Proceedings of Machine Learning Research vol 120:1–10, 2020 2nd Annual Conference on Learning for Dynamics and Control

Toward fusion plasma scenario planning for NSTX-U using machine-learning-accelerated models

Mark D. Boyer MBOYER@PPPL.GOV
Princeton Plasma Physics Laboratory

Editors: A. Bayen, A. Jadbabaie, G. J. Pappas, P. Parrilo, B. Recht, C. Tomlin, M.Zeilinger

Abstract

One of the most promising devices for realizing power production through nuclear fusion is the tokamak. To maximize performance, it is preferable that tokamak reactors achieve advanced operating scenarios characterized by good plasma confinement, improved magnetohydrodynamic (MHD) stability, and a largely non-inductively driven plasma current. Such scenarios could enable steady-state reactor operation with high *fusion gain* — the ratio of produced fusion power to the external power provided through the plasma boundary. Precise and robust control of the evolution of the plasma boundary shape as well as the spatial distribution of the plasma current, density, temperature, and rotation will be essential to achieving and maintaining such scenarios. The complexity of the evolution of tokamak plasmas, arising due to nonlinearities and coupling between various parameters, motivates the use of model-based control algorithms that can account for the system dynamics. In this work, a learning-based accelerated model trained on data from the National Spherical Torus Experiment Upgrade (NSTX-U) is employed to develop planning and control strategies for regulating the density and temperature profile evolution around desired trajectories. The proposed model combines empirical scaling laws developed across multiple devices with neural networks trained on empirical data from NSTX-U and a database of first-principles-based computationally intensive simulations. The reduced execution time of the accelerated model will enable practical application of optimization algorithms and reinforcement learning approaches for scenario planning and control development. An initial demonstration of applying optimization approaches to the learning-based model is presented, including a strategy for mitigating the effect of leaving the finite validity range of the accelerated model. The approach shows promise for actuator planning between experiments and in real-time.

Keywords: Fusion, Plasma Control, Reduced Modeling, Actuator Planning, Genetic Algorithms

1. Introduction

Fusion is the process by which two light nuclei combine to form a heavier nucleus and convert a small amount of mass into a large amount of energy. This energy release makes it a potential means for producing electrical power. The most likely fuels for fusion are plentiful: deuterium can be extracted from sea water and tritium can be bred from lithium. Furthermore, there is no associated risk of a runaway nuclear reaction, no generation of high-level nuclear waste or weapon-grade material, and no emission of green house gases or air pollution.

While these advantages make fusion an attractive alternative to the use of fossil fuels or nuclear fission for power production, fusion is extremely challenging from both scientific and technical perspectives. Fusion reactions require extremely high temperatures (on the order of 100 million degrees) to occur frequently enough to make a reactor economically viable. The fuel mixture be-comes a plasma at these temperatures — an ionized gas of independent negatively and positively charged particles. The tokamak (Wesson, 2004), one of the most promising devices for confining fusion plasmas, closes and twists an applied magnetic field lines into a helical structure, trapping the plasma inside a toroidal vessel and creating a magneto-hydrodynamic equilibrium.

The NSTX-U spherical torus is a low-aspect-ratio tokamak (Menard et al., 2017; Battaglia et al., 2018), which enables a more compact device that makes efficient use of magnetic fields for confinement. Experiments on NSTX-U aim to build on the results of NSTX (Kaye et al., 2015) to explore how aspect ratio affects the scaling of confinement quality with power, magnetic field, and plasma current (Kaye et al., 2007) and guide the design of future pilot plants (Menard et al., 2016). ITER, which is the next experimental step for tokamak fusion research currently under construction in France, will attempt to be the first device to achieve a burning plasma (one in which a majority of the heat needed to sustain the plasma comes from fusion reactions) and to show the scientific feasibility of a commercial nuclear fusion power plant. NSTX-U aims to study key elements of burning plasma physics, which will be critical to ITER and future reactors, including control of non-inductive scenarios (Gerhardt et al., 2011), fast ion instabilities (Fredrickson et al., 2018), plasma boundary physics, heat load management (Vail et al., 2019a,b), and disruption prediction (Gerhardt et al., 2013).

To facilitate and optimize ITER and future nuclear reactor operations, several challenging control problems must be addressed. An overview of these control problems can be found in (Pironti and Walker, 2005; Walker et al., 2006). Numerous actuators, including magnetic coils, neutral beams, antennae, and gas injection valves are used to initiate, shape, fuel, and heat the plasma. Many critical tokamak plasma-control problems can be posed as optimization problems, and numerical methods for solving such problems have been explored in (Wehner et al., 2019; Felici and Sauter, 2012; Xu et al., 2010; Ou et al., 2008). These approaches typically make use of either high-fidelity simulations from integrated tokamak modeling codes, like TRANSP (Poli et al., 2018; Hawryluk, 1981), which have significant computational costs, or reduced models, which sacrifice accuracy for speed. Recent work has shown that machine learning approaches can be used to generate highly accelerated surrogate models of computationally intensive fusion simulation codes that retain high accuracy (Boyer et al., 2019; Meneghini et al., 2017; Citrin et al., 2015). Machine learning can also enable training of highly flexible empirical models for aspects of tokamak plasma evolution that are not well modeled from first-principles. Machine learning simulation acceleration and system identification thus enables fast, high fidelity nonlinear modeling spanning a large range of operating space, relaxing the trade-off of fidelity and speed in optimization and real-time control applications. The development of high accuracy and high speed models motivates studying the application of sample inefficient but highly flexible and powerful global methods for optimization, planning, and feedback approaches, like genetic optimization and reinforcement learning.

In this work, a learning-based data-driven modeling approach is applied to the temperature, density, and fast ion profile evolution on NSTX-U. The modeling approach includes an accelerated first-principles-based neutral beam injection model, and data-driven models of the response of stored energy, temperature, and density profile evolution. The model is shown to be able to reproduce results of experimental discharges and executes in a time on the order of 1s per simulation (in contrast with hours that would be needed to complete a TRANSP simulation). The model is used here to study application of sequential-quadratic-programming and genetic algorithms to actuator trajectory optimization. Constraints on actuator magnitudes and rate limits are considered, and the uncertainty of the accelerated model predictions is included as a penalty term, enabling experiment operators to vary the trade off of exploration and exploitation through varying the weight on the penalty term. While the temperature and density profiles modeled in this work represent only a subset of the plasma states that need to be predicted for detailed scenario development, the demonstrated modeling approach is actively being extended to predict other important quantities, like the safety factor profile and equilibrium evolution. Other machine learning approaches are being developed for predictions of plasma stability on NSTX-U (Piccione et al., 2020), which will eventually be used as constraints in the trajectory optimization problem.

2. Modeling

2.1. Neutral beam modeling

One of the most time consuming calculations in the integrated modeling code TRANSP is NUBEAM, a Monte Carlo code that calculates the influence of neutral beam injection on plasma heating, current drive, and torque, among other effects. While the code is highly accurate, the computational cost imposes a burden on its application in optimization and real-time applications. In (Boyer et al., 2019) an accelerated surrogate model for NUBEAM, referred to here as NubeamNet, was developed by generating a large database of NUBEAM calculations for a range of plasma conditions relevant to the NSTX-U operating space. To handle the high-dimensionality of the inputs and outputs of the regression problem (inputs consist of 15 scalar + 4 profiles each sampled at 20 spatial locations, outputs consist of 4 scalars + 7 profiles each sampled at 20 points), the modeling approach makes use of principal component analysis. Effects on the plasma are determined by how the fast injected beam ions slow down over time as they interact with the bulk plasma in the reactor. The slowing down process depends on plasma and beam parameters, which can vary throughout a discharge. To account for this in the model, inputs to the neural network are first filtered with a bank of three separate low pass filters with time constants chosen to span the expected range of beam slowing-down times. This gives the neural network a reduced representation of the relevant history of those inputs to adjust the output appropriately. Finally, an ensemble approach is used, in which multiple (3 in this work) neural networks each using 2 layers of 100 nodes each are trained on different subsets of the training data, and the output of the model uses the average of the neural network predictions. Large disagreement between models within the ensemble was shown to be a good indicator of extrapolation. This metric can be used to assess reliability of predictions or guide expansion of the training dataset. In this work, the metric is used to discourage trajectory optimization algorithms from returning results from unreliably modeled operating space. A comparison of NUBEAM predicted heating and current profiles with those predicted by the neural network is shown in Figure 1 showing good agreement during a TRANSP run that was not in the training dataset.

2.2. Energy and density evolution

While a great deal of progress has been made in theoretical understanding and computational modeling of the turbulence that dominates transport in tokamak plasmas, there is still much work to be done to enable consistent, accurate predictions of the evolution of these profiles from first-principles models. Accuracy aside, these models are far too computationally intensive for use in optimization and real-time applications. While recent work has shown that models for transport coefficients can be accelerated through the use of neural networks (Meneghini et al., 2017; Citrin et al., 2015), we

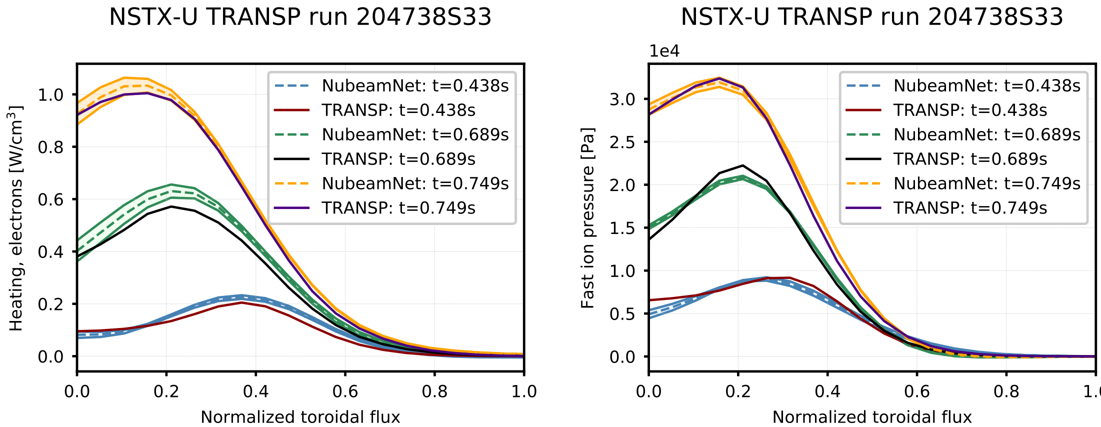

Figure 1: Comparison of NUBEAM calculated beam heating to electrons and fast ion pressure profiles predicted by the neural network. The shaded regions represent the standard deviation of the ensemble of neural network predictions around the mean.

take an alternative data-driven approach that is fast and can reproduce experimental profile evolution well enough for the shot planning and control applications it is developed for.

Since the shape of the temperature and density profile (distribution along the radial direction of the plasma cross-section) are typically observed to be 'stiff', i.e., insensitive to the detailed distribution of sources, the electron temperature and pressure profiles are written as

$$n_e = \langle n_e \rangle \widehat{n_e} \tag{1}$$
$$n_e T_e = \langle n_e T_e \rangle \widehat{n_e T_e} \tag{2}$$

where $\langle \cdot \rangle$ represents the volume average and $\widehat{\cdot}$ denotes the volume-average normalized distribution profile shape. The profile shapes are modeled in this work by training a neural network to predict the shape of the electron density and temperature profiles measured during the 2016 NSTX-U experimental campaign. The model uses plasma current, plasma boundary shaping parameters (major radius, minor radius, elongation, etc.), volume-averaged electron density, and volume-averaged electron pressure as input and is developed using the same approach as described in (Boyer et al., 2019): principal component analysis is used to reduce dimensionality of the output profiles, and an ensemble of neural networks composed of 3 fully connected layers of 100 nodes are trained. As demonstrated in the example results shown in Figure 2, the model is able to accurately reproduce the shape of the electron density and pressure profiles for a shot not included in the training dataset.

Since the plasma density evolution on NSTX-U is typically dominated by recycling from the plasma facing components and not well controlled by available actuators, the volume-averaged electron density $\langle n_e \rangle$ is taken as a prescribed input to the model, where $\langle \cdot \rangle$ represents the volume average. The impurity and deuterium ion densities are calculated based on quasi-neutrality, and

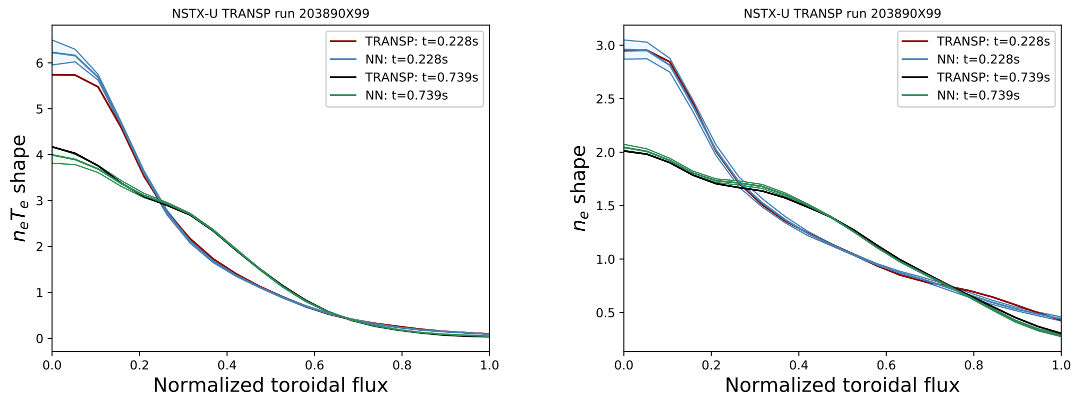

Figure 2: Comparison of experimental electron density and pressure profile shapes and those predicted by the neural network.

assuming a flat effective charge, Z_{eff}, profile and a single impurity species with atomic number Z_I:

$$n_I = n_e \frac{Z_{eff} - 1}{Z_I (Z_I - 1)} \tag{3}$$

$$n_D = n_e \frac{Z_I - Z_{eff}}{Z_I - 1}. \tag{4}$$

Impurity and deuterium ion temperature is assumed to scale with electron temperature, i.e.,

$$T_I = T_D = \alpha T_e. \tag{5}$$

The evolution of the plasma stored energy W is modeled as

$$\dot{W} = -\frac{W}{\tau_e} + P. \tag{6}$$

The absorbed power P is modeled as the volume integrated power from the NubeamNet model. For simplicity of presentation, terms due to radiation losses, and Ohmic heating are neglected. The ITER confinement scaling (ITER Physics Experts Groups, 1999) is used for τ_e:

$$\tau_{98y,2} = H_{98y,2} 0.0562 I_p^{0.93} B_T^{0.15} \bar{n}_e^{0.41} P_{Loss,th}^{-0.69} R_0^{1.97} \epsilon^{0.58} \kappa^{0.78}. \tag{7}$$

In this expression, I_p is the plasma current in MA, B_T is the toroidal magnetic field in T, $\bar{n}_e$ is the line-averaged electron density in #/m$^3 \times 10^{19}$, R_0 is the major radius in m, ϵ is the inverse aspect ratio, and κ is the elongation. Plasma boundary shaping parameters like major radius and elongation are considered to be prescribed in this work, but are controllable parameters that will be considered in the scenario optimization problem in future work. The loss power $P_{Loss,th}$ is in MW and taken here to be the total absorbed beam heating power less $\dot{W}$. The factor $H_{98y,2}$ is used to account for deviation of achieved confinement from the scaling model and is nominally set to 1.

The volume-averaged energy density $\langle E\rangle = W/V$ can then be written as

$$
\begin{aligned}
\frac{2}{3}\langle E\rangle &= \langle (n_I + n_D)T_I + n_e T_e\rangle & (8)\\
&= \langle \alpha(n_I + n_D)T_e + n_e T_e\rangle & (9)\\
&= \langle \alpha n_e T_e Z_{fac} + n_e T_e\rangle & (10)
\end{aligned}
$$

where $Z_{fac} = \left[\frac{Z_{eff}-1}{Z_I(Z_I-1)} + \frac{Z_I - Z_{eff}}{Z_I-1}\right]$. Assuming a flat Z_{eff} profile allows the energy density to be written as

$$
\begin{aligned}
\frac{2}{3}\langle E\rangle &= \langle n_e T_e\rangle \left(1 + \alpha Z_{fac}\right) & (11)\\
\langle n_e T_e\rangle &= \frac{2}{3}\langle E\rangle / \left(1 + \alpha Z_{fac}\right). & (12)
\end{aligned}
$$

2.3. Model results

The inputs to the model described by (1)-(12) are the time histories of injected beam powers and the total toroidal plasma current trajectory as well as model parameters H_{98y2}, Z_{eff}, Z_I, and α. While the main impurity species is typically carbon on a device with graphite tiles like NSTX-U, the parameters H_{98y2}, Z_{eff}, and α are expected to vary from discharge to discharge due to variations in experimental conditions, e.g., interactions with plasma facing wall components. For the purposes of control and optimization, they can either be treated as uncertain parameters, can be fit to reproduce the observables of a specific experimental discharge, or can be estimated in real-time using a dynamic observer. To demonstrate the behavior of the model, the free parameters were iteratively selected to approximately reproduce the results of NSTX-U discharge 204118. The predicted electron density and temperature profile evolution is compared to the measured evolution in Figure 3. While the predictions due not match the experiment perfectly, they are anticipated to be suitable for use in between-shots and real-time optimization, and require a calculation time of < 1s, in contrast to hours required to simulate the same discharge with the TRANSP code.

3. Trajectory optimization

While the model described here only predicts a subset of the quantities that define the plasma state in a tokamak experiment, we use it here to demonstrate an envisioned application of learning based models for scenario optimization and control. We consider the problem of optimizing the cost

$$
\begin{aligned}
J_0 = \int_{t_i}^{t_f} &\left[(n_e(0) - n_{e,t}(0))^2 + (n_e(0.5) - n_{e,t}(0.5))^2\right.\\
&\left.+(P_{fi}(0) - P_{fi,t}(0))^2 + (P_{fi}(0.5) - P_{fi,t}(0.5))^2\right] dt, \qquad (13)
\end{aligned}
$$

which weights the error between electron pressure n_e and fast ion pressure P_{fi} and time-varying target trajectories $n_{e,t}$ and $P_{fi,t}$ at two radial locations in the plasma (normalized radius $\rho = 0.0$ and $\rho = 0.5$). The evolution of these states is subject to the dynamic equations (1)-(12). The decision variables are formed by parameterizing the injected beam power and plasma current trajectories as piecewise linear functions, i.e., they are the amplitude of the inputs at a finite set of points in time, as well as the time between those points. Decision variables are constrained to enforce magnitude

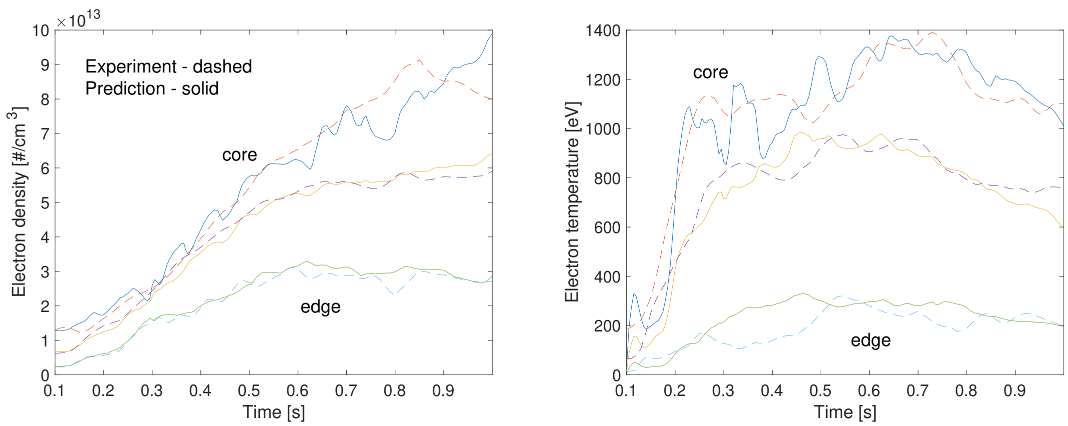

Figure 3: Comparison of experimental and predicted electron density and temperature profiles for discharge 204118, showing good agreement in both the core and edge of the plasma.

and rate limitations of physical actuators and/or plasma stability considerations, e.g., the maximum available power from each of the 6 beams on NSTX-U is approximately 2MW, and ramping the plasma current too fast can lead to MHD plasma instabilities that could terminate the discharge.

Because the models are learned from a finite set of training data, it is possible that a planning algorithm could lead to predictions that are unreliable. To avoid this, each of the learning based models uses an ensemble approach. The uncertainty in the model is approximated by the variance of the predictions of the ensemble, $\sigma^2_{n_e}$ and $\sigma^2_{P_{fi}}$, and the time integral of the uncertainty is weighted in an augmented cost function:

$$J_1 = J_0 + \int_{t_i}^{t_f} [\lambda_{n_e}(\sigma^2_{n_e}(0) + \sigma^2_{n_e}(0.5) + \lambda_{P_{fi}}(\sigma^2_{P_{fi}}(0) + \sigma^2_{P_{fi}}(0.5)]dt \tag{14}$$

By changing the weights λ_{n_e} and $\lambda_{P_{fi}}$, the optimizer can be made to exploit only the region of operating space that the models are trained on, or to enable exploration of the uncertain operating space. The latter could be used to help guide experimental exploration or, in the case of models trained on simulated data, can define regions of interest for conducting additional simulations and expanding the training dataset.

For planning problems with significant nonlinearities and constraints, gradient-based optimization is likely to find only a local optimal solution. Genetic algorithms can overcome this issue at the expense of the need for a large number of samples. By accelerating the simulation time by several orders of magnitude compared to the integrated modeling code TRANSP, the model described here makes it tractable to perform genetic optimization to find a good candidate solution, followed by sequential quadratic programming to refine the solution.

A population of 60 individuals is initialized. For each generation, 3 elite individuals are retained, while of the remaining individuals, 80% are generated through a cross-over operation and 20% are generated through a mutation operation. After several generations, the best solution is then refined through sequential quadratic programming. This process takes only a few minutes, and could be greatly accelerated by exploiting parallel computation of individual simulations in each generation, and for calculating numerical derivatives during the refinement step.

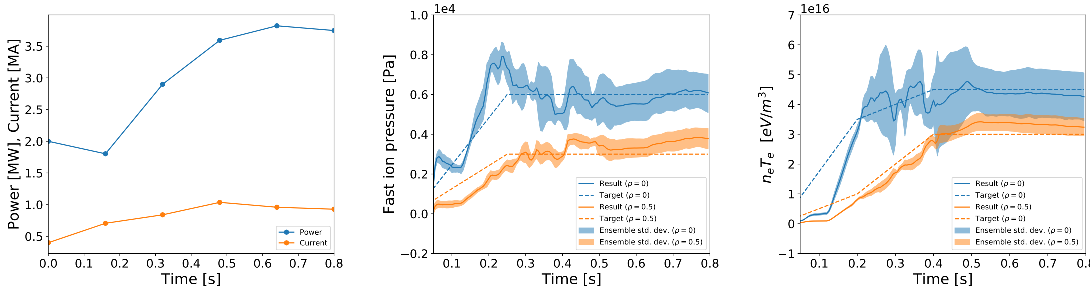

Figure 4: (left) Optimized actuator trajectories, (center) achieved and targeted fast ion pressure, and (right) achieved and targeted electron pressure.

Results of applying the optimization approach are shown in Figure 4. The plasma current and total beam power requests are ramped up to around 1MA and 3.6MW, respectively, and result in good tracking of the target trajectories for both fast ion pressure and electron pressure. As noted earlier, the optimization penalizes large ensemble standard deviations (depicted as shaded regions) to help ensure the obtained results are from the reliably modeled operating space.

4. Discussion

An approach to modeling a subset of tokamak profiles combining learning from empirical data and simulation acceleration was proposed. Application of optimization algorithms to an actuator planning problem was enabled by the fast execution time of the modeling approach. The approach is actively being developed to model additional profiles and behaviors, including the evolution of the plasma equilibrium, the current profile, and rotation profiles. The more advanced model is anticipated to enable a highly accelerated alternative to computationally intensive integrated models for scenario optimization and control. The rapid execution time will also facilitate studying application of reinforcement learning approaches for between shots planning and real-time applications.

The digital data for the figures in this paper can be found in:
http://arks.princeton.edu/ark:/88435/dsp011v53k0334.

Acknowledgments

This work was supported by the US Department of Energy Grant under contract number DE-AC02-09CH11466. 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

D.J. Battaglia, M.D. Boyer, S. Gerhardt, D. Mueller, C.E. Myers, et al. Scenario development during commissioning operations on the National Spherical Torus Experiment Upgrade. *Nuclear Fusion*, 58(4), 2018. ISSN 17414326. doi: 10.1088/1741-4326/aaa6e0.

M. D. Boyer, S. Kaye, and K. Erickson. Real-time capable modeling of neutral beam injection on NSTX-U using neural networks. *Nuclear Fusion*, 59(5), 2019. ISSN 17414326. doi: 10.1088/1741-4326/ab0762.

J. Citrin, S. Breton, F. Felici, et al. Real-time capable first principle based modelling of tokamak turbulent transport. *Nuclear Fusion*, 55(9):1–10, 2015. ISSN 17414326. doi: 10.1088/0029-5515/55/9/092001.

F Felici and O Sauter. Non-linear model-based optimization of actuator trajectories for tokamak plasma profile control. *Plasma Physics and Controlled Fusion*, 54(2):025002, feb 2012. ISSN 0741-3335. doi: 10.1088/0741-3335/54/2/025002. URL http://stacks.iop.org/0741-3335/54/i=2/a=025002?key=crossref.7771d3637a1052accfa9dff36ffbb072.

E.D. Fredrickson, E.V. Belova, N.N. Gorelenkov, et al. Global alfvén eigenmode scaling and suppression: experiment and theory. *Nuclear Fusion*, 58(8):082022, 2018. URL http://stacks.iop.org/0029-5515/58/i=8/a=082022.

S.P. Gerhardt, D.a Gates, S.M. Kaye, et al. Recent progress towards an advanced spherical torus operating point in NSTX. *Nuclear Fusion*, 51(7):073031, jul 2011. ISSN 0029-5515. doi: 10.1088/0029-5515/51/7/073031. URL http://stacks.iop.org/0029-5515/51/i=7/a=073031?key=crossref.69d73cd4231c28ed0f2b7f78a3460839.

S.P. Gerhardt, D.S. Darrow, R.E. Bell, et al. Detection of disruptions in the high-beta spherical torus NSTX. *Nuclear Fusion*, 53(6):063021, 2013. URL http://stacks.iop.org/0029-5515/53/i=6/a=063021.

RJ Hawryluk. An Empirical Approach To Tokamak Transport. In B. Coppi, editor, *Physics of Plasmas Close to Thermonuclear Conditions: Proceedings of the Course Held in Varenna, Italy, 27 August-8 September 1979*, volume 1, pages 19–46, Varenna, Italy, 1981. Elsevier Ltd. URL http://w3.pppl.gov/~pshare/transp/papers/Hawryluk.pdf.

ITER Physics Experts Groups. Chapter 2: Plasma Confinement and Transport. *Nuclear Fusion*, 39:2175, 1999. URL http://iopscience.iop.org/0029-5515/47/6/S02.

S.M Kaye, F.M Levinton, D Stutman, et al. Confinement and local transport in the National Spherical Torus Experiment (NSTX). *Nuclear Fusion*, 47(7):499–509, jul 2007. ISSN 0029-5515. doi: 10.1088/0029-5515/47/7/001. URL http://stacks.iop.org/0029-5515/47/i=7/a=001?key=crossref.369bd284c7af713934252da6cc883cae.

S.M. Kaye, T. Abrams, J.-W. Ahn, et al. An overview of recent physics results from NSTX. *Nuclear Fusion*, 55(10):104002, 2015. URL http://stacks.iop.org/0029-5515/55/i=10/a=104002.

J. E. Menard, T. Brown, L. El-Guebaly, M. D. Boyer, J. Canik, et al. Fusion nuclear science facilities and pilot plants based on the spherical tokamak. *Nuclear Fusion*, 56(10):106023, 2016. URL http://iopscience.iop.org/article/10.1088/0029-5515/56/10/106023/meta.

J. E. Menard, J. P. Allain, D. J. Battaglia, F. Bedoya, R. E. Bell, et al. Overview of NSTX Upgrade initial results and modeling highlights. *Nuclear Fusion*, 57(10):102006, 2017. URL http://iopscience.iop.org/article/10.1088/1741-4326/aa600a/meta.

O. Meneghini, S.P. Smith, P.B. Snyder, et al. Self-consistent core-pedestal transport simulations with neural network accelerated models. *Nuclear Fusion*, 57(8):086034, 2017. URL http://stacks.iop.org/0029-5515/57/i=8/a=086034.

Y. Ou, C. Xu, E. Schuster, et al. Design and simulation of extremum-seeking open-Loop optimal control of current profile in the DIII-D tokamak. *Plasma Physics and Controlled Fusion*, 50(11):115001, nov 2008. ISSN 0741-3335. doi: 10.1088/0741-3335/50/11/115001. URL http://stacks.iop.org/0741-3335/50/i=11/a=115001?key=crossref.bd7bf783b98b99157ddf908378d3ba8b.

A. Piccione, J.W. Berkery, S.A. Sabbagh, and Y. Andreopoulos. Physics-guided machine learning approaches to predict the ideal stability properties of fusion plasmas. *Nuclear Fusion*, 60(4):046033, 2020. ISSN 0029-5515. doi: 10.1088/1741-4326/ab7597.

A Pironti and M. Walker. Fusion, tokamaks, and plasma control: an introduction and tutorial. *IEEE Control Systems Magazine*, 25(5):30–43, 2005. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1512794.

F. Poli, J. Sachdev, J. Breslau, M. Gorelenkova, and X. Yuan. TRANSP. [Computer Software] doi:10.11578/dc.20180627.4, Jun 2018. URL https://dx.doi.org/10.11578/dc.20180627.4.

P. J. Vail, M. D. Boyer, A. S. Welander, and E. Kolemen. Design and simulation of the snowflake divertor control for NSTX-U. *Plasma Physics and Controlled Fusion*, 61(3), 2019a. ISSN 13616587. doi: 10.1088/1361-6587/aaf94a.

P. J. Vail, O. Izacard, and E. Kolemen. Optimization of the snowflake divertor for power and particle exhaust on NSTX–U. *Nuclear Materials and Energy*, 19(April):516–523, 2019b. ISSN 23521791. doi: 10.1016/j.nme.2019.03.003. URL https://doi.org/10.1016/j.nme.2019.03.003.

M. L. Walker, D. A. Humphreys, D. Moreau D. Mazon, M. Okabayashi, T. H. Osborne, and E. Schuster. Emerging applications in tokamak plasma control. Control solutions for next-generation tokamaks. *IEEE Control System Magazine*, 26(2):35–63, April 2006.

W. P. Wehner, E. Schuster, M. D. Boyer, and F. M. Poli. Transp-based optimization towards tokamak scenario development. *Fusion Engineering and Design*, 146(December 2018):547–550, 2019. ISSN 09203796. doi: 10.1016/j.fusengdes.2019.01.019. URL https://doi.org/10.1016/j.fusengdes.2019.01.019.

J Wesson. *Tokamaks*. Clarendon Press, Oxford, 3rd edition, 2004.

C. Xu, Y. Ou, J. Dalessio, et al. Ramp-up-phase current-profile control of tokamak plasmas via nonlinear programming. *IEEE Transactions on Plasma Science*, 38(2):163–173, feb 2010. ISSN 0093-3813. doi: 10.1109/TPS.2009.2037626. URL http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5395621.