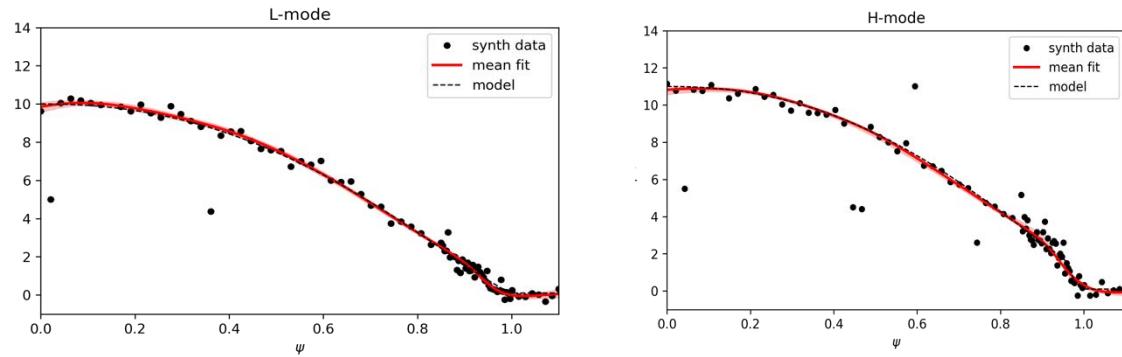


Fig. 13. Example fits for L-mode (left) and H-mode (right) using GPR with the developed kernel and a student-T likelihood function that mitigates the effects of outliers.

VI. 3D PERTURBED EQUILIBRIUM DATABASE AND MOR REPRESENTATION

Next, we report our first attempt to generate a large numerical database of perturbed 3D equilibria from various tokamak devices, utilizing the MARS-F 3D MHD code [59]. A 3D equilibrium is obtained when the plasma, magnetically confined in tokamak, responds to an externally applied 3D magnetic field perturbation. Such a magnetic perturbation is typically generated by currents flowing in window-frame coils installed outside the plasma, often for the purpose of controlling certain instabilities (e.g. the so-called edge localized mode, ELM) that are of critical importance for the tokamak operation.

The macroscopic plasma response is computed by the MARS-F code, solving the MHD equations in toroidal tokamak geometry and including important physics effects such as the finite resistivity of the plasma (as a conducting fluid) and the finite flow velocity of the plasma along the toroidal angle of the torus [60]. The MHD solution for the perturbed 3D fields, superposed with the 2D equilibrium fields discussed in the previous sections, form a perturbed 3D equilibrium. Essential quantities representing a perturbed 3D equilibrium thus include the perturbed plasma current (in particular the parallel current), the 3D displacement of the plasma, among other fields, as a result of the plasma response to the externally applied 3D fields. These 3D quantities associated with the MHD solution form the key part of the perturbed 3D equilibrium database.

A database of over 1000 3D equilibria has been numerically generated, covering 4 tokamak devices: DIII-D, MAST, ASDEX Upgrade, and ITER. Each datapoint is associated with (i) the 2D equilibrium, (ii) the specific toroidal row of the ELM-control coils that is utilized to generate the external 3D field, (iii) the harmonic number n (which we choose from 1 to 4) specifying the periodicity of the applied perturbation along the toroidal angle. It is important to note that the plasma response with different rows of ELM control coils can be straightforwardly combined (via linear superposition) to form a new 3D equilibrium. Therefore, the present database can be easily expanded without resorting to additional MARS-F computations. In terms of the machine memory requirement, each 3D equilibrium contains about 16 MB data, with the fundamental database

(i.e. without the aforementioned linear superposition) containing about 15 TB of data. We also note that this database is already stored in memory-optimized format. For instance, the spatial variation of the perturbed fields along the poloidal angle of the torus is stored in Fourier harmonics, which is a much more efficient representation than that in the real space.

The large amount of the 3D equilibrium data suggests that MOR techniques may play an important role. This serves two purposes: (i) to reduce the huge size of the data in the full 3D representation (by more than one order of magnitude), and more importantly (ii) to facilitate training of the ML algorithms for achieving real-time reconstruction of 3D equilibria [61], based on the 2D equilibrium input.

In this study, we apply the singular-value-decomposition (SVD) to each datapoint. Figure 14 shows one example, where we apply SVD to the MARS-F computed perturbed parallel plasma current density δJ^{\parallel} . δJ^{\parallel} provides the major contribution to the perturbed 3D magnetic field (as part of the 3D equilibrium) due to the plasma response. Moreover, δJ^{\parallel} provides the shielding current that is responsible for screening of the resonant magnetic perturbations [60]. It is evident, from Fig. 14(a), that the singular value decays fast, with the first two playing the dominant role for the considered example. Retaining only the first two SVD-eigenstates, we find that δJ^{\parallel} is well reconstructed [Fig. 14(c)] as compared to the true quantity [Fig. 14(b)]. The relative error is about 10%. This shows that low-order SVD can indeed be efficient for MOR-representation of the 3D equilibrium quantities, despite the fact that the latter quantities substantially vary in space.

The example shown above does not represent the best nor the worst situations with respect to the SVD truncation error. Application of the same technique through all data points, we find that it is generally desirable to retain more SVD-eigenstates, in order to achieve a reasonable accuracy of the MOR-representation for the whole database. An illustrative example of truncation is reported in Fig. 15, while retaining the first 5 eigenstates. Plotted here is the relative error of the SVD-based MOR-representation versus the ratio of the fifth to the first singular values, S_5/S_1 , for

δJ^{\parallel} of the plasma. The SVD is applied to each individual 3D equilibrium from the database. As a general trend, we note that the relative error decreases with decreasing S_5/S_1 . This is understandable, since smaller values of S_5/S_1 imply faster convergence of the SVD-series. Consequently, retaining the first five eigenstates allows more accurate MOR-representation of the data. For the whole database [Fig. 15(a)], about 75% of data points can be represented by the first five SVD-eigenstates with the relative error of 20%. The percentages of course increase (decrease) with increasing (decreasing) the level of the relative error.

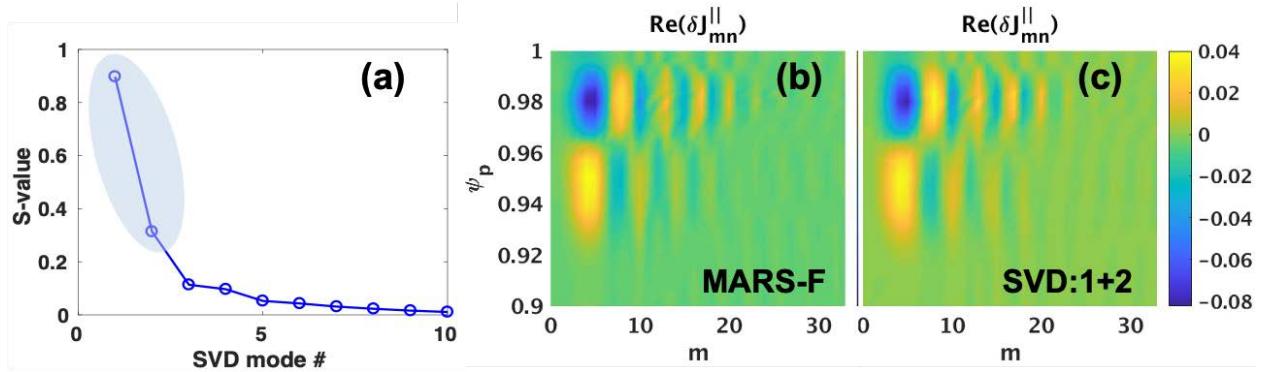


Fig. 14. An example of the SVD-based MOR-representation of the perturbed 3D equilibrium, for the poloidal Fourier harmonics of normalized δJ^{\parallel} in the plasma edge region: (a) the first 10 singular values, (b) the MARS-F computed δJ^{\parallel} as a 2D function of the plasma radial coordinate (labeled by the normalized equilibrium poloidal flux along the vertical axis) and the poloidal Fourier harmonic (denoted as m along the horizontal axis), and (c) the reconstructed δJ^{\parallel} from the first two SVD-eigenstates indicated by the shaded region in (a). The plasma response is obtained assuming 1 kA-turn of the coil current in the upper row of the I-coils in the $n=3$ configuration for a DIII-D plasma from discharge 157376.

As an interesting observation, it is evident from Fig. 15(a) that the relative error is generally larger for the MAST data. The most likely reason is that MAST is a low aspect ratio ($R/a \sim 1.4$) device, resulting in a stronger toroidal coupling between Fourier harmonics. It is thus reasonable that the convergence of the SVD-series is also generally slower. By excluding the MAST data from the database, i.e. only keeping the data from the conventional aspect ratio devices [Fig. 15(b)], we find that about 95% of data points can be represented by the first five SVD-eigenstates with the relative error of 20%. Similar level of accuracy is achieved for another key perturbed

quantity associated with the perturbed 3D equilibrium, i.e. the radial displacement of the plasma [61].

The SVD-reduced model allows efficient training of neural networks for fast reconstruction of 3D equilibria. For instance, the output of the neural network can be the first few singular values of the decomposition and the associated eigenvectors. This output can then be used to approximately reconstruct a 3D perturbed equilibrium, using as input the 2D equilibrium information in combination with known 3D source fields. The final product of the described workflow can then be potentially incorporated into a real-time plasma control system for many purposes, including monitoring of asymmetric events during discharges and ultimately controlling these 3D events.

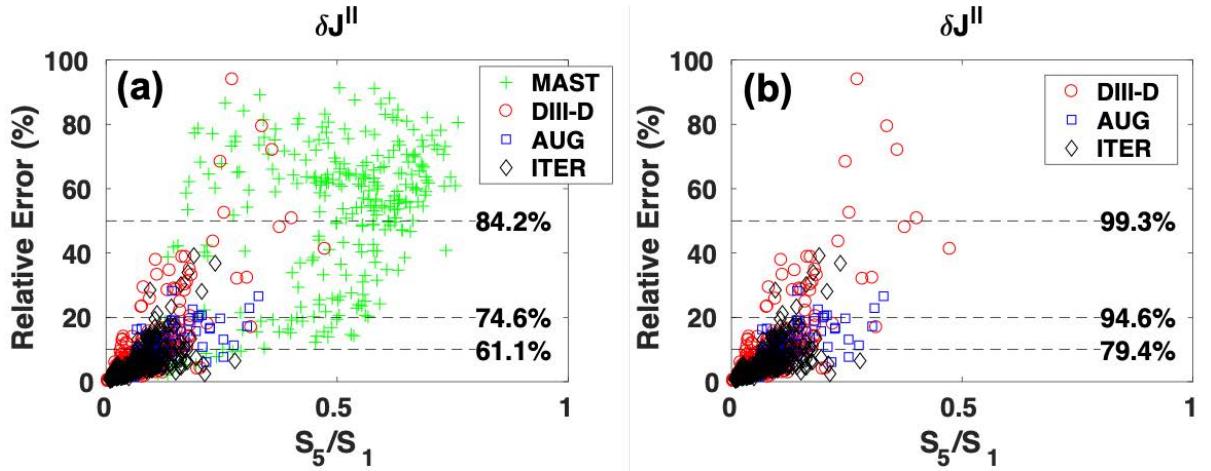


Fig. 15. Relative error obtained with the first five SVD eigenstates, to represent the MARS-F computed $\delta J^{\parallel\parallel}$ versus the ratio S_5/S_1 of the fifth to the first singular values. Compared are the results for (a) the whole database of perturbed 3D equilibria including all four devices, and (b) the database covering only the three conventional aspect ratio devices excluding the spherical tokamak MAST. The horizontal dashed lines indicate the percentage of the data points (from the respective database) achieving the corresponding levels (10%, 20%, 50%) of relative error.

VII. SUMMARY

In this paper, an overview of the recent progress in the application of ML/AI algorithms to extend EFIT equilibrium reconstruction capability is presented. This includes creation of a device-independent portable core equilibrium solver to facilitate adaptation of ML/AI algorithms, construction of a large DIII-D EFIT database for developments of EFIT MOR surrogate models to reconstruct approximate solutions. A neural-network (NN) MOR surrogate model has been successfully trained and tested using the magnetically reconstructed datasets with very encouraging results. Other progress includes developments of a GP Bayesian framework that can adapt its many hyperparameters to improve processing of experimental input data and a 3D perturbed equilibrium database from toroidal full MHD linear response modeling using the MARS-F MHD code for developments of 3D-MOR surrogate models.

This progress lays a solid foundation moving forward to create the full EFIT-AI: equilibrium reconstruction making maximum use of experimental data in a burning plasma environment to understand, control, and optimize tokamaks. We fully expect EFIT-AI to become routine for between-shot and post-shot analysis. The degree to which EFIT-AI can serve for real-time control depends on how much faster we can make EFIT-AI, especially in the area of the Bayesian experimental analysis. We are increasing the amount of work to be done in analyzing the data (to handle noisy data more accurately), as well as speeding up the process through the use of neural networks for initial conditions and in hyperparameter selection. Similarly, incorporating 3D effects, will improve the accuracy of the equilibrium reconstruction and provide useful information for interpretation of important 3D effects such as local divertor heat fluxes in tokamaks, but the extent to which it will impact real-time control will depend on the speed and accuracy of EFIT-AI. As demonstrated by the references, this effort of improving equilibrium reconstruction and data analysis using machine learning is of high interest within the fusion community. EFIT-AI is systemically learning from these efforts, and from our applied math and computer science colleagues, to integrate this knowledge into the production-level reconstruction that has been the

hallmark of the EFIT code.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award(s) DE-SC0021203, DE-FC02-04ER54698, DE-FG02-95ER54309, and GA IR&D.

DISCLAIMER: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

- [1] E.J. Strait, “Stability of High Beta Tokamak Plasmas,” *Phys. Plasmas*, **1**, 1415 (1997);
<https://doi.org/10.1063/1.870691>.
- [2] F.M. Levinton, M.C. Zarnstorff et al., “Improved Confinement with Reversed Magnetic Shear in TFTR,” *Phys. Rev. Lett.*, **75**, 4417 (1995);
<https://doi.org/10.1103/PhysRevLett.75.4417>.
- [3] E.J. Strait, L.L. Lao et al., “Enhanced Confinement and Stability in DIII-D Discharges with Reversed Magnetic Shear.” *Phys. Rev. Lett.*, **75**, 4421 (1995);
<https://doi.org/10.1103/PhysRevLett.75.4421>.
- [4] L.L. Lao et al., “Reconstruction of Current Profile Parameters and Plasma Shapes in Tokamaks,” *Nucl. Fusion*, **25**, 11 (1985); <https://doi.org/10.1088/0029-5515/25/11/007>.
- [5] L.L. Lao et al., “Equilibrium Analysis of Current Profiles in Tokamaks,” *Nucl. Fusion*, **30**, 1035 (1990); <https://doi.org/10.1088%2F0029-5515%2F30%2F6%2F006>.
- [6] L.L. Lao et al., “MHD Equilibrium Reconstruction in the DIII-D Tokamak,” *Fusion Sci. Technol.*, **48**, 968 (2005); <https://doi.org/10.13182/FST48-968>.
- [7] D.P. O’Brien et al., “Equilibrium Analysis of Iron Core Tokamaks Using a Full Domain Method,” *Nucl. Fusion* **32**, 1351 (1992); <https://doi.org/10.1088/0029-5515/32/8/I05>.
- [8] B. Nelson et al., “Formation and Sustainment of a 150 kA Tokamak by Coaxial Helicity Injection,” *Phys. Rev. Lett.* **72**, 3666 (1994); <https://doi.org/10.1103/PhysRevLett.72.3666>.
- [9] A. Sykes et al., “The Spherical Tokamak Programme at Culham,” *Nucl. Fusion*, **39**, 1271 (1999); <https://doi.org/10.1088/0029-5515/39/9Y/305>.
- [10] L.C. Appel, M.K. Bevir, and M.J. Walsh, “Equilibrium Reconstruction in the START Tokamak,” *Nucl. Fusion*, **41**, 169 (2001); <https://doi.org/10.1088/0029-5515/41/2/303>.

[11] L.C. Appel et al., “A Unified Approach to Equilibrium Reconstruction,” 33rd EPS Conf. on Controlled Fusion and Plasma Physics (Rome, Italy) vol 30C (ECA) P-2.184.

[12] B. Lee, N. Pomphrey, and L.L. Lao. “Physics Design of Poloidal Field, Toroidal Field, and External Magnetic Diagnostics in KSTAR,” *Fusion Tech.* **36**, 278 (1999); <https://doi.org/10.13182/FST99-A108>.

[13] Y. In et al., “Resistive $n = 1$ Modes in Reversed Magnetic Shear Alcator C-Mod Plasmas,” *Nucl. Fusion* **40**, 1463 (2000); <https://doi.org/10.1088/0029-5515/40/8/304>.

[14] T. Oikawa et al., “Heating and Non-inductive Current Drive by Negative Ion Based NBI in JT-60U,” *Nucl. Fusion* **40**, 435 (2000); <https://doi.org/10.1088/0029-5515/40/3Y/301>.

[15] S.A. Sabbagh et al., “Equilibrium Properties of Spherical Torus Plasmas in NSTX,” *Nucl. Fusion* **41**, 1601 (2001); <https://doi.org/10.1088/0029-5515/41/11/309>.

[16] W. Zwingmann, “Equilibrium Analysis of Steady State Tokamak Discharges,” *Nucl. Fusion* **43**, 842 (2003); <https://doi.org/10.1088/0029-5515/43/9/308>.

[17] J. Li et al., “Long Pulse Enhanced Confinement Discharges in the HT-7 Superconducting Tokamak by Ion Bernstein Wave Heating and Lower Hybrid Wave Current Drive,” *Phys. Plasmas* **10**, 1653 (2003); <https://doi.org/10.1063/1.1556297>.

[18] H. Hongda, J. Zhang, J. Dong, and Q. Li, “Study of MHD Equilibrium in HL-2A Tokamak,” *Plasma Sci. Technol.*, **8**, 397 (2006); <https://doi.org/10.1088/1009-0630/8/4/06>.

[19] J. Qian et al., “Equilibrium Reconstruction in EAST Tokamak,” *Fusion Sci. Tech.* **11**, 142 (2009); <https://doi.org/10.1088/1009-0630/11/2/03>.

[20] Y.S. Park et al., “KSTAR Equilibrium Operating Space and Projected Stabilization at High Normalized Beta,” *Nucl. Fusion* **51**, 053001 (2011); <https://doi.org/10.1088/0029-5515/51/5/053001>.

[21] Y.G. Li et al., “EFIT Equilibrium Reconstruction Including Polarimetry Measurements on Tore Supras,” *Fusion Sci. Technol.*, **59**, 397 (2011);
<https://doi.org/10.13182/FST11-A11654>.

[22] G.Q. Li et al., “Kinetic Equilibrium Reconstruction on EAST Tokamak,” *Plasma Phys. Control. Fusion* **55**, 125008 (2013); <https://doi.org/10.1088/0741-3335/55/12/125008>.

[23] Y. Huang et al., “GPU-Optimized Fast Equilibrium Reconstruction in Fine Grids for Real-Time Control and Data Analysis,” *Nucl. Fusion* **60**, 076023 (2020);
<https://doi.org/10.1088/1741-4326/ab91f8>.

[24] E. Xue, X. Zhang, K. Nakamura et al., “Equilibrium Reconstruction and Equilibrium Properties in QUEST Tokamak,” *J. Fusion Energy* **38**, 244 (2019);
<https://doi.org/10.1007/s10894-019-00209-5>.

[25] J.W. Berkery, S. A. Sabbagh, L. Kogan, D. Ryan, J. M. Bialek, Y. Jiang, D. J. Battaglia, S. Gibson, and C. Ham. "Kinetic Equilibrium Reconstructions of Plasmas in the MAST Database and Preparation for Reconstruction of the First Plasmas in MAST Upgrade," *Plasma Phys. Control. Fusion* **63**, 055014 (2021);
<https://doi.org/10.1088/1361-6587/abf230>.

[26] Y. Jiang, S. A. Sabbagh, Y. S. Park, J. W. Berkery, J. H. Ahn, J. D. Riquezes, J. G. Bak et al. , "Kinetic Equilibrium Reconstruction and the Impact on Stability Analysis of KSTAR Plasmas," *Nucl. Fusion* **61**, 116033 (2021); <https://doi.org/10.1088/1741-4326/ac26a4>.

[27] H. Grad and H. Rubin, “Hydromagnetic Equilibria and Force-Free Fields,” Proceedings of the 2nd UN Conf. on the Peaceful Uses of Atomic Energy, Vol. 31, Geneva: IAEA p. 190 (1958).

[28] V.D. Shafranov, “Plasma equilibrium in a magnetic field,” *Reviews of Plasma Physics*, Vol. 2, New York: Consultants Bureau, p. 103 (1966).

[29] B. Ph. van Milligen, V. Tribaldos, and J. A. Jiménez, “Neural Network Differential Equation and Plasma Equilibrium Solver.” *Phys. Rev. Lett.* **75**, 3594 (1995);
<https://doi.org/10.1103/PhysRevLett.75.3594>.

[30] C. Rea, R.S. Granetz, K. Montes, R.A. Tinguely, N.E. Eidietis, J.M. Hanson, B. Sammuli, Disruption prediction investigations using machine learning tools on DIII-D and Alcator C-Mod,” *Plasma Phys. Control. Fusion* **60**, 084004 (2018);
<https://doi.org/10.1088/1361-6587/aac7fe>.

[31] K.J. Montes, C. Rea, R.S. Granetz, R.A. Tinguely, N.W. Eidietis, O. Meneghini, D. Chen, B. Shen, B.J. Xiao, K. Erickson, “Machine learning for disruption warnings on Alcator C-Mod, DIII-D, and EAST,” *Nucl. Fusion* **59**, 096015 (2019);
<https://doi.org/10.1088/1741-4326/ab1df4>.

[32] J. Kates-Harbeck, A. Svyatkovskiy, W. Tang, “Predicting disruptive instabilities in controlled fusion plasmas through deep learning”, *Nature* **568**, 526 (2019);
<https://doi.org/10.1038/s41586-019-1116-4>.

[33] D. R. Ferreira, P. J. Carvalho and H. Fernandes, "Deep Learning for Plasma Tomography and Disruption Prediction From Bolometer Data," *IEEE Transactions on Plasma Science*, **48**, 36 (2020); <https://doi.org/10.1109/TPS.2019.2947304>.

[34] G. Dong, X. Wei, J. Bao, G. Brochard, Z. Lin and W. Tang. “Deep learning based surrogate models for first-principles global simulations of fusion plasmas,” *Nucl. Fusion* **59**, 096015 (2019) **61**, 126061 (2021); <https://doi.org/10.1088/1741-4326/ac32f1>.

[35] B H Guo, et al., “Disruption prediction on EAST tokamak using a deep learning algorithm.” *Plasma Phys. Control. Fusion* **63**, 115007 (2021);
<https://doi.org/10.1088/1361-6587/ac228b>.

[36] S. Joung, J. Kim, S. Kwak, J. Bak, S. Lee, H. Han, H. Kim, G. Lee, D. Kwon, and Y.-C. Ghim, “Deep neural network Grad–Shafranov solver constrained with measured magnetic signals,” *Nuclear Fusion* **60**, 016034 (2019). <https://doi.org/10.1088/1741-4326/ab555f>.

[37] D. A. Kaltsas and G. N. Throumoulopoulos. “Neural network tokamak equilibria with incompressible flows,” arXiv preprint arXiv:2109.12850 (2021);
<https://arxiv.org/pdf/2109.12850.pdf>.

[38] Yuri. V. Mitrishkin, Pavel S. Korenev, Artem E. Konkov, Valerii I. Kruzhkov, and Nicolai E. Ovsianikov, “New Identification Approach and Methods for Plasma Equilibrium Reconstruction in D-Shaped Tokamaks,” *Mathematics* **10**, 40 (2022);
<https://doi.org/10.3390/math10010040>.

[39] J.T. Wai, M.D. Boyer, and E. Kolemen, “Neural Net Modeling of Equilibria in NSTX-U,” arXiv: 2202.13915v1 (2012), submitted to *Phys. Plasmas*;
<https://arxiv.org/pdf/2202.13915.pdf>.

[40] Q. Peng, et al., “A Linux cluster for between-pulse magnetic equilibrium reconstructions and other processor bound analyses,” *Rev. Sci. Instrum.* **72**, 3277 (2001);
<https://doi.org/10.1063/1.1384451>.

[41] J. R. Ferron, et al., “Real time equilibrium reconstruction for tokamak discharge control,” *Nucl. Fusion* **38**, 1055 (1998); <https://doi.org/10.1088/0029-5515/38/7/308>.

[42] CMake cross platform, open-source build system, <http://www.cmake.org/>.

[43] Thomas W. Fredian and Joshua A. Stillerman “MDSplus. Current developments and future directions,” *Fusion Engineering and Design* **60**, 229 (2002).

[44] O. Meneghini, G. Snoep, B.C. Lyons, et al., “Neural-Network Accelerated Coupled Core-Pedestal Simulations with Self-Consistent Transport of Impurities and Compatible with ITER IMAS,” *Nucl. Fusion* **61**, 026006 (2021);
<https://doi.org/10.1088/1741-4326/abb918>.

[45] O. Meneghini, S.P. Smith, L.L. Lao, et al., “Integrated Modeling Applications for Tokamak Experiments with OMFIT,” *Nucl. Fusion* **55**, 083008 (2015);
<https://doi.org/10.1088/0029-5515/55/8/083008>.

[46] L.L. Lao, H.E. St. John, R.D. Stambaugh, et al., “Separation of β_p and ℓ_i in tokamaks of non-circular cross-section,” *Nucl. Fusion* **25**, 1421(1985);
<https://doi.org/10.1088/0029-5515/25/10/004>.

[47] Diederik P. Kingma and Jimmy Ba, Adam, “A Method for Stochastic Optimization,” arXiv: 1412.6980 (2017).

[48] Jakob Svensson and European Fusion Development Agreement and JET Project, “Non-Parametric Tomography Using Gaussian Processes. EFDA (2011);
<https://books.google.com/books?id=g2SCMwEACAAJ>.

[49] Dong Li, Y. B. Dong, Wei Deng, et al., “Bayesian Tomography and Integrated Data Analysis Infusion Diagnostics,” *Rev. Sci. Instrum.* **87**, 11E319 (2016).

[50] Semin Joung, Jaewook Kim, Sehyun Kwak, Kyeoreh Park, S. H. Hahn, H. S. Han, H. S. Kim, J. G. Bak, S. G. Lee, and Y. c. Ghim, “Imputation of Faulty Magnetic Sensors with Coupled Bayesian and Gaussian Processes to Reconstruct the Magnetic Equilibrium in Real Time.,” *Review Sci. Instrum.* **89**, 10K106 (2018).

[51] Sehyun Kwak, Jakob Svensson, S Bozhenkov, Joanne Flanagan, Mark Kempenaars, Alexandru Boboc, Y-C Ghim, and JET Contributors, “Bayesian Modelling of Thomson Scattering and Multichannel Interferometer Diagnostics Using Gaussian Processes,” *Nucl. Fusion*, **60**, 046009 (2020).

[52] Sehyun Kwak, J. Svensson, S. Bozhenkov, H. Trimino Mora, U. Hoefel, A. Pavone, M. Krychowiak, A. Langenberg, Y. c. Ghim, and W7-X Team, “Bayesian Modelling of Multiple Plasma Diagnostics at Wendelstein 7-X,” arXiv: 2103.07582 (2021); <https://arxiv.org/abs/2103.07582>.

[53] Boyer, Mark D., and Jason Chadwick, "Prediction of Electron Density and Pressure Profile Shapes on NSTX-U Using Neural Networks," *Nucl. Fusion* **61**, 046024 (2021).

[54] M.A. Chilenski, M. Greenwald, Y. Marzouk, N.T. Howard, A.E. White, J.E. Rice, and J.R. Walk, “Improved Profile Fitting and Quantification of Uncertainty in Experimental Measurements of Impurity Transport Coefficients Using Gaussian Process Regression. *Nucl. Fusion* **55**, 023012 (2015).

[55] Aaron Ho, Jonathan Citrin, Fulvio Auriemma, Clarisse Bourdelle, Francis J Casson, Hyun-Tae Kim, Pierre Manas, Gabor Szepesi, Henri Weisen, and JET Contributors, “Application of Gaussian Process Regression to Plasma Turbulent Transport Model Validation via Integrated Modelling,” *Nucl. Fusion* **59**, 056007 (2019).

[56] Eric C Howell and JD Hanson, “Development of a Non-Parametric Gaussian Process Model in the Three-Dimensional Equilibrium Reconstruction Code V3FIT,” *Journal Plasma Phys.* **86**, 905860102 (2020); <https://doi.org/10.1017/S0022377819000813>.

[57] Sehyun Kwak, J. Svensson, O. Ford, L. Appel, Y. c. Ghim, and JET Contributors, “Bayesian Equilibria of Axisymmetric Plasmas,” arXiv: 2103.07581 (2021);
<https://arxiv.org/abs/2103.07581>.

[58] Liu, Zijie, Zhengping Luo, Tianbo Wang, Yao Huang, Yuehang Wang, Qingze Yu, Qiping Yuan, Bingjia Xiao, and Jiangang Li. "Plasma Current Profile Reconstruction for EAST Based on Bayesian Inference," Fusion Eng. Des. **172**, 112722 (2021).

[59] Y.Q. Liu, et al., “Feedback Stabilization of Nonaxisymmetric Resistive Wall Modes in Tokamaks. I. Electromagnetic Model,” Phys. Plasmas **7**, 3681(2000);
<https://doi.org/10.1063/1.1287744>.

[60] Y.Q. Liu, et al., “Full Toroidal Plasma Response to Externally Applied Non-Axisymmetric Magnetic Fields,” Phys. Plasmas **17**, 122502 (2010); <https://doi.org/10.1063/1.3526677>.

[61] Y.Q. Liu, et al. “Surrogate Models for 3-D Perturbed Magnetohydrodynamic Equilibria in Tokamaks,” to be submitted to Nucl. Fusion (2022).