

A general infrastructure for data-driven control design and implementation in tokamaks

Joseph Abbate^{1†}, Rory Conlin², Ricardo Shousha², Keith Erickson¹,
and Egemen Kolemen^{1,2‡}

¹Princeton Plasma Physics Laboratory, Princeton, NJ 08540, USA

²Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08540, USA

(Received xx; revised xx; accepted xx)

A general infrastructure for tokamak controllers based on data-driven neural net models is presented. The paradigm allows for more flexible choices of both the underlying model and the desired controlled variables and targets. The system is implemented and tested on the DIII-D tokamak, doing simultaneous pressure and temperature control via a finite-set model-predictive controller. Traditional control methods such as PID have proven effective for decoupled control tasks, but scale poorly when trying to achieve more complicated goals such as full state control. This is exactly where model-based controllers succeed.

1. Introduction

Modern tokamaks tend to use proportional-integral-derivative (PID) controllers, rather than more advanced controllers such as model predictive control (MPC), H-infinity, or reinforcement learning. PID controllers need trial-and-error tuning to obtain desired results, which requires many shots even when pre-tuned with simulations offline. However, operators largely are willing to pay this cost since it is infeasible to debug more sophisticated control schemes under a time-crunch. Creating a more accurate yet still intuitive and interpretable control would save shots and time, thereby increasing experimental efficiency.

The most robust control mechanisms at the moment are for plasma current, radial position, vertical position, and (to a lesser extent) Last Closed Flux Surface (LCFS) shape. On top of these base controls, other control algorithms are also sometimes employed. Error fields and Resistive Wall Modes (RWM) have been controlled via non-axisymmetric magnetic coils; Sawteeth and Neoclassical Tearing Modes (NTMs) can be controlled via steerable Electron Cyclotron Current Drive (see Walker *et al.* 2020).

One of the most difficult control problems, however, is full-profile magnetic and kinetic control. Theoretically, many of the modes and parameters that determine fusion performance are derived from these profiles; so by controlling the profiles the operator is directly controlling the quantities that matter. For example, pressure gradients around rational flux surfaces trigger MHD instabilities; kinetic profiles largely determine bootstrap current; and RWM/NTM stability depends on both pressure and rotational velocity. In modern tokamaks, some of the primary knobs for profile control include neutral beams, electron cyclotron waves, ion cyclotron waves, lower hybrid waves, helicon,

† Abbate and Conlin contributed equally to this paper, order is alphabetical

‡ Email address for correspondence: ekolemen@princeton.edu

cold gas fueling, pellet fueling, current ramp rate of the central solenoid, and 3D-coils (Humphreys *et al.* 2015).

A 2004 JET experiment (Laborde *et al.* 2005) used lower hybrid current drive (LHCD), ion-cyclotron resonance heating (ICRH), and neutral beam injection (NBI) in order to control 8 basis functions characterizing the safety factor q and the normalized electron-temperature-gradient ρ_{T_e} . Both profile and actuator quantities were assumed to be linear perturbations around an experimentally determined steady-state reference. The power deposition profiles from the actuators were assumed fixed (and determined from the reference around which they were linearized). A linear, time-independent response of the profiles to the three actuator values was also assumed. The system response matrix was empirically determined from previous open-loop experiments perturbing the plasma around the reference. This matrix was then used for proportional-integral feedback in an experiment simultaneously controlling q and ρ_{T_e} . A variety of other initiatives in the late 1990s / early 2000s also sought to control a limited number of proxy-quantities for the profile shape. Some used similar system identification methods for locally linearized model-based control: e.g. Moreau *et al.* (2011) for DIII-D and JT-60U. Others used simpler feedback (mostly Proportional) control: Ferron *et al.* (2006) for DIII-D, Wijnands *et al.* (1997) for Tore Supra, Fujita & team (2006) for JT-60U.

However, more complete control over all profiles (rather than a small number of proxies for individual profile shapes) and global models (rather than carefully tuned local linearizations) are required for some control applications. Through the 2010s a first-principles, model-driven approach was developed for controlling the full q profile (Barton *et al.* 2013; Wehner *et al.* 2016). The density profile is assumed to maintain a reference shape while varying linearly with the target density actuator \bar{n}_e (which is achieved via gas puffing by a separate control loop). The temperature profile is also assumed to maintain a reference shape, and scale like $\left(\frac{I_p}{I_p^{\text{ref}}}\right)^\gamma \left(\frac{P_{\text{tot}}}{P_{\text{tot}}^{\text{ref}}}\right)^\epsilon \left(\frac{\bar{n}_e}{\bar{n}_e^{\text{ref}}}\right)^\zeta$ for empirically determined coefficients γ , ϵ , and ζ ; and actuators of plasma current I_p , total power P_{tot} , and target density \bar{n}_e . The ion and electron temperatures are assumed equal, and the plasma Z_e is assumed constant in space and time. Noninductive current drive from each electron cyclotron gyrotron and each neutral beam is assumed to scale with a reference deposition profile times its power. Using simulations, intuition, and experimental data from the scenario of interest, the reference profiles, reference actuators, empirical coefficients, and other parameters are chosen. All of these assumptions combined allow for a linear poloidal flux diffusion (and in turn safety factor q) equation with an effective actuator set $\vec{u} = f(I_p, \bar{n}_e, \vec{P})$ for \vec{P} the individual power from each neutral beam and each electron cyclotron heating gyrotron, and f a nonlinear mapping from real to effective actuators.

Via finite difference, the flux diffusion equation in turn allows a linear state-space mapping $\vec{\theta}(t+1) = \mathbf{A}\vec{\theta}(t) + \mathbf{B}\vec{u}(t)$ with $\vec{\theta} = f(\vec{q})$ for f a nonlinear mapping. This state-space form, where the state evolution is a linear mapping of the present state and actuators, allows an abundance of existing control techniques to be employed. Recent experiments involved simultaneous control of the interior points of the q profile via linear model-predictive control to actuate individual powers of neutral beams and the target density, while constraining the model to use a feedforward (i.e. pre-programmed) electron cyclotron heating power. The total power from the beams is constrained to hit a desired total stored energy. The stored energy is assumed to evolve like $\frac{dE}{dt} = -\frac{E}{\tau_E} + P$, where the energy confinement time τ_E is obtained from the H_{98} (ITER Physics Expert Group on Confinement and Transport 1999) scaling; and the injected power P is assumed to consist only of the neutral beam heating since that is the dominant component. A proportional-integral (PI) controller is used to determine the best total injected neutral beam power

to hit a target stored energy. Meanwhile, I_p is excluded from the model-predictive control framework and set by a separate linear-quadratic-integral (LQI) control loop (described in (Boyer *et al.* 2013)) to hit a desired edge q profile value q_{95} . In essence then, the group has now demonstrated simultaneous q profile plus scalar kinetic (via the stored energy) control on DIII-D in a pseudo-first principles framework, which was tested specifically for the high- q_{\min} scenario (Holcomb *et al.* 2015).

An otherwise similar approach to first-principles control instead considers the full nonlinear diffusion equations and *locally* linearizes at each timestep. This method was used in a model-predictive control framework to do offline simulations (Felici & Sauter 2012), scenario optimization (Teplukhina *et al.* 2017), online state estimation (Felici *et al.* 2011), and realtime q profile plus β_n control on the TCV tokamak (Maljaars *et al.* 2015, 2017). This local linearization makes it easier to start from arbitrarily difficult equations, which means more naturally including kinetic profile diffusion (Te , Ti , and ne to date) in addition to current profile diffusion (Felici *et al.* 2018).

Unfortunately, after many decades of work developing transport equations the dynamics are still not well-captured, especially for kinetic profiles. The density profile is notoriously hard to predict; and all setups require some information at the edge (usually values at $\rho \sim 0.8$) to give decent results. In all the aforementioned methods, there are ad hoc assumptions required to get the diffusion equations to give semi-accurate results. As more progress continues to be made on the above models, we instead propose a new paradigm for control that uses a fully data-driven model to predict dynamics, and use that model to control all points on all profiles simultaneously. Specifically, we seek to use machine learning and other techniques to generate a model for the plasma state from a large database of experimental data. We then seek to use this model in a model-driven control framework. Machine learning has already been implemented for other areas of tokamak control, especially for disruption prediction (Montes *et al.* 2019; Fu *et al.* 2020; Rea *et al.* 2019). In this paper, we build on the work of (Abbate *et al.* 2021) and (Jalalvand *et al.* 2021) to do plasma profile control. We describe our workflow to train and deploy deep-learning plasma evolution models in control algorithms on the DIII-D tokamak (Luxon & Davis 1985), within the infrastructure of the machine’s standard Plasma Control System (PCS) (Penaflor *et al.* 2001). The framework is well-modularized so that anyone could train a model with the offline dataset, and then plug their control algorithm into the PCS via these tools. In this paper, we describe this infrastructure, and demonstrate it with an initial, simple test on the device: neural-net based finite set control of the core pressure and electron temperature via beam power modulation.

2. Infrastructure design

Controllers can be viewed as black boxes that map states to actions. These can be anything from simple PID controllers to much more advanced reinforcement learning algorithms.

We have built a basic infrastructure for doing model-based, data-driven control of plasma profiles. As shown in Figure 1, the infrastructure includes

- a workflow for generating data both offline and within the plasma control system of the same form so that models can be trained offline and deployed on the same type of data online.
- a user-interface in the control system to specify both the value and importance (“weight”) of each controlled state. At present, this includes 33 spatial components from each plasma profile.
- a code for converting neural networks written with the python deep learning library

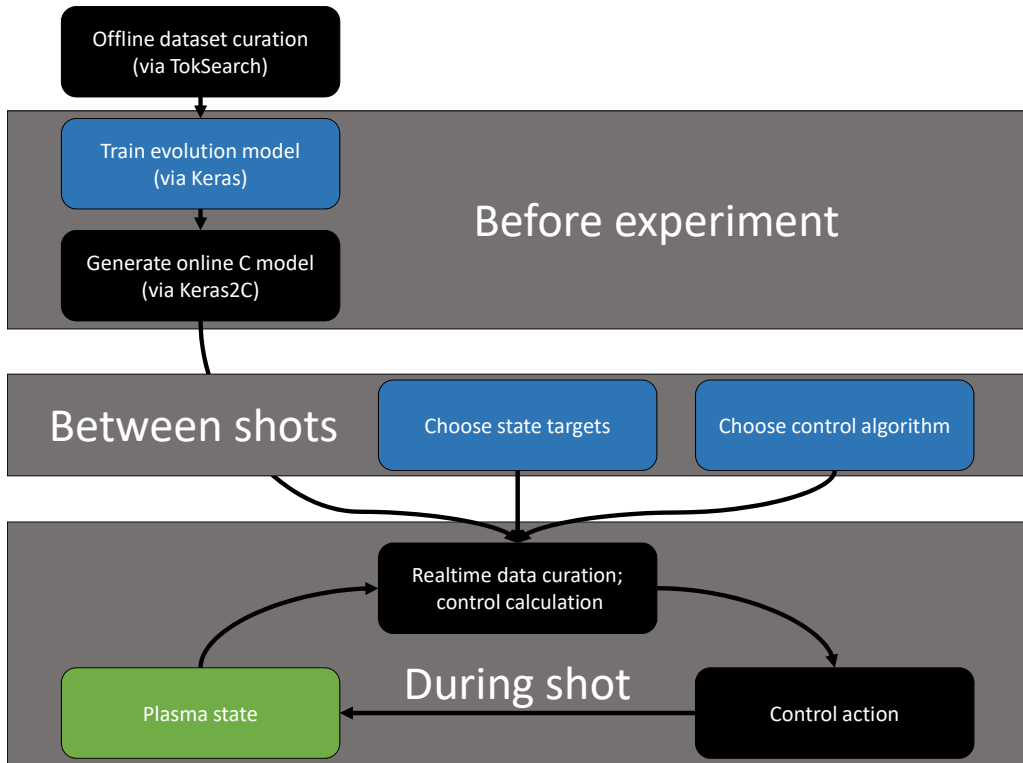


Figure 1. Diagram of new infrastructure for machine learning control on DIII-D. Historical data is used to train a machine learning model to predict plasma evolution, which is then converted to PCS code using Keras2C. The model is used in real time to predict plasma evolution under different actuator options, and the option that results in predictions closest to user defined targets is used to evolve the plasma state at each timestep. Blue corresponds to the parameters a user should change before experiments and during shots (as we describe in the Initial experimental control test section); black corresponds to infrastructure described in this section that a user generally would not change; and green corresponds to the state, which is defined by the behavior of the plasma.

Keras (Chollet 2015) into realtime-compatible C code to be used in plasma control systems, called Keras2C (Conlin *et al.* 2021).

These basic components comprise a code infrastructure within the DIII-D plasma control system modular enough to easily train a plasma evolution model, and swap in any control algorithm that uses such a model. Each of the components is described in this section.

2.1. Dataset curation

The MDSplus accessing tool TokSearch (Sammuli *et al.* 2018) is leveraged to generate an offline database of timetraces from previous experiments (about 150,000 samples from about 3,500 shots run between 2013 and 2018). Before the experiment, one can use this offline database to train a plasma evolution model and Keras2C to convert it into realtime-capable form. A pipeline written in C within the DIII-D Plasma Control System interface is used to generate similar samples “online” for use in control. Table 1 summarizes the signal sources used for online deployment vs offline training. As described in more detail in (Abbate *et al.* 2021), Zipfit is a standard DIII-D workflow

	Name	Offline Source	Online Source	Units
Profiles	n_e	Zipfit	mtanh	$10^{19}/\text{m}^3$
	T_e	Zipfit	mtanh	keV
	Rotation	Zipfit	spline	kHz
	$\iota = 1/q$	EFITRT2	EFITRT2	-
	Pressure	EFITRT2	EFITRT2	Pa
Actuators	Injected Power	pi nj	VEP	MW
	Injected Torque	ti nj	VEP	N·m
	Target Current	ipti pp	ipti pp	A
	Target Density	dstdenp	dstdenp	$10^{19}/\text{m}^3$
	Line Avg. Density	dssdenest	dssdenest	$10^{19}/\text{m}^3$

Table 1. Signals available both in realtime and offline for training control models and algorithms. “Source” denotes the MDS+ pointname or algorithm used to generate the data. `ipti pp` is the plasma current setpoint. `dstdenp` is the interferometer average density target setpoint. VEP is a subcontroller that sets individual beam powers to achieve the algorithm’s requested `pi nj` and `ti nj`.

that tries a variety of fitting mechanisms and chooses the option with the lowest χ^2 values; and EFITRT2 is the RealTime, kinetically-and-MSE (Motional Stark Effect) constrained version of the Grad-Shafranov solver EFIT (Lao *et al.* 2005). As is standard for many transport applications on DIII-D, thomson scattering is used to infer the electron temperature and density; Charge Exchange Recombination (CER) infers the ion temperature and rotation; \dot{B} probes, MSE, and pressure fits via Thomson scattering are primarily used for the equilibrium reconstruction; and an interferometer is used for the line average density. Note that while the safety factor q is more common in tokamak physics for describing the ratio of poloidal to toroidal magnetic field, due to the singularity at the separatrix, the rotational transform $\iota = 1/q$ was used in order to avoid numerical difficulties. The following subsection describes the profile-fitting mechanisms in more detail.

Note that `ipti pp` (the plasma current setpoint) and `dstdenp` (the interferometer average density target setpoint) are the exact commands used within a simple PI controller on the tokamak to set a desired current/density by adjusting the solenoid current ramp rate / gas valve flow (voltage). The offline database uses the true values of `pi nj` and `ti nj` sent into the device; while the online controller instead takes in the values requested by the Variable Energy and Perveance (VEP) algorithm. For control, energy and perveance are kept fixed, and VEP merely varies the duty cycle of the beams. VEP coordinates with each of the 8 beams to hit the specified total injected power and torque, given other constraints on the beams (Boyer *et al.* 2019). For example, in the control experiment described later in this paper one of the beams is configured to remain “on” in order to acquire MSE data; and 2 other beams to stay on to acquire CER data.

For more details on the data excluded, time-averaging, and smoothing used for the offline model, see (Abbate *et al.* 2021).

Realtime profile fitting

The profiles provided by rtEFIT are already fit as smooth functions of the normalized poloidal flux ψ . The CER (measuring rotation) and Thomson (measuring electron temperature and density) data are given as point measurements at fixed locations in the lab frame. This requires fitting the data to obtain the values at fixed locations in magnetic coordinates. To accomplish this, a real time fitting algorithm is used. The mapping

between the lab frame and flux coordinates is provided by rtEFIT, which calculates the poloidal flux on a grid in (R, Z) coordinates. These mapped points are then fit. A modified hyperbolic tangent with a linear core, described in (Groebner *et al.* 2001), is used for temperature and density profiles. Smoothing splines (Wand 2000) are used for the rotation profile.

The modified hyperbolic tangent basis was chosen to accurately capture the pedestal region at the expense of slightly less accuracy in the core, where the profile is constrained to be linear. A future upgrade will extend this to higher order polynomials in the core. Smoothing splines (equation 2.1) were chosen for the rotation profile as they are a flexible basis capable of fitting a wide range of shapes.

$$\hat{f} = \arg \min_{\hat{f}} \sum_{i=1}^N (y_i - \hat{f}(x_i))^2 + \lambda \int \hat{f}''(x)^2 dx \quad (2.1)$$

In contrast to standard interpolating splines, smoothing splines include a term to penalize the curvature of the resulting function, so in general they do not pass directly through the function values, which is desired when working with real time measurements that may be noisy.

These fitting techniques are placed in a separate algorithm, and are part of an ongoing effort to develop a real-time version of the Consistent and Automatic Kinetic EFIT (CAKE) code (Xing *et al.* 2021).

Buffering of fitted points

Each signal is obtained on its own sampling time, generally 1-10ms between measurements. Values are averaged in time in a boxcar fashion over 50ms windows, a window long enough to ensure multiple datapoints lie in the window but short enough that the plasma does not change too significantly (the energy confinement time at DIII-D is 100s of ms).

2.2. Plasma control system user interface

Between shots during an experiment, the user can input a (potentially time-varying) set of desired profiles within the plasma control system: electron temperature, electron density, plasma rotation, $\iota = 1/q$, and total pressure. The input at present is 33 evenly spaced points between $\psi_n = 0$ and $\psi_n = 1$, for ψ_n the normalized poloidal flux coordinate. The user also can input time-varying weights corresponding to each of the points within each of the profiles. Control algorithms within the framework are built to minimize the weighted least-squares error between the user-specified targets and the actual profiles:

$$\text{Error}(\vec{y}_1, \dots, \vec{y}_{\text{Num profiles}}) = \sum_{i=1}^{\text{Num profiles}} |\vec{w}_i \cdot (\vec{y}_i - \vec{y}_i^{\text{target}})|^2 \quad (2.2)$$

for \vec{w}_i the vectors of user-specified weights for each profile, $\vec{y}_i^{\text{target}}$ the values of the individual profiles, and $\vec{y}_i^{\text{target}}$ the user-specified target profiles.

2.3. Actuator commands

The controllers used in this framework can be thought of as high level controllers in a cascaded control scheme. After the high level controller determines the correct actuator values at the current timestep, they are sent to the respective lower level controllers which interface directly with the hardware. Neutral beam power and torque requests are sent to VEP as described above, which calculates the required duty cycle then interfaces directly




Figure 2. Inputs (blue proposals and green information from present and previous timesteps) and output (black profiles 200ms in the future) of the model. Note that each timestep of data is a boxcar average over all values between the timestep it represents and 50ms beforehand. Dashed line represents the timestep at which the prediction is made.

with the master beam algorithm. Requests for plasma current and average density are sent to their respective PI controllers to set the solenoid current ramp rate / gas valve flow (voltage).

3. Initial experimental control test

3.1. Control algorithm design

The previous sections discuss a new control system infrastructure for training a data-driven model offline and using it for multiple-profile control online. The plasma control system infrastructure allows users to specify desired target profiles and weights, and the control algorithm chooses actuator values at each cycle time to minimize the error.

In this section, a simple initial test deployed at DIII-D using this infrastructure is described. The system dynamics are identified via a smaller version of the neural network described in (Abbate *et al.* 2021). Figure 2 shows the structure of the model inputs and outputs. The model takes as input the present state defined by the plasma profiles. The model also takes as input the recent history of how a set of scalar parameters has been responding to the actuator commands. This empirically increases the model accuracy, and can be seen as a way of adapting the model to brand-new regimes (e.g. if the model sees the plasma β has stayed constant despite an increasing injecting power, the model can guess that a further increase in injected power may not be effective even if the model has not been trained on data from this specific scenario). Finally, the model takes as input proposals for actuators at each of four timesteps into the future (i.e. to the right of the dashed line). These are the values the control algorithm will be allowed to set. Given all of these inputs (the green and blue rectangles), the neural network predicts the future state (defined by plasma profiles) 200ms into the future (the black rectangle). The specific signals are shown in Table 1.

The model is used to choose actuator values via Finite Set Model-Predictive Control.

↑ pi nj				
Actuator	50ms	100ms	150ms	200ms
Δ target pi nj (MW)	0.31	0.63	0.94	1.26
Δ target ti nj (N m)	0.27	0.54	0.81	1.08
Δ target density ($10^{19}m^{-3}$)	0	0	0	0
Δ target current (MA)	0	0	0	0

constant				
Actuator	50ms	100ms	150ms	200ms
Δ target pi nj (MW)	0	0	0	0
Δ target ti nj (N m)	0	0	0	0
Δ target density ($10^{19}m^{-3}$)	0	0	0	0
Δ target current (MA)	0	0	0	0

↓ pi nj				
Actuator	50ms	100ms	150ms	200ms
Δ target pi nj (MW)	-0.31	-0.63	-0.94	-1.26
Δ target ti nj (N m)	-0.27	-0.54	-0.81	-1.08
Δ target density ($10^{19}m^{-3}$)	0	0	0	0
Δ target current (MA)	0	0	0	0

Table 2. Set of three “proposals” used to generate the inputs for the model in Finite Set Model-Predictive Control: desired (signed) changes in pi nj , ti nj , density, and current for each 50ms window in the 200ms time horizon over which profiles are predicted. In this case, pi nj and ti nj are proposed to linearly ramp while density and current are proposed to stay constant. For each proposal in the set, a model would predict the single-timestep evolution of the plasma state 200ms into the future. The first column of the proposal with the lowest predicted error 200ms from now (between the user-specified target and the prediction) is chosen as the control setpoint for the next timestep. These are the actual proposals used for DIII-D shot 187076.

As described in (Cortes *et al.* 2008), finite set control is one of the simplest options for control given a dynamics model of any form. In this control, at every timestep, the model predicts the future state of the plasma conditioned on each of a finite set of control options (i.e. “proposals”). In the implementation of this algorithm, a set of (potentially time-varying) proposals is specified in the user interface in the form shown in table 2. At each cycle within the shot, the neural net is evaluated to predict the resulting profiles from each proposal. The algorithm then selects the proposal that led to the smallest error between the profiles predicted and the target profiles specified by the user. This optimal proposal’s first column (i.e. the actuator value for the first timestep in the proposal’s trajectory) is used as the actuator setpoint for the upcoming timestep.

Figure 3 demonstrates the mechanism of finite set control for a timeslice in the experiment described in the following section. At this time, the proposals were those in table 2. The targets were the temperature and pressure values at $\psi_N = 0, 0.03,$ and 0.06 with weights of 1 (and all other weights 0). In this case, every target point for both pressure and temperature are lower than the corresponding present values by a significant margin. The algorithm predicts that the proposal with increasing pi nj would maintain or increase the pressure and temperature, while a decreasing pi nj would bring the values down very close to the targets over the coming 200ms. So the control algorithm in this case would choose to use the proposal with decreasing pi nj . Note that a full time-trace of pi nj is used as input to the model, 300ms into the past and 200ms (proposed) into the future. This time-dependent input is shown, along with the input profiles of temperature and pressure. For simplicity, the other inputs (table 1) are not shown in this figure.




Figure 3. Demonstration of the mechanism of finite set control during DIII-D shot 187076. At this timestep, three points in the core were targeted for both temperature and pressure (weight=1 for $\psi_N = 0, 0.03$; and 0.06 and 0 for all other points). p_{inj} was proposed to linearly decrease, remain constant, or linearly increase. The top two panels show the input to the model (grayed black) and the predictions given the three proposals. The black Xs mark the targets, demonstrating that the chosen proposal (the one with the lowest mean-squared distance between the target and the prediction) will be the decreasing p_{inj} proposal. The bottom panel shows the p_{inj} input to the model: the gray portion is the historical input, and each of the three proposals is included as input to the model for the corresponding predictions on the top plots.

3.2. Experimental checkout

As part of the initial test, core-profile control is attempted by varying injected power. Throughout the test, three proposals as shown in table 2 are given: a proposal with linearly decreasing injected power (π_{inj}), constant π_{inj} , and linearly increasing π_{inj} . Injected torque (τ_{inj}) is given a similar ramp to π_{inj} within each of the three proposals, which should make it easier for the VEP beam controller to properly fulfill the π_{inj} and τ_{inj} requests (especially since the machine’s beams were all pointing in the same direction on the day of the experiment). We do control in flattop only, loading in the rampup, shape parameters, and other settings from a plasma startup shot in 2020 (182500). We chose this shot because it had been an experimental test case for the realtime kinetic equilibrium reconstruction. Note that the model was trained on shots from campaign years 2010 through 2018, so that this shot was not seen during training. During the flattop control, all profile targets, when active, have the same shape from a timeslice just after rampup in the reference shot (182500 at 1.50s). All points on the profile are scaled by the same scalar value so that choosing a target is effectively just choosing the core value for each profile.

All experimental results in this section correspond to DIII-D shot 187076. The following subsections are best read while referencing the annotated experimental trajectory shown in figure 4 which plots timetraces of pressure and temperature at $\psi_N = 0.03$.

Phase I

The algorithm takes full control of the beams at 1.35s. In this first phase, the core values (from $\psi_N = 0$ to $\psi_N = 0.06$) of pressure and temperature are controlled simultaneously. The same pressure profile from the baseline slice is used, but the temperature profile is scaled from a 3.3keV core down to a 2.0keV core (so that effectively the algorithm will attempt to achieve a combination of pressure and temperature outside the range of its training data). Both pressure and temperature (the thick gray timetraces) are initially above the targets (the black dashed lines). However, from the beginning of control at 1.35s until 1.73s, an error in the beam setup meant that Motional Stark Effect (MSE) data was not available, which in turn meant that the realtime (MSE-constrained) equilibrium fitting code was not returning a pressure and safety factor (q) profile. The controller properly floated previous values of pressure and q in light of the error, but this meant that the controller thought the pressure was above the target when really it was below by about 1.6s. The algorithm therefore continued decreasing beam power for much of this red shaded region despite that the true pressure profile was below the target.

At 1.73s, rtEFIT recovers, and immediately the controller begins to increase π_{inj} . At 1.97s, the controller begins to decrease the injected power to settle down on the pressure and temperature targets. At this point, the profiles are still below the targets, so the controller (like a derivative controller) is effectively looking ahead to slow down the upward trajectory as the inputs come in to the final value. At 2s the phase switches to a new set of targets.

Phase II

During this next phase, pressure control is turned off and the temperature target is scaled up to have a core value of 3keV. Once again (likely due to π_{inj} having been increasing through the previous timesteps), all proposals are predicted to yield an increasing temperature. The controller sensibly chooses to uniformly increase π_{inj} between 2.0 and 2.42s, and the temperature increases. A slight overshoot occurs, but (like a proportional controller) the algorithm begins decreasing the injected power right as the temperature passes the target. The value flattens out quite close to the targeted

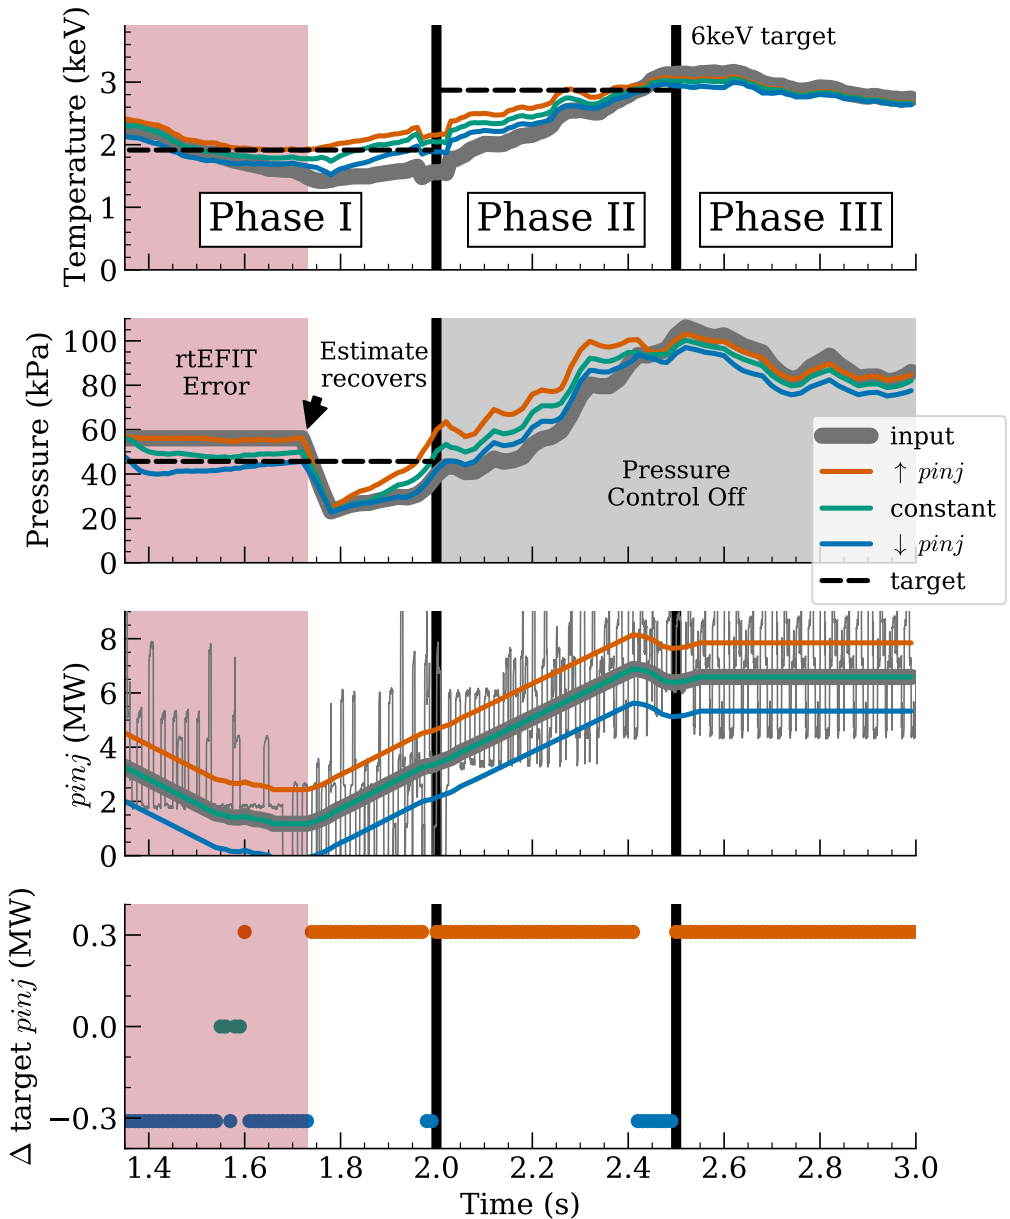


Figure 4. DIII-D shot 187076. On top two panels, true values of pressure and electron temperature at $\psi_N = 0.03$ are plotted over time in gray. In blue, green, and orange the model's prediction 200ms into the future is tracked proposing a decreasing, constant, and increasing injected power (respectively). The bottommost plot shows the change in injected power (p_{inj}) requested based on the winning proposal, and the plot above shows the p_{inj} actually implemented.

value. For convenience, a zoom-in of the relevant information is shown in Figure 5. Note the mechanistic reason for the behavior: the algorithm correctly predicts that even if beam power is kept constant, the temperature will continue increasing, so that the optimal action to maintain the target is to decrease the beam power right away rather than holding it steady.

Phase III

At 2.5s the final phase of the shot begins: scaling the temperature target profile up to a core value of 6keV. Towards the end of the previous phase, an $n = 2$ MHD mode began to grow, and at the beginning of the 6keV phase, the beams are already firing at maximum power, so the proposals to increase cannot be fulfilled. The temperature therefore hovers around the 3keV where it had started.

Analysis

In summary, the finite set controller is able to perform with some of the beneficial characteristics of a stable proportional-derivative controller without any hand-tuning (though more work would be needed to compare the controllers quantitatively). While the phases of the experiment were not long enough to see the full dynamics, we suspect that the controller would have properly settled down close to the target values in the first and second phases. Although the finite set controller is the simplest possible "model-predictive" control algorithm and just a starting point, the benefit of the lookahead to more efficiently achieving a desired state is demonstrated.

Each neural net evaluation takes about a millisecond, and these calculations are done in serial. The DIII-D control system, like most realtime systems, requires that the algorithm run on a fixed cycle time throughout the shot. To ensure it completes at each cycle with room for error, a cycle time of 10ms was used.

4. Conclusion and future work

We have built the infrastructure and done a successful initial test for a new paradigm of fully data-driven, model-based tokamak state control. Moving forward, we plan to use more intelligent control schemes. One idea is to locally linearize the machine learning model that predicts future states so that we have a form $\vec{x}(t+1) = \mathbf{A}\vec{x}(t) + \mathbf{B}\vec{u}(t)$ amenable to traditional control techniques. One way to do this is to get a local Jacobian from the machine learning predictor we already have, in which case the state \vec{x} would be the true state of the system. Another mechanism is to attempt to encode the true plasma state into a latent state where the dynamics are globally linear. In this case, we would be able to overcome the (generally poor) assumption that the real-state dynamics must be linear in a large enough region around the current state to be useful for model-based control. We have also been developing an offline testbed for simulating controllers via the ASTRA transport modeling suite (Pereverzev & Yushmanov 2002). This will give us the means to evaluate many controllers offline and decide the best few to try during real experiments.

We also at present are using a limited set of actuators (total injected power and torque from beams, target density setpoint, and plasma current setpoint). We intend to gradually add actuators to the base infrastructure over time (e.g. individual beam power and possibly voltages + perveances; electron cyclotron gyrotron angles and powers; shape coils; 3D coils; etc). We also intend to add more state information that may improve predictive models and hence controllability, e.g. high-resolution turbulence information from C02 interferometers and Beam Emission Spectroscopy.

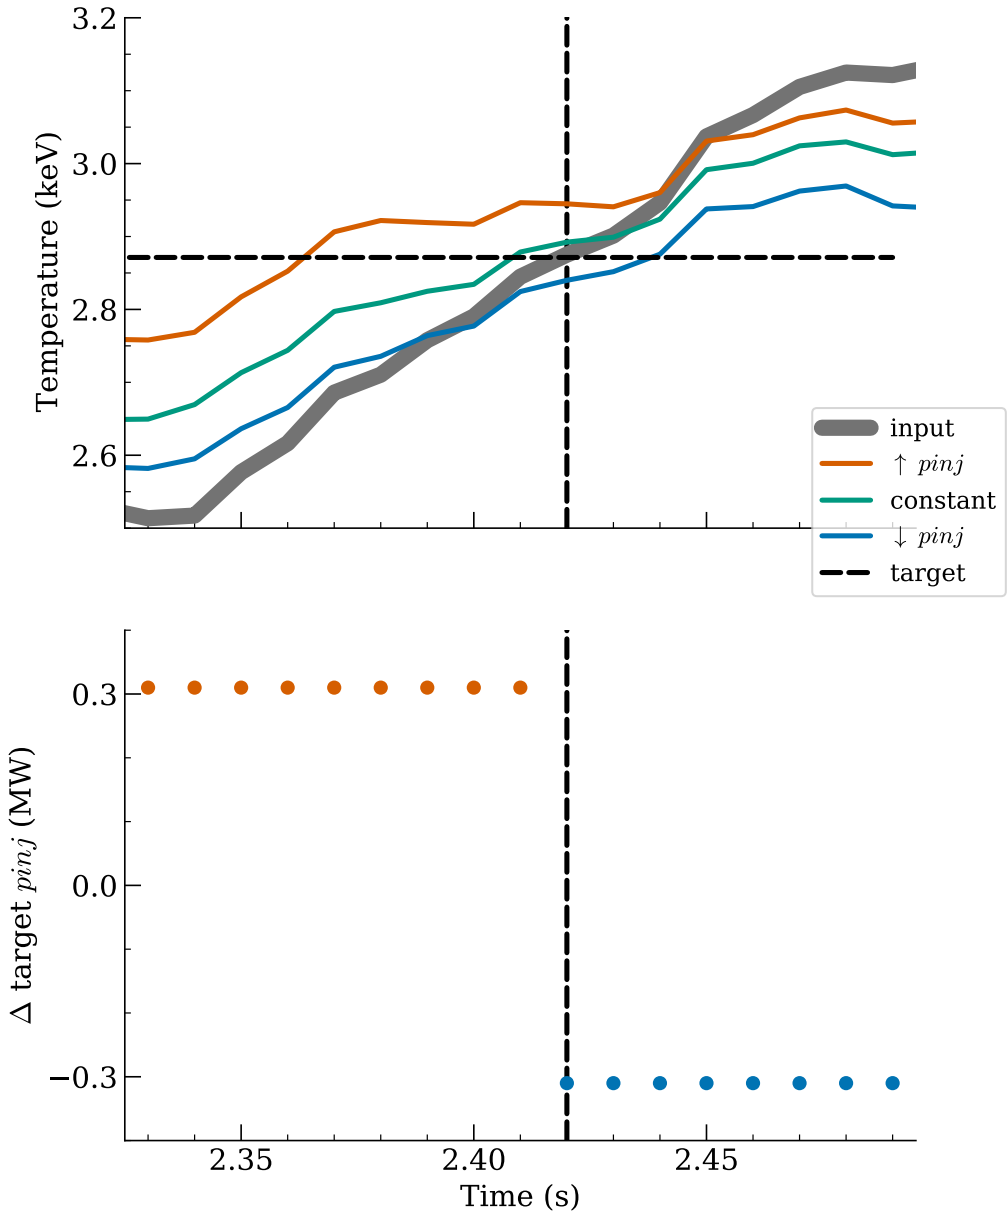


Figure 5. Zoomed display of Phase II in Figure 4, showing the derivative-like action attempting to mitigate overshooting the temperature target by decreasing the injected power right as the target is achieved (at the vertical dashed line).

Thanks to Al Hyatt, Dan Boyer, Jayson Barr, David Eldon, and the rest of the DIII-D operations and control team. Thanks to Josiah Wai for helpful discussions on model-predictive control. 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 Awards DE-AC02-09CH11466, DE-SC0021275, and DE-FC02-04ER54698.

Disclaimer

This report is 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 hereof. 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

- Abbate, J., Conlin, R. & Kolemen, E. 2021 Data-driven profile prediction for DIII-D 61 (4), 046027.
- Barton, Justin E., Shi, Wenyu, Besseghir, Karim, Lister, Jo, Kritz, Arnold, Schuster, Eugenio, Luce, Tim C., Walker, Michael L., Humphreys, David A. & Ferron, John R. 2013 Physics-based control-oriented modeling of the safety factor profile dynamics in high performance tokamak plasmas. In *52nd IEEE Conference on Decision and Control*, pp. 4182–4187. IEEE.
- Boyer, M.D., Erickson, K.G., Grierson, B.A., Pace, D.C., Scoville, J.T., Rauch, J., Crowley, B.J., Ferron, J.R., Haskey, S.R., Humphreys, D.A., Johnson, R., Nazikian, R. & Pawley, C. 2019 Feedback control of stored energy and rotation with variable beam energy and perveance on DIII-d 59 (7), 076004.
- Boyer, Mark D, Barton, Justin, Schuster, Eugenio, Luce, Tim C, Ferron, John R, Walker, Michael L, Humphreys, David A, Penafior, Ben G & Johnson, Robert D 2013 First-principles-driven model-based current profile control for the DIII-d tokamak via LQI optimal control 55 (10), 105007, tex.ids= boyer_first-principles-driven_2013-1.
- Chollet, F 2015 Keras .
- Conlin, Rory, Erickson, Keith, Abbate, Joseph & Kolemen, Egemen 2021 Keras2c: A library for converting keras neural networks to real-time compatible c 100, 104182.
- Cortes, P., Kazmierkowski, M.P., Kennel, R.M., Quevedo, D.E. & Rodriguez, J. 2008 Predictive control in power electronics and drives 55 (12), 4312–4324.
- Felici, F., Citrin, J., Teplukhina, A.A., Redondo, J., Bourdelle, C., Imbeaux, F., Sauter, O., JET Contributors & the EUROfusion MST1 Team 2018 Real-time-capable prediction of temperature and density profiles in a tokamak using RAPTOR and a first-principle-based transport model 58 (9), 096006.
- Felici, F & Sauter, O 2012 Non-linear model-based optimization of actuator trajectories for tokamak plasma profile control 54 (2), 025002.
- Felici, F., Sauter, O., Coda, S., Duval, B.P., Goodman, T.P., Moret, J-M., Paley, J.I. & the TCV Team 2011 Real-time physics-model-based simulation of the current density profile in tokamak plasmas 51 (8), 083052, tex.ids= felici_real-time_2011-1.
- Ferron, J.R, Gohil, P, Greenfield, C.M, Lohr, J, Luce, T.C, Makowski, M.A,

- Mazon, D, Murakami, M, Petty, C.C, Politzer, P.A & Wade, M.R 2006 Feedback control of the safety factor profile evolution during formation of an advanced tokamak discharge 46 (10), L13-L17.
- Fu, Yichen, Eldon, David, Erickson, Keith, Kleijwegt, Kornee, Lupin-jimenez, Leonard, Boyer, Mark D., Eidietis, Nick, Barbour, Nathaniel, Izacard, Olivier & Kolemen, Egemen 2020 Machine learning control for disruption and tearing mode avoidance 27 (22501).
- Fujita, T & team, the JT-60 2006 Steady state operation research in JT60u with extended pulse length 46 (3), S3-S12.
- Groebner, R.J, Baker, D.R, Burrell, K.H, Carlstrom, T.N, Ferron, J.R, Gohil, P, Lao, L.L, Osborne, T.H, Thomas, D.M, West, W.P, Boedo, J.A, Moyer, R.A, McKee, G.R, Deranian, R.D, Doyle, E.J, Rettig, C.L, Rhodes, T.L & Rost, J.C 2001 Progress in quantifying the edge physics of the h mode regime in DIII-d 41 (12), 1789-1802.
- Holcomb, C. T., Heidbrink, W. W., Ferron, J. R., Van Zeeland, M. A., Garofalo, A. M., Solomon, W. M., Gong, X., Mueller, D., Grierson, B., Bass, E. M., Collins, C., Park, J. M., Kim, K., Luce, T. C., Turco, F., Pace, D. C., Ren, Q. & Podesta, M. 2015 Fast-ion transport in $q_{min} > 2$, high-beta steady-state scenarios on DIII-d 22 (5), 055904.
- Humphreys, D., Ambrosino, G., de Vries, P., Felici, F., Kim, S. H., Jackson, G., Kallenbach, A., Kolemen, E., Lister, J., Moreau, D., Pironti, A., Raupp, G., Sauter, O., Schuster, E., Snipes, J., Treutterer, W., Walker, M., Welander, A., Winter, A. & Zabeo, L. 2015 Novel aspects of plasma control in ITER 22 (2), 021806.
- ITER Physics Expert Group on Confinement and Transport 1999 Chapter 2: Plasma confinement and transport 39, 2175, `tex.ids=iter_physics_expert_group_on_confin_transport_chapter_1999`.
- Jalalvand, Azarakhsh, Abbate, Joseph, Conlin, Rory, Verdoolaege, Geert & Kolemen, Egemen 2021 Real-time and adaptive reservoir computing with application to profile prediction in fusion plasma p. 12.
- Laborde, L, Mazon, D, Moreau, D, Murari, A, Felton, R, Zabeo, L, Albanese, R, Ariola, M, Bucalossi, J, Crisanti, F, Baar, M de, Tommasi, G de, Vries, P de, Joffrin, E, Lennholm, M, Litaudon, X, Pironti, A, Tala, T & Tuccillo, A 2005 A model-based technique for integrated real-time profile control in the JET tokamak 47 (1), 155-183.
- Lao, L. L., St John, H. E., Peng, Q., Ferron, J. R., Strait, E. J., Taylor, T. S., Meyer, W. H., Zhang, C. & You, K. I. 2005 MHD equilibrium reconstruction in the DIII-d tokamak 48, 968-977.
- Luxon, J. L. & Davis, L. G. 1985 Big dee - a flexible facility operating near breakeven conditions 8 (1), 441-449.
- Maljaars, E., Felici, F., Blanken, T.C., Galperti, C., Sauter, O., de Baar, M.R., Carpanese, F., Goodman, T.P., Kim, D., Kim, S.H., Kong, M., Mavkov, B., Merle, A., Moret, J.M., Nouailletas, R., Scheffer, M., Teplukhina, A.A., Vu, N.M.T., The EUROfusion MST1-team & The TCV-team 2017 Profile control simulations and experiments on TCV: a controller test environment and results using a model-based predictive controller 57 (12), 126063.
- Maljaars, E., Felici, F., De Baar, M. R., Van Dongen, J., Hogeweij, G. M.D., Geelen, P. J.M. & Steinbuch, M. 2015 Control of the tokamak safety factor profile with time-varying constraints using MPC 55 (2), 23001, publisher: IOP Publishing.
- Montes, K. J., Rea, C., Granetz, R. S., Tinguely, R. A., Eidietis, N., Meneghini, O. M., Chen, D. L., Shen, B., Xiao, B. J., Erickson, K. & Boyer, M. D. 2019 Machine learning for disruption warnings on Alcator CMod, DIII-D, and EAST 59 (9).
- Moreau, D., Mazon, D., Walker, M.L., Ferron, J.R., Burrell, K.H., Flanagan, S.M., Gohil, P., Groebner, R.J., Hyatt, A.W., La Haye, R.J., Lohr, J., Turco, F., Schuster, E., Ou, Y., Xu, C., Takase, Y., Sakamoto, Y., Ide, S., Suzuki, T. & ITPA-IOG group members and experts 2011 Plasma models for real-time control of advanced tokamak scenarios 51 (6), 063009, `tex.ids=moreau_plasma_2011-1`.
- Penaflo, B.G, Ferron, J.R, Walker, M.L, Piglowski, D.A & Johnson, R.D 2001

- Real-time control of DIII-d plasma discharges using a linux alpha computing cluster 56-57, 739–742.
- Pereverzev, G V & Yushmanov, P N 2002 ASTRA automated system for TRansport analysis .
- Rea, C., Montes, K. J., Erickson, K. G., Granetz, R. S. & Tinguely, R. A. 2019 A real-time machine learning-based disruption predictor in DIII-D 59 (9), 096016.
- Sammuli, B.S., Barr, J.L., Eidielis, N.W., Olofsson, K.E.J., Flanagan, S.M., Kostuk, M. & Humphreys, D.A. 2018 TokSearch: A search engine for fusion experimental data 129, 12–15.
- Teplukhina, A A, Sauter, O, Felici, F, Merle, A, Kim, D, the TCV Team, the ASDEX Upgrade Team & the EUROfusion MST1 Team 2017 Simulation of profile evolution from ramp-up to ramp-down and optimization of tokamak plasma termination with the RAPTOR code 59 (12), 124004.
- Walker, Michael L., De Vries, Peter, Felici, Federico & Schuster, Eugenio 2020 Introduction to tokamak plasma control. In *2020 American Control Conference (ACC)*, pp. 2901–2918. IEEE, tex.ids= walker_introduction_2020-1.
- Wand, M. P. 2000 A comparison of regression spline smoothing procedures 15 (4), 443–462.
- Wehner, William, Lauret, Menno, Schuster, Eugenio, Ferron, John R., Holcomb, Chris, Luce, Tim C., Humphreys, David A., Walker, Michael L., Penafior, Ben G. & Johnson, Robert D. 2016 Predictive control of the tokamak q profile to facilitate reproducibility of high-qmin steady-state scenarios at DIII-d. In *2016 IEEE Conference on Control Applications (CCA)*, pp. 629–634. IEEE.
- Wijnands, T, Houtte, D. Van, Martin, G, Litaudon, X & Froissard, P 1997 Feedback control of the current profile on tore supra 37 (6), 777–791.
- Xing, Z.A., Eldon, D., Nelson, A.O., Roelofs, M.A., Eggert, W.J., Izacard, O., Glasser, A.S., Logan, N.C., Meneghini, O., Smith, S.P., Nazikian, R. & Kolemen, E. 2021 CAKE: Consistent automatic kinetic equilibrium reconstruction 163, 112163.