Fusion Engineering and Design 163 (2021) 112163

Contents lists available at ScienceDirect

Fusion Engineering and Design

journal homepage: www.elsevier.com/locate/fusengdes

# CAKE: Consistent Automatic Kinetic Equilibrium reconstruction

Z.A. Xing [a,*], D. Eldon [a,b], A.O. Nelson [a], M.A. Roelofs [d], W.J. Eggert [a], O. Izacard [a], A.S. Glasser [a], N.C. Logan [e], O. Meneghini [b], S.P. Smith [b], R. Nazikian [c], E. Kolemen [a,c,*]

[a] Princeton University, Princeton, NJ 08544, USA
[b] General Atomics, P.O. Box 85608, San Diego, CA 92186-5608, USA
[c] Princeton Plasma Physics Laboratory, P.O. Box 451, Princeton, NJ 08543-0451, USA
[d] Eindhoven University of Technology, Eindhoven 5612 AZ, The Netherlands
[e] Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, CA, 94551, USA

ARTICLE INFO

*Keywords:*
MHD
Stability
Equilibrium reconstruction

ABSTRACT

Magnetohydrodynamic equilibrium reconstructions with kinetic constraints are an essential input of many tokamak stability and transport analysis workflows. The reconstructions themselves also contain useful information about pressure and current density distributions within the plasma which can be used to drive other calculations. While standardized magnetic-only, and partially kinetic equilibrium reconstruction is widely available and used, reconstructions for well diagnosed machines such as DIII-D, that incorporates a large and heterogeneous set of internal diagnostic measurements yet robustly reaching high numerical accuracy have been a time intensive, manual process. The Consistent Automatic Kinetic Equilibrium reconstruction tool (CAKE) has been developed and implemented at DIII-D for producing low error, kinetically constrained reconstructions without human intervention, and to identify strategies for doing so that would be broadly applicable. The techniques employed to allow CAKE to handle data of limited quality and availability, as well as the full automation of the workflow may assist future efforts to form kinetic constraint profiles in real time plasma control applications.

## 1. Introduction

Magnetohydrodynamic (MHD) equilibrium reconstruction forms a starting point of analysis for tokamak transport analysis and stability studies; and provides inputs to plasma control systems for both current and future devices. Stability codes such as DCON [1] and ELITE [2] depend on reliable magnetic equilibria as inputs. Therefore accurate equilibrium reconstruction is highly desired. In particular the reconstruction techniques would ideally be able to robustly produce unique equilibria of good resolution and accurately reflect the measurements of available diagnostics. In practice achieving such results is challenging. Automatic procedures for robustly producing equilibrium reconstructions based on magnetic measurement only, as well as those incorporating some kinetic constraints have been available for sometime [15]. However challenges remain for robustly producing kinetic equilibira that incorporates measurements of internal magnetic field (via motional Stark effect spectroscopy), as well as temperature and densities (via Thomson scattering, and charge exchange spectroscopy for example), while ensuring the result is calculated to a high degree of numerical accuracy on a large grid. Producing these well constrained equilibrium reconstruction often requires the intervention of specialists during the workflow and is typically time consuming and vulnerable to inconsistencies in technique and judgement.

In contrast, the Consistent Automatic Kinetic Equilibria (CAKE) reconstruction workflow presented here is fully automated with no requirement for mid-process inspections or adjustments from the user, yet able to produce equilibrium reconstructions incorporating a comprehensive set of diagnostics with low numerical error. A universal set of rules and tuning constants governs every aspect of the procedure. Such an approach offers no opportunity for cherry-picking and no case-specific subjectivity in outlier identification,[1] while enabling the robust production of low convergence error equilibria needed for stability calculations. The CAKE workflow is able to produce quality kinetic equilibria to support large scale statistical studies and plasma

* Corresponding authors.
*E-mail addresses:* zxing@princeton.edu (Z.A. Xing), ekolemen@princeton.edu (E. Kolemen).

[1] The possibility of some subjective bias entering into the choice of global CAKE settings is of course a concern. Best efforts have been made to provide justification for the choices made in CAKE's default settings, and they are at least consistent even if they are ultimately found to be imperfect. This reduces relative errors.

https://doi.org/10.1016/j.fusengdes.2020.112163
Received 17 July 2020; Received in revised form 11 October 2020; Accepted 30 November 2020
Available online 24 December 2020

experiments.

CAKE has been tested with DIII-D [3,4] data, including comparison to carefully prepared kinetic equilibrium reconstructions using other procedures and shown to produce consistent and reasonable results in many common scenarios. The remainder of this paper is organized as follows: Section 2 describes the problem of kinetic equilibrium reconstruction, required diagnostic inputs, and CAKE's role in relation to larger context. Section 3 describes CAKE's method in gathering, mapping, filtering and sanitizing needed data with a focus on robustness and consistency of method. Section 4 addresses the calculation and application of kinetic constraints, as well as the tuning of the EFIT [5] settings to achieve low numerical error results. Section 5 examines the quality and limitations of CAKE results, discusses generalizable techniques, and provides examples of applications of the work flow. Finally, summary and conclusion is given in Section 6.

## 2. Background

Magnetic equilibrium reconstruction has been a subject of extensive study [6,5,7]. Typical magnetic equilibrium reconstruction seeks a solution to the static Grad-Shafranov equation [8]:

$$\begin{aligned} \Delta^{*}\psi &= -\mu_0 R J_\phi \\ J_\phi(R,\psi) &= R\frac{\mathrm{dp}}{d\psi} + \frac{F}{4\pi^2\mu_0 R}\frac{\mathrm{dF}}{d\psi} \end{aligned} \tag{1}$$

where $p$ is the isotropic pressure and $F$ is the paramagnetic current defined as $F = 2\pi R\frac{B_\phi}{\mu_0}$. If the pressure and paramagnetic current are known sufficient accuracy and resolution, then solving the Grad-Shafranov equation would be trivial. Unfortunately, the inverse problem must be solved, based on boundary magnetic measurements in combination with internal measurements of temperature, density, magnetic pitch angle and so on, variables not directly appearing in the standard equation. In ideal situations, it is more desirable to solve the full (non-static) version of the problem, and account for the non-isotropic pressure. There have been attempts to partially address this with diffusive calculations [9,10], but they have not been robustly applied to large databases. The work presented here similarly restricts itself to the static form of the GS equation above.

Reconstructions performed with internal profile data in addition to edge magnetic data are referred to as kinetic equilibrium reconstructions, and these are needed to reproduce internal structures, both those resulting from the pedestal and boot-strap current, as well resulting from fast ions, transport barriers and current drive [11].

To be useful for MHD stability calculations, it is further desirable that the normalized error (GS error) in the numerical solution to the Grad-Shafranov equation (Eq 1) be low. Additionally, depending on the flexibility and resolution of the numerical parameterization used to perform the reconstruction, it is not always advantageous to add more diagnostic constraints as it may make the problem over determined, and make it difficult to obtain low numerical error. Though ideally the reconstruction would accommodate as much measurement input as robustly as possible.

One approach used by code such as INDENTD and EQNOX parameterizes the toroidal current profile purely numerically (effectively using peicewise-linear basis function) which then necessitates the addition of smoothing constraints to the cost function in addition to $\chi^2$[7]. Alternatively, a set of basis functions can be used to parameterize $J_\phi$, which is used by codes such as EFIT [5] and CLISTE [12]. EFIT in particular parameterizes $J_\phi$ using either spline or polynomial basis function for $P$ and $FF'$ [13], and uses Picard iterations to minimize $\chi^2$ associated with a given set of input constraints. This parameterization trades some of the flexibility afforded by the purely numerical method, in exchange for a reduced sensitivity to the smoothing factor, and theoretically higher achievable resolution due to the analytical basis functions. However, the choice of basis function for EFIT can bias the possible solutions to the inversion. For example, experience on DIII-D shows that polynomial basis functions tends to be robust and reliable especially when sparse internal measurements are available, but spline basis function is typically needed to reproduce sharp features such as bootstrap current, or steep pedestals. However, this means the choices of knot locations can bias the fit just sufficiently that the Picard iteration is not able to reach the desired high numerical accuracy.

Automated workflows for kinetic and partial kinetic equilibrium reconstruction is not a new concept. EFIT can consistently and automatically produce equilibrium reconstructions constrained by magnetic probes only or by a combination of probes and Motion Stark Effect spectroscopy (MSE), and these reconstructions are useful for many applications and for preliminary analysis [13]. The most advanced of these are the technique reported by Sabbagh et al. [14], which uses Thomson Scattering data to constrain internal pressure profile in EFIT. Such reconstructions have been featured extensively in RWM studies done on NSTX and KSTAR [15,16]. Using directly applied pressure constraints in $R$, $Z$ coordinates, and polynomial basis functions, this method sidesteps the problems of mapping and fitting of measured profiles. However, this makes it difficult to incorporate multiple, differently located, pressure-relevant measurements such as TS and CER, or fast ion pressure estimates with significantly differing beam geometry. The choice of a polynomial basis function also makes it a challenge to fit distinct pedestal features commonly seen on H-mode tokamak plasmas.

Due to these limitations, kinetic equilibira are still made 'by hand' at many devices either as the only kinetic equilibria or to supplement those from automatic workflows. In these manual workflows, production of kinetic constraints depends on trained operators to guide the analysis at every step by inspecting profile fits and adjusting settings in order to consistently produce results [17,13,18]. The user manually manages profile fitting of all the key physics quantities, such as: electron temperature $T_e$ and density $n_e$ from Thomson scattering [19,20], ion temperature $T_i$, density $n_i$, and rotation velocity components $v_\phi$, and $v_\theta$ from CER [21], and sometimes radiated power $P_{rad}$ from bolometry [22] ($P_{rad}$ may be used for transport code runs [23] which are used to determine fast ion profiles). Profile analysis is used to ease the combination of diagnostic measurements from different locations and to calculate additional constraint inputs such as fast ion pressure and boot-strap current. Profiles of $T_e$, $n_e$, etc. are passed to a transport code such as ONETWO [24] or TRANSP [23] to calculate fast ion pressure and density profiles as well as current density relaxation. Managing profile analysis includes identification and rejection of outliers and selection of profile fit settings, both of which can be subjective. Differences in profile analysis setup between different analysts are a known source of systematic variation in kinetic EFIT results [25], which can complicate comparisons of results from different sources.

The CAKE automatic workflow aims to be as automatic and robust as the magnetic-only reconstructions or partial kinetic EFITs, but incorporating the large selection of diagnostics that is commonly seen in manual workflows. It should be able to robustly produce equilibrium reconstructions on a large grid (129 by 129 for DIII-D, with cell sizes of 1.33 cm by 2.5 cm, for example) to a high degree of numerical precision (low GS error) needed for reliable stability calculations while also being able to fit features of interest such as a distinct pedestal feature or hollow current profiles where applicable. Further, it should be able to achieve this over a range of typical plasmas with consistent tuning parameters and rules.

CAKE does this by first automatizing both the profile analysis, including ELM filtering, data quality checks and outlier rejection, and profile realignment necessitated by differing measurement locations and misalignment due to 3D effects. CAKE then calculates the internal constraints on pressure and current, their uncertainties, and finally tuns the numerical settings of EFIT, particularly of the basis function. CAKE is implemented within the OMFIT framework [26], allowing it to easily call upon a wide range of codes and utilities. The core of the CAKE

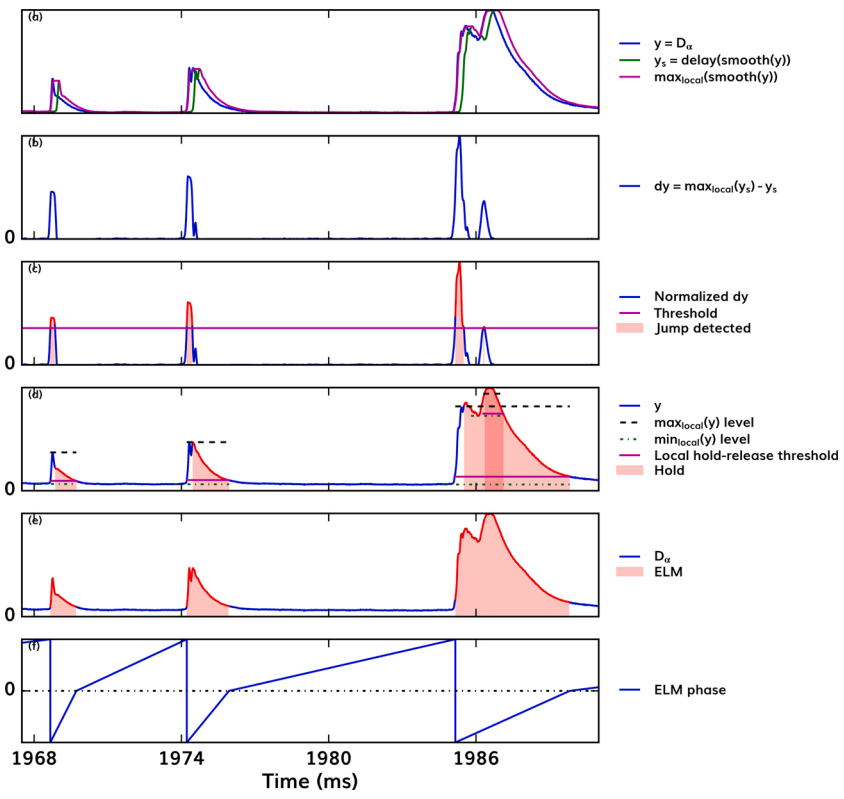


**Fig. 1.** Demonstration of ELM detection. (a) The original signal $y$ is smoothed (short timescale) to get $y_s'$. $max_{local}$ is the local maximum of the last few samples of $y_s'$. $y_s'$ is delayed to get $y_s$. (b) The difference $dy$ between $max_{local}$ and $y_s$ is found. (c) $dy$ is normalized and compared to a threshold to detect jumps. (d) The flag is held until $y$ falls below a threshold. The threshold is defined as a fraction of the difference between the minimum and maximum $y$ from a short window around the jump. (e) Final ELM detection. (f) ELM phase evolves from −1 to 0 during ELMs and from 0 to 1 between ELMs.

workflow is powered by the flexible equilibrium reconstruction code EFIT [5] which is iterated upon to reach the desired goals. To ensure a degree of self consistency, the CAKE workflow by default iterates upon itself by using the output of the previous iteration to re-map diagnostic measurements.

## 3. Data filtering and mapping

CAKE includes standardized profile analysis with robust fail-safes to ensure reliable delivery of useful constraint profiles to EFIT. There are consistent mathematical rules for outlier rejection and for producing profile fits. Additionally, the analysis uses consistent uncertainty propagation rules to generate profile uncertainties which are used as the weights for the constraints input into EFIT. All control parameters are set globally before the workflow begins, which is required for expansion to real time analysis, and prevents cherry-picking and inconsistency. These parameters are listed in Tables A.1, A.3 and A.4 in the appendix.[2]

CAKE uses the standard, automatic, magnetics-only equilibrium reconstruction to perform the initial mapping of diagnostic measurements. From there CAKE uses its own kinetic equilibria to re-map measurements to flux space and perform a second pass. Iterating the mapping in this way is a standard method of improving kinetic equilibrium accuracy and is done because the initial equilibrium used for mapping does not have an accurate flux map, $\psi_N(R, Z)$. Each iteration improves the approximation, but experience on EAST suggest that beyond the second iteration the effect is small [18]. A comparison between iterations is shown in Fig. 6. Further, CAKE, by default, will realign measured profiles by DIII-D's TS system to address issues that arises from 3D effects.

### 3.1. ELM detection and removal

CAKE's first stage includes gathering and pre-processing of data, including the filtering for data quality and edge localized modes (ELMs). ELM events are a known source of noise for many diagnostics, and the measurements taken during an ELM are unreliable. Also, the inter-ELM periods are cyclic in nature, characterized by an ELM phase. In profile reconstruction, especially for equilibrium applications, it is desirable to only use data from a set range of ELM phase to reduce scatter in the data measured. CAKE defaults to removing data measured during the early ELM phase, when the plasma is rapidly recovering, as well as the ELM event itself.

The ELM filter depends on detecting ELMs using an algorithm inspired by the real-time ELM detector used to enable detachment control in ELMing H-mode [27] and explained by Fig. 1. The deuterium Balmer alpha line $D_\alpha$ (656.1 nm) is measured by filterscopes [28], which are photodiodes behind optical band-pass filters and sampled at 20 kHz. Several relevant filterscope signals (selected automatically based on the plasma shape) are summed to increase signal-to-noise ratio of the $D_\alpha$ measurement passed to the ELM detector.[3]

Then, the calculations based on the filterscope measurements are as follows (see Fig. 1): The total $D_\alpha$ measurement $y$ is smoothed to get $y_s'$ and local maximum of this signal from a short window is recorded as $max_{local}$. $y_s'$ is delayed in time by a short interval to get $y_s$. The difference $dy$ between $max_{local}$ and $y_s$ is normalized to its own all-time maximum and compared to a threshold. This detects sudden jumps in $D_\alpha$, which are taken to be the leading edges of ELMs. The ELM flag is held until $y$ drops back to nearly its original value, or until a timer expires, marking the end of the ELM.

[2] CAKE can automatically generate settings tables to track exact setup, such as the tables shown in the appendix: Tables A.1–A.4 were generated by CAKE while loaded with default settings.

[3] Low signal from ELMs is a concern in detached plasmas, where the $D_\alpha$ can actually decrease near the strike point during ELMs for deeply detached plasmas. Using several filterscopes together (not all at the strike point) compensates for dimming of the ELM flash near the strike point in such cases.

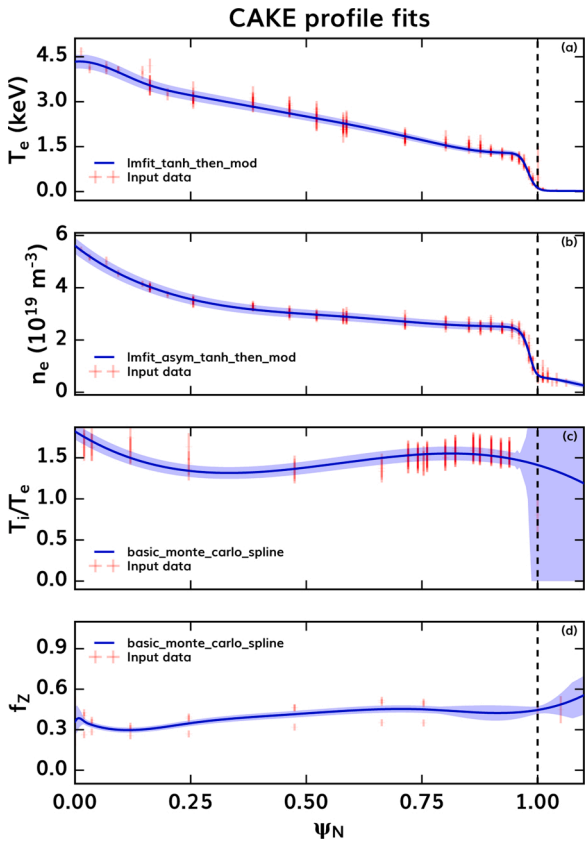


**Fig. 2.** Example of profile fits to (a, b) electron temperature and density from Thomson scattering [19] and (c,d) $\frac{T_i}{T_e}$ and $f_z = \frac{n_C}{n_e}$ from CER [21] (c,d). Extra constraints with very low weight are added to the far scrap-off layer as a safeguard against bad behavior from severely under-constrained fits. They are outside of the plot range.

Using ELM timing data from ELM detection, the ELM filter acts to select a specific part of the ELM cycle. The first criterion is ELM phase, which is defined as −1 at the start of an ELM, increasing to 0 at the end of the ELM, then increasing from 0 to 1 throughout the inter-ELM period, as shown in Fig. 1(f). Long ELM-free periods (such as might occur during ELM suppression [29]) are handled by setting a delay since the last ELM after which all data pass the ELM filter. A minimum end-of-ELM delay must also pass before any data are accepted; this rejects data between double peaks in $D_\alpha$ where ELM phase would rapidly sweep from 0 to 1. Thus, the ELM filter responds appropriately to mixtures of ELMing and non-ELMing phases and is robust to noise that might otherwise appear as a series of short inter-ELM periods. Typical analysis of the pedestal structure would set the acceptable ELM phase range to [0.80, 0.99], which represents the pedestal just before an ELM [25]. For more general stability calculations, a wider phase range are used. The default settings for the ELM filter are given in Table A.1.

### 3.2. Data quality control and filtering

The data quality filter identifies and rejects outliers and unphysical values. CAKE enforces minimum uncertainty values (both absolute and percentage) based on estimates of systematic uncertainty of DIII-D diagnostics. Specific limits can be found in Table A.3. Additionally, CAKE filters diagnostic fitting results (for example, individual $T_i$ measurements via CER) based on their reduced $\chi^2$ (for example, from the Gaussian fit to the spectrometer intensity measurements). These values reflect cases where the diagnostic fitting algorithms have failed to fit the detector observations, and since at DIII-D, reported diagnostic uncertainty is calculated by the co-variance method, high reduced $\chi^2$ would also indicate that the reported uncertainty is unreliable.

### 3.3. Profile fitting and spatial alignment corrections

After ELM and quality filters are applied, data are fit vs. $\psi_N$ using a variety of fit functions. The original $\psi_N$ map used for fitting is taken from a magnetic-only or magnetic-and-MSE-only equilibrium reconstruction calculated by EFIT [5]. A variety of fits have been tested in CAKE, including splines, polynomials, and a modified hyperbolic tangent fit similar to the one presented in Groebner et al. [30]. The default functions used to fit the various quantities are given in Appendix B. An example set of profile fits is shown in Fig. 2. These fits are the basis for computation of fast ion profiles and neoclassical calculations for current profiles [31,32] using various methods.

A complication that occurs in DIII-D is misalignment of the Thomson profiles with the base equilibrium reconstruction and other diagnostics. There is two major causes. First, there are 3D error fields which are not accounted for in the 2D reconstruction. On DIII-D, the equilibrium reconstructions are dominated by a poloidal array of magnetic probes at a different toroidal angle from Thomson. When combined with the presence of 3D perturbations, this makes it difficult for the axisymmetric reconstruction to find the correct separatrix at the Thomson location.

This is a known issue that has been coped with at DIII-D before [33]. Figures in Wilcox et al. [34] show just how much 3D fields (in that case applied intentionally) can affect alignment of profiles relative to 2D equilibrium reconstructions. To compensate for potential misalignment between Thomson and the initial equilibrium mapping, the electron profiles (both $T_e$ and $n_e$ together) are shifted such that the separatrix location inferred from the $T_e$ profile matches the separatrix in the initial mapping. The location of the separatrix can be estimated from the $T_e$ profile by one of two methods: firstly, the lower knee of a hyperbolic tangent fit [30] in the pedestal tends to line up well with the apparent position of the separatrix, or, alternatively, the $T_e$ value at the separatrix can be assumed. The assumption is based on two-point-modelling of the scrape-off-layer which typically returns an estimate near ~80 eV at the separatrix in DIII-D H-mode [35]. Inaccuracy in the assumed value of $T_{e,sep}$ will not cause large error in the position adjustment as long as the $T_e$ gradient is large near the assumed $T_{e,sep}$, which is a reasonable assumption in H-mode. CAKE first attempts to use the tanh lower knee method, but if the required correction estimated by this method is too large or the resulting $T_e$ at the new separatrix is too high, then the constant $T_e$ method is used instead. CAKE does not yet evaluate the two-point-model to make case-by-case estimates of $T_{e,sep}$, instead using a constant value. An example of profile alignment analysis is shown in Fig. 3. Because neither of the separatrix position estimates works well in L-mode (although two-point-modelling could work in L-mode if it were evaluated for L-mode), and because L-mode profiles are less sensitive to the exact position of the separatrix because gradients are lower, no profile alignment is attempted for time-slices which are in L-mode.

There is further the issue of CER alignment, or rather the relative alignment between electron and ion profiles and pedestal location. DIII-D's Thomson scattering diagnostic [19] makes measurements of the pedestal near the top of the machine, rather than the midplane, where most diagnostics (including CER [21]) are located. Thus they are subject to different degrees of misalignment due to both toroidally and poloidally varying uncertainty in the 2D boundary. In cases of a well defined pedestal, the misalignment of electron and carbon profiles cause them to appear to have two distinct pedestal locations that are separated by up to a pedestal width. Aside from being nonphysical, this relative misalignment will generate challenging constraint profiles if not corrected. Since CAKE uses CER measurements of $C^{6+}$ density ($n_C$) and assumes quasi-neutrality to calculate deuterium density, relative misalignment causes bumps and ripples in the calculated thermal pressure near the pedestal, which translates to even more significant ripples in the bootstrap current constraint. In more hands-on workflows, this problem is overcome by hand examination and adjustment of alignment, especially

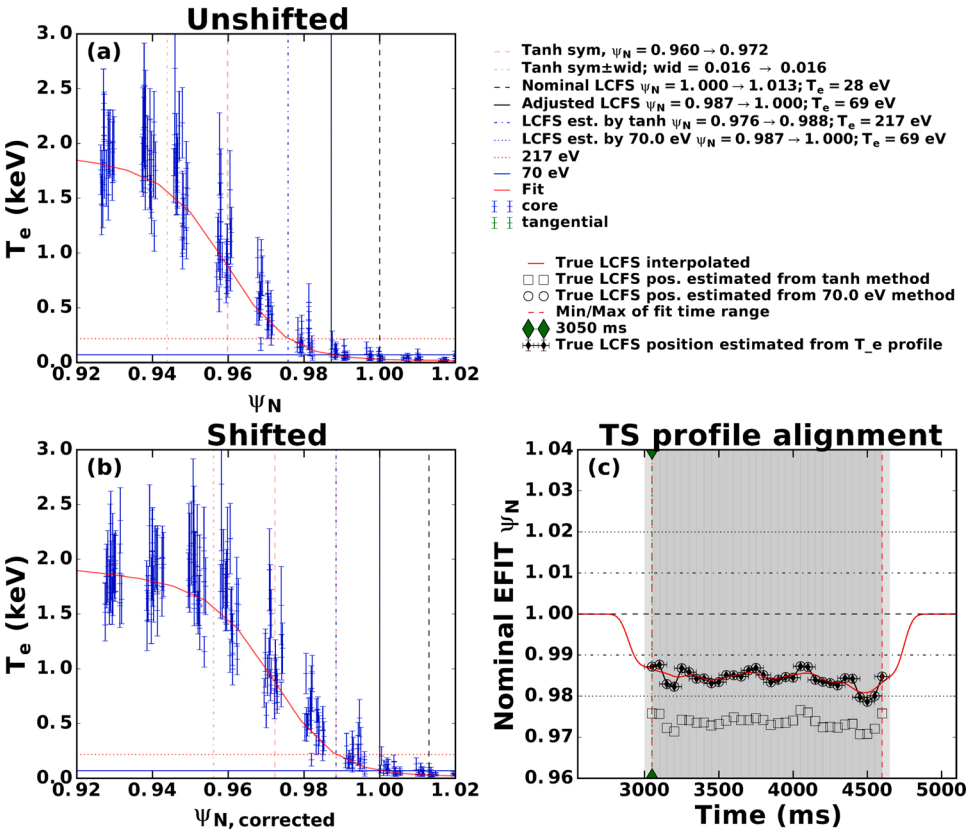


**Fig. 3.** Thomson profile alignment demonstration. (a) $T_e$ vs. $\psi_N$ from the nominal EFIT with no correction and (b) the same data after correcting the $\psi_N$ values. On (a) and (b), vertical lines mark: symmetry point of the hyperbolic tangent fit [30] (red dashed), one half width inside of the symmetry point (red dashed-dotted), one half width outside of the symmetry point (blue dashed-dotted; this is the true LCFS indicated by the tanh method), the original LCFS from the nominal EFIT (black dashed), and the final estimate of the true LCFS from this analysis (black solid). (c) Position of the estimated true LCFS in terms of nominal $\psi_N$ vs. time. Estimates from the tanh method are marked with squares, estimates from the separatrix temperature method with circles, and the selected best answer with black diamonds. The selected answer is smoothed to give the solid red line. The shaded regions indicate time ranges where data are considered when forming kinetic constraints. The dashed-dotted green vertical line with green diamonds marks the time-slice shown in (a,b). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

of CER measurements. In order to achieve a high level of robustness automatically, CAKE assumes that the ratios $\frac{T_i}{T_e}$ and $\frac{n_C}{n_e}$ do not change significantly across the pedestal. This is implemented by fitting $\frac{T_i}{T_e}$ and $\frac{n_C}{n_e}$ as a spline curve instead of the ion quantities directly, and removing the measurement points in the pedestal region from the fit. Due to potential alignment differences, the ratios of measurements in the pedestal region can be nonphysically high/low. By fitting ratios, the workflow automatically enforces the relative alignment of the pedestal, at a cost of potential structural details in the pedestal ion profiles, such as expected in QH mode discharges.

### 3.4. Detection of H-mode for the purpose of guiding profile analysis choices

To determine which time windows to treat as a H-mode pedestal, a inverse scale lengths ratio (ISLR) of electron pressure is considered:

$$\text{ISLR} = \frac{max\left(\left|\frac{\nabla p_e}{p_e}\right|_{0.75\le\psi_N\le1.0}\right)}{\left|\left\langle\frac{\nabla p_e}{p_e}\right|_{0.4\le\psi_N\le0.75}\right\rangle|} \tag{2}$$

with ISLR $\geq$9 identifying H-mode. The numerator of Eq. (2) is simply an attempt to find the peak of $\nabla p_e/p_e$ in the pedestal; the range $0.75 \leq \psi_N \leq 1.0$ is intended to be wide enough to capture even an unusually wide or badly misaligned pedestal. The denominator normalizes the pedestal peak to help determine if the $p_e$ profile is qualitatively shaped like it has a strong pedestal. The average in the denominator does not include the region in which the peak is sought, nor does it include data near the magnetic axis. The neighborhood of the magnetic axis is excluded because $\partial p_e/\partial\psi_N$ does not generally tend towards 0 or an especially meaningful value at the axis. This function operates on profile fits of $p_e$. The selection of the threshold value (9) and the lower bound of the interior averaging region (0.4) was based on comparing L-H and H-L transition times between those determined by (ISLR > threshold) vs. transition times determined by manually checking changes in $D_\alpha$, $\partial\langle n_e\rangle/\partial t$, and other quantities. This method can find L-H and H-L transition times with accuracy of about 30 ms and low sensitivity to exact choice of threshold for a wide range of DIII-D scenarios, but its real value is in classifying which time windows should be treated as if they had a strong pedestal and informing choices of profile alignment and fitting functions. For these applications, the time windows under consideration are often greater than ~30 ms, so its accuracy is adequate.

### 3.5. Coping with incomplete input data sets

There are many ways that an experiment might be missing some kinetic measurements that would be needed for kinetic equilibrium reconstruction. In particular to DIII-D experiments, sometimes the neutral beam sources that enable CER measurements near the magnetic axis are off for extended periods of time, leading to a lack of ion data in the core region even when edge data may be available. There are also times when CER measurements are entirely unavailable for certain time period or entire shots. Because CAKE prioritizes reliability, methods are provided for handling smaller gaps in avaliable datasets with some assumptions.

For the former case of regionally missing or spare CER data, fits can be attempted with the avalible data (in the edge for example). However, the fit through the region with missing data can extrapolate to nonphysical values with nothing to anchor the fit. Thus, fits to CER data are given anchor points at the magnetic axis with very low fit weight such that they have negligible effect when data are present near the axis, but prevent wild extrapolation otherwise. This can be handled by assuming a relationship between thermal ion density and electron density and between thermal ion temperature and electron temperature, as long as Thomson measurement of electron quantities are available. The nominal value of the $T_i/T_e$ anchor generated from assumptions

given in Table A.4. The nominal value of the $n_C/n_e$ anchor is calculated from the $n_C$ (carbon density) measurement closest to the axis, based on the observation that density profiles often have very shallow slopes in the core.

Cases where ion measurements are entirely missing are not handled automatically, avoid mixing equilibrium reconstruction made with and with out ion measurements. However, CAKE can at the scientist's discretion, produce variants where the entire ion temperature and density profiles are estimated by using $T_i/T_e$ assumptions. The uncertainty in $T_e$ is also copied and scaled by $T_i/T_e$, but also enhanced by an additional factor to denote the uncertainty in the assumption. This general concept has also been employed at EAST to compensate for incomplete coverage of $T_i$ [18].

## 4. Robust production of kinetic equilibria

For reliable calculation of MHD stability parameters such as $\delta W$ and $\Delta'$, the magnetic equilibrium used need to reach a high degree of numerical accuracy, more so than other uses for equilibrium reconstructions. A commonly used rule-of-thumb regarding output of the popular EFIT algorithm is that the GS error should reach below $10^{-8}$. The strategy used by CAKE to achieve robust reconstruction of low error equilibria is described below.

### 4.1. Forming constraint profiles

CAKE uses 3 main sets of internal constraints: pressure, edge current, and MSE measurements. The core current profile in DIII-D is well constrained by MSE measurements, and EFIT can internally process and use MSE data to constrain the reconstruction [36]. For the profile constraints that are produced from CAKE's profile analysis, experimental uncertainties are propagated through the calculations to form the input weighting. While CAKE does not attempt to calculate a high accuracy profile of the radial electric field $E_r$, some estimate of $E_r$ is needed to adjust MSE measurements, which are sensitive to $E_r$ [37]. For this purpose, it is reasonable to use estimated $E_r$ calculated from $V_{tor} \times B_{pol}$, as this forms the dominant component of $E_r$ outside of the pedestal, where MSE is most useful. Further, the best choices for edge constraints and basis function, discussed below, depends on if MSE constraints are available. Standard CAKE operation expects the availability of MSE measurements, but incorporates the appropriate alternative settings that allow processing of shots with no MSE data available, at reduced physics quality.

The general equation for the total current is $J_{tot} = J_{Ohm} + J_{BS} + J_{NBCD} + J_{ECCD}$, where $J_{Ohm}$ is the ohmically drive current, $J_{BS}$ is the boot-strap current, $J_{NBCD}$ is the neutral beam current drive, and $J_{ECCD}$ is the electron-cyclotron current drive. For the edge current constraint, CAKE calculates $J_{BS}$ using Sauter's formula [31,32] with the method for extending to multiple ion species given by Koh et al. [38]. Koh's other modifications to the $J_{BS}$ formula are not used because it was shown that the modifications typically produced small changes and were less consistent with simulations of DIII-D and NSTX using the first-principles kinetic code NEO [39]. As part of the CAKE project, a native calculation of $J_{BS}$ and neoclassical conductivity $\sigma_{NC}$ was implemented within the OMFIT [26] framework and found to be consistent with the implementation of the Sauter model in NEO [40]. CAKE estimates the inductively driven current using $J_{Ohm} \propto \sigma_{NC}$ and scales $J_{Ohm}$ to produce the correct total current $I_p$ (The neoclassical conductivity $\sigma_C$ formula is taken from Sauter [31]). In other words, it is assumed that the loop voltage has no spatial variation and the current is fully relaxed. This is obviously mistaken during the rampup phase of the discharge, but probably is not a major source of error once the current profile has had time to relax. Sensitivity to the distribution of $J_{Ohm}$ is reduced as this calculation is only used to produce edge current constraint. Thus, a bad estimate of $J_{Ohm}$ in the core is inconsequential since it is not even passed to EFIT. CAKE does not yet attempt to estimate $J_{NBCD}$ or $J_{ECCD}$, but as long as these terms are small at the edge where CAKE places its constraint points, this will produce only small error (if they are large in the core, then MSE will measure their effects and include them in constraints passed to EFIT). The total current constraint profile calculated by CAKE is shown in Fig. 4, including the points at which the curve is evaluated and passed to EFIT. This edge current constraint is naturally only applied in the edge, but the width of this edge is determined by the width of the pedestal. Different DIII-D operations scenarios have potentially very different pedestal widths. The width of the edge constraint applied by CAKE is $w_{\text{constraint}} = \min(2 * w_{\text{ped}}, 0.3)$ to accommodate this. The width is limited to maximum of 0.3 to prevent the over application of the edge current constraint in cases where the pedestal width determination fails to produce a reasonable value, which occasionally happens when if the pedestal is poorly defined. For L mode reconstructions, the edge current constraint is not as important as it does not have the distinct boot-strap current peak. For these cases, the width of the edge constraint is also set to 0.3, which approximately corresponds to the edge region where there is little MSE coverage on DIII-D.

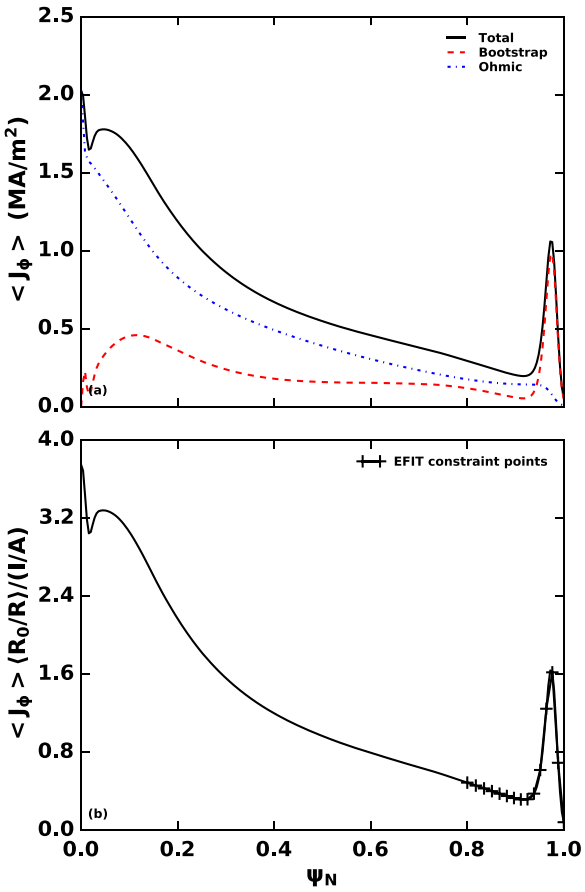


**Fig. 4.** Example current constraint profile. (a) Bootstrap (red dashed), Ohmic (blue dashed-dotted), and total (black solid) current density. (b) Normalized current density for input to EFIT. $R_0$ is the nominal major radius of the vacuum vessel center (1.6955 m for DIII-D), $R$ is the major radius, $I$ is the total current, and $A$ is the cross sectional area of the LCFS. The overall constraint profile is only evaluated at a selected set of points. These are placed mostly at the edge where MSE has trouble measuring and where the amount of un-relaxed Ohmic current is expected to be low. The edge is also important because of an often substantial amount of bootstrap current which EFIT would miss without this constraint. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The pressure constraint is calculated as $p_{tot} = p_e + p_i + p_Z + p_{fast}$, where $p_e$, $p_i$, and $p_Z$ are the electron, ion, and impurity thermal pressures, and $p_{fast}$ is the fast ion pressure. The thermal pressure terms determined by profile fits or filled in with assumptions, fast ion density and pressure

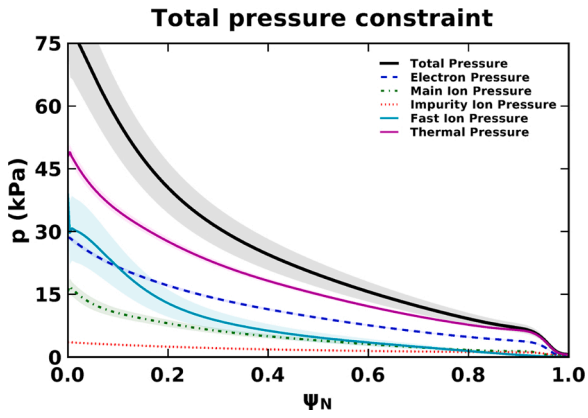


**Fig. 5.** Example of the total pressure constraint profile (solid black line) with a breakdown of the individual components. Total pressure from the thermal particle and fast ion populations are shown with solid magenta and cyan lines. Electron pressure is shown with a dashed blue line, thermal main ion (deuterium) pressure by a dashed-dotted green line, and thermal impurity ion pressure by a dotted red line. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

are calculated by running ONETWO [24] by default, but can alternatively be calculated via TRANSP and NUBEAM [41], or crudely estimated using assumptions in Table A.4. The resulting total pressure constraint is the black curve in Fig. 5, with uncertainty given by the shaded region. An example is shown in Fig. 5.

The set of pressure constraint locations used depends on the availability of MSE constraints. In cases where MSE constraint is available, then the pressure constraint is applied for nearly the entire extent of the plasma (where $\psi_n \geq 0.2$). The constraint is not applied to the very core due to empirical observation that the constraint at $\psi_n = 0$ make it less likely for the result to reach the low GS error goal that we set. This is coupled with the fact that the typical fast ion pressure dominates the very core, where estimation methods used by CAKE (ONETWO, or fraction of electron pressure) are reduced models as compared to more reliable but slower options like NUBEAM [41]. In the cases where MSE constraints are missing, the pressure constraint is further peeled back so that the core pressure is not constrained from kinetic observations. The lack of MSE data results in a lack of effective core current constraints, which makes it difficult to apply core pressure constraints also, as the existence of core pressure constraints without core current constraints often results in nonphysical current profiles.

### 4.2. Consideration for fitting basis function and knot locations

EFIT can use two types of basis function to fit the $FF'$ and $P'$ terms in the Grad-Shafranov equation, polynomial or cubic tension splines [5]. In order to reliably fit the sharp pedestal present in many DIII-D discharges, CAKE uses the spline basis function. There needs to be a concentration of knots in the pedestal region for the same reason. The knot locations for the $P'$ basis function is chosen by the automatic knot location chooser available in OMFIT [26] based on the pressure constraint profile between $0.7 \leq \psi_n \leq 1$. For a specified number of knots, this algorithm finds the knot locations that would allow cubic splines to best fit the reference profile, in this case for the pressure constraint profile. The initial $FF'$ knot placement is made with the knowledge and expectation of the pedestal in mind, as the auto knot derived from the edge current profile is not always adequate or robust for actual equilibrium reconstruction due to the fact that the current profile is but one of many constraints that the $FF'$ function needs to fit. Instead the $FF'$ knots are chosen based on physics heuristics. Two knots are used to accommodate the bootstrap current constraint by locating one knot at the top of the current constraint peak, and one near the turning point where the bootstrap current peak ends. Two more knots are placed on the "shoulder" of the pedestal, defined in

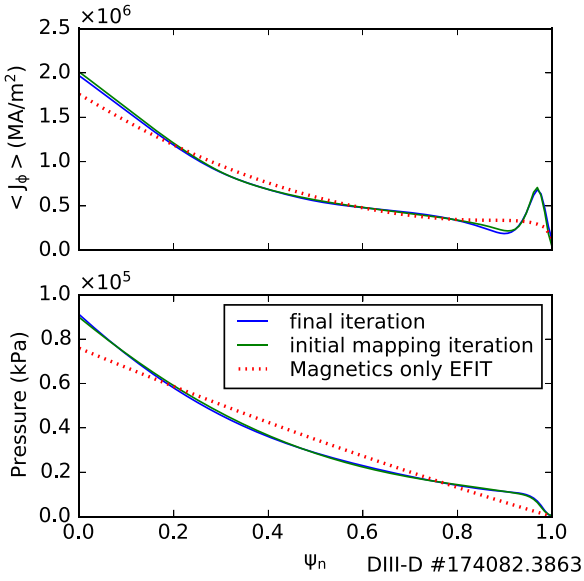


**Fig. 6.** CAKE final iteration (blue) vs. mapping iteration (green) and magnetics only equilibrium (dotted). Results from mapping iteration is used to remap the diagnostic measurements for the final result. The different between the iterations is small compared to the improvement over the magnetics only reconstruction. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

CAKE as a full pedestal width inside of the pedestal, in such a way the knots evenly divide the "shoulder" width. Two final knots are placed in the core, according to the FF' curvature determined by the reference EFIT equilibria produced at DIII-D that prioritizes MSE fitting (i.e. EFIT02).

Unfortunately, EFIT convergence to low GS error levels are often impacted by numerical limitations in addition to physics considerations. From experience, the precise knot location can often have a significant effect on the GS level reached. To address this, CAKE uses a non linear minimizer that makes fine adjustment to the knot locations with in a narrow window ($\Delta\psi_n = 0.02$). The cost function being minimized is a composite of reduced $\chi^2$, the GS error, and the $\beta_p + L_i$ criteria, defined by:

$$C = \max(\chi_r^2, 1) \times (1 + 10\log_{10}\left(\frac{\text{GSerror}}{\text{GSgoal}}\right)) \times (1 + 400\left(\frac{(\beta_p + \frac{l_i}{2})_{\text{CAKE}} - (\beta_p + \frac{l_i}{2})_{\text{EFIT01}}}{(\beta_p + \frac{l_i}{2})_{\text{EFIT01}}}\right) \quad (3)$$

where GS goal is set to $10^{-8}$ by default. The GS error term is scaled by 10 to encourage progress towards the goal. The $\beta_p + L_i$ requirement is from observations that this value can be very reliably calculated from external measurements [8] and if CAKE's output result in a deviation from that of existing equilibrium reconstructions (in this case the EFIT01 equilibrium standard at DIII-D) it would indicate catastrophic failure on CAKE's part. As such this term is weighted heavily. All contributing terms are limited to a minimum value of 1 so that the minimizer does not reward them beyond their goal.

### 4.3. Further GS error improvements

The EFIT code provides a mode of operating, called equilibrium mode, where the equilibrium can be resolved with fixed boundary. These solutions are made by resolving the Grad-Shafranov equation in $\psi$ coordinates based on the direct solution. The process is often found to improve the numerical accuracy of the result and CAKE has incorporated

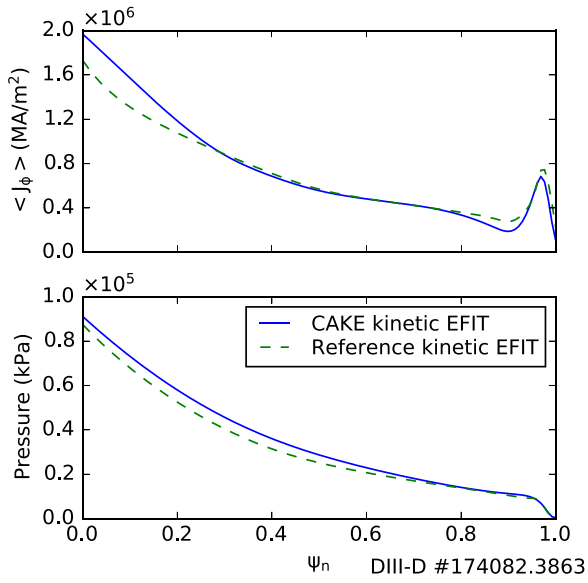


**Fig. 7.** Sample kinetic EFIT result for DIII-D#174082. The reference kinetic EFIT is a high quality kinetic EFIT made using T. Osborne's procedures [25].

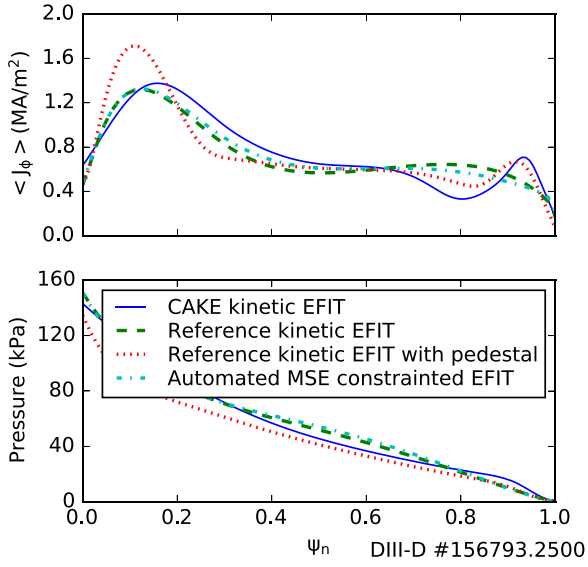


**Fig. 8.** Sample kinetic EFIT results for DIII-D#156793. The shot was part of reversed q profile experiment which are especially challenging. The kinetic EFIT produced by CAKE is shown in solid. The reference kinetic EFIT plotted in green dashed line is one of series hand made kinetic equilibria that was used in stability analysis. Fig. 10 presents a comparison between calculations using this series of equilibria and CAKE equilibria. The red dotted profiles are from a manual reconstruction with emphasis on the pedestal shape. This also illustrates the inconsistencies that some times arise in difficult to fit scenarios due to choices in parameters, which CAKE aims to alleviate. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

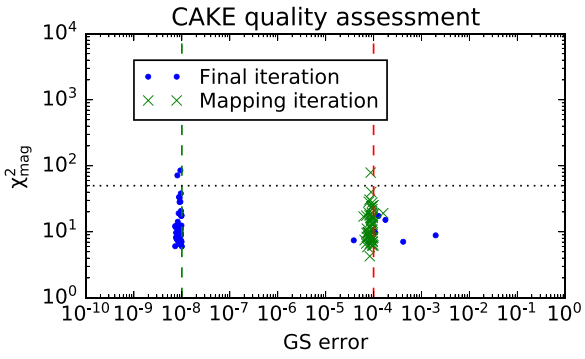


**Fig. 9.** CAKE quality assessment. Each blue dot represent the output from a time slice plotted according to the GS error and magnetic $\chi^2$. The green crosses represent the result of the mapping iteration, which have reduced requirements on GS error. The data presented represents 55 time slices that have MSE data from a cross section of 9 different shots representing a range of H mode operation on DIII-D. >80% of time slices here achieved the GS error goal with good $\chi^2$. If MSE constraints are relaxed, >95% of time slices reach the GS goal. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

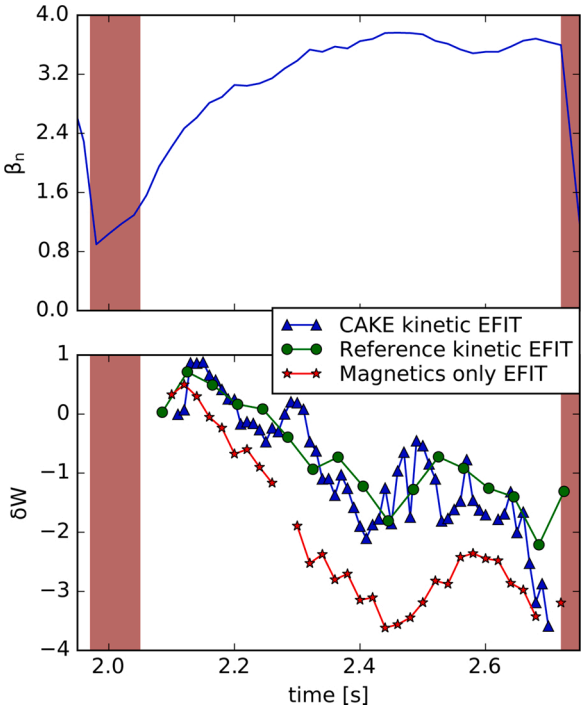


**Fig. 10.** DCON calculation of $\delta W$ using CAKE results as inputs. The red shaded times represent the observations of major instabilities (likely RWMs). $\delta W$ values made using the reference kinetic equilibria (see Fig. 8), as well as those from automatic, magnetic only equilibria, are plotted in comparison. CAKE results are found to replicate the trend obtained by the reference decreasing as the plasma recovers from the previous instability driven relaxation. The missing points from the magnetics-only EFIT results indicates poor convergence for one of the times. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

**Table A.1**
Input filtering settings for CAKE.

| Setting | Default value | Description |
|---|---|---|
| **ELM processing** | | |
| *ELM detection* | | |
| $F_{smo}$ | butter_smooth | Smoothing function |
| $\tau_1$ | 0.50 ms | Timescale for mild smooth |
| $F_{\tau 2}$ | 5.00 | Factor for determining heavy smooth ($F_{\tau 2} = \tau_2/\tau_1$) |
| $F_{\tau d1}$ | 0.00 | Factor for determining mild smooth for derivative from $\tau_1$ |
| $F_{\tau d2}$ | 4.00 | Factor for determining heavy smooth for derivative from $\tau_1$ |
| $DS_{lim}$ | 0.05 | Threshold on difference of smooths normalized to max. |
| $DD^{+}_{lim}$ | 0.14 | Threshold on difference of smooth derivs. when $DD > 0$ |
| $DD^{-}_{lim}$ | 0.14 | Threshold on difference of smooth derivs. when $DD < 0$ |
| $f_{DD,lim}$ | 1.00 | Multiplier for $DD^{\pm}_{lim}$ when $DS < 0$ |
| $b$ | 7.00 samples | Averaging neighborhood width for debouncing |
| *ELM filtering* | | |
| $\phi_{ELM,min}$ | 0.50 | Minimum ELM phase |
| $\phi_{ELM,max}$ | 1.00 | Maximum ELM phase |
| $t_{ELM,accept}$ | 25.00 ms | Delay after ELM ends before accepting data regardless of $\phi$ |
| $t_{ELM,reject}$ | 0.10 ms | Period after ELM ends when data are rejected regardless of $\phi$ |
| **Initial data sanitization** | | |
| *Thomson scattering* | | |
| $\chi^2_{red,max}$ | 8.00 | Max. acceptable reduced $\chi^2$ |
| $(\sigma_{Te}/T_e)_{max}$ | 0.50 | Max. acceptable fractional uncertainty in $T_e$ |
| $(\sigma_{Te}/T_e)_{cold,max}$ | 0.95 | Max. acceptable fractional uncertainty in $T_e$ at low $T_e$ |
| $T_{e,cold}$ | 50.00 eV | Threshold for using cold limit on $\sigma_{Te}/T_e$ |
| $(\sigma_{ne}/n_e)_{max}$ | 0.50 | Max. acceptable fractional uncertainty in $n_e$ |
| *CER* | | |
| $(\sigma_{Ti}/T_i)_{max}$ | 0.11 | Max. acceptable fractional uncertainty in $T_i$ |
| $(\sigma_{nC}/n_C)_{max}$ | 0.25 | Max. acceptable fractional uncertainty in $n_C$ |

**Table A.2**
Default fitting parameters for CAKE.

| Setting | Default value | Description |
|---|---|---|
| **Profile fitting** | | |
| *Fitting functions* | | |
| $F_{Te}$ | lmfit_tanh_then_mod | Function to use for fitting $T_e$ |
| $F_{ne}$ | lmfit_asym_tanh_then_mod | Function to use for fitting $n_e$ |
| $F_{Ti}$ | basic_monte_carlo_spline | Function to use for fitting $T_i$ |
| $F_{fz}$ | basic_monte_carlo_spline | Function to use for fitting $f_z$ |
| *Sanitization of uncertainties in fit inputs* | | |
| $\sigma_{min,Te,fit}$ | 1.00 eV | Min. uncertainty for input when fitting $T_e$ |
| $\sigma_{min,ne,fit}$ | $1\text{E}{+}18\ \text{m}^{-3}$ | Min. uncertainty for input when fitting $n_e$ |
| $\sigma_{min,fz,fit}$ | 0.03 | Min. uncertainty for input when fitting $f_z$ |
| *Miscellaneous profile fitting settings* | | |
| $S_{ne}$ | 1E-19 | Scaling factor for $n_e$ measurements before fitting |
| $S_{other}$ | 1.00 | Scaling factor for other measurements |

this as an additional strategy to reach the GS error goal. Additionally, EFIT provides for a back-averaging parameter that controls the degree that the solution at each numerical step is averaged with the previous, in order to avoid EFIT's numerical solver getting stuck in a numerical loop. In particular difficult cases, CAKE will further adjust this back-averaging parameter in order to further improve the results. These two methods used to improve the numerical accuracy of the result does not visibly change the output profile to a discernible extent. However, these techniques, especially in combination with the nonlinear optimization of knot locations mentioned above, typically cause CAKE run time to increase by a factor of 10–15. Hence, an option is provided to skip such optimization in cases where high numerical quality is not needed.

## 5. Results and discussion

CAKE reliably produces reasonable kinetic equilibrium reconstructions, as demonstrated in Fig. 8 and in 7 . The figure shows that CAKE has reproduced structure in the pedestal which is commonly missed by magnetics-only and magnetics + MSE reconstructions. Comparison of constraint profiles and outputs between CAKE and other workflows shows that CAKE produces reasonable estimates in most cases. Compared to similar reconstructions performed manually, CAKE's main advantage is the ability to process large numbers of cases automatically while applying exactly the same rigorously defined treatment. This level of consistency cannot be produced when a user's intuition is relied upon to complete the analysis.

To improve profile fitting and constraint calculations, CAKE iterates its equilibria process by using the equilibrium from the previous iteration to remap the diagnostic to $\psi_n$ coordinates. In theory, additional iterations of this remapping process will progressively improve the equilibria by improving the fit and mapping of data. But recent experience on EAST indicates that 2 iterations is sufficient [18], and this is implemented in CAKE as default. Since only the last iteration produces the final equilibria to be used for further analysis, the GS error requirements for the earlier (mapping) iterations can be relaxed for speed. For a comparison of the profiles from the iteration process see Fig. 6.

CAKE can reach GS error goal for most of the timeslices produced, allowing CAKE equilibria to be used for MHD stability calculations. About 80% of time slices produced with MSE constraints reaches the goal GS error of $1 \times 10^{-8}$ (see Fig. 9), while if MSE constraints are ignored, 95% will reach the goal. Comparatively speaking, the application of MSE constraints reduces the rate, likely due to the increased fitting criteria, but their inclusion is important for knowledge of internal current profiles. Fortunately, as shown below in Fig. 10, CAKE results produces predominantly good MHD calculations at the quality level achieved. Without MSE data available, CAKE will tend to produce relatively peaked current profiles that may not reflect reality. However, expansion of the edge current constraint to include additional current drive mechanism would improve CAKE's performance in this respect in the future.

### 5.1. Applications of CAKE

One of the primary goals of CAKE is to produce kinetic equilibria that can be used for MHD stability calculations. This is demonstrated by running the DCON code on the outputs produced [1]. In particular, CAKE was used to produce kinetic equilibria for every 10 ms, and the result is used to calculate the time-varying ideal MHD stability parameter $\delta W$ with DCON. The result is compared to that calculated with reference kinetic EFITs in Fig. 10 and found to correspond well. This process involved minimal intervention, taking advantage of the OMFIT framework [26] and its GPEC module. It succesfully replicates the results from manually made equilibria showing the drop in stability prior to the minor disruption, making it a promising tool for database studies or to guide real-time implementations.

The CAKE workflow performs automatic profile analysis as part to support constraint calculations (see Fig. 2). However, this also means CAKE produces plasma profiles of $T_e$, $T_i$, $n_e$, $n_i$, and $V_{tor}$ mapped to the equilibrium produced. These can be used for further analysis. Initial tests have been conducted on automatic workflows that links CAKE as presented here to analysis codes including TRANSP [23] and UEDGE [42], though the OMFIT [26] interface. Efforts are also ongoing to

**Table A.3**
Default tuning for CAKE.

| Setting | Default value | Description |
|---|---|---|
| **Constraint setup** | | |
| *Pressure constraint* | | |
| $\sigma_{min,pe,frac}$ | 0.02 | Min. fractional uncertainty in $p_e$ term in constraint |
| $\sigma_{min,pe}$ | 100.00 Pa | Min. absolute uncertainty in $p_e$ term in constraint |
| $n_{pc}$ | 80 | Number of $p$ constraint points |
| $\langle\psi_{pc,MSE}\rangle$ | 0.80 | Average $\psi_N$ position of $p$ constraint |
| $\psi_{pc,min,MSE}$ | 0.20 | $\psi_N$ position of innermost $p$ constraint |
| $\psi_{pc,max,MSE}$ | 1.00 | $\psi_N$ position of outermost $p$ constraint |
| $n_{pc}$ | 50 | Number of $p$ constraint points, no MSE |
| $\langle\psi_{pc}\rangle$ | 0.83 | Average $\psi_N$ position of $p$ constraint points, no MSE |
| $\psi_{pc,min}$ | 0.50 | $\psi_N$ position of innermost $p$ constraint, no MSE |
| $\psi_{pc,max}$ | 1.00 | $\psi_N$ position of outermost $p$ constraint, no MSE |
| *Current constraint* | | |
| $c_{j,\mathrm{const}}$ | 2 | Ratio of edge $J$ constraint width to $w_{\mathrm{ped}}$ |
| $w_{j,\max}$ | 0.3 | Maximum allowed edge current constraint width |
| *Miscellaneous constraint settings* | | |
| $\sigma_{OT,frac}$ | 0.20 | Fractional error, fast ion results from ONETWO |
| *GS error minimization settings* | | |
| $GS_{\mathrm{goal},i}$ | 0.00 | GS error goal, initial iteration |
| $GS_{\mathrm{goal},f}$ | 1E−08 | GS error goal, final iteration |
| $\Delta\psi_n^{\mathrm{knot}}$ | 0.04 | Allowed window for knot optimization |
| $n_{\mathrm{maxtrials}}$ | 40.00 | Maximum number of knot optimization iterations |

**Table A.4**
Assumptions setup for CAKE.

| Setting | Default value | Description |
|---|---|---|
| **Assumptions for filling in incomplete datasets** | | |
| *Assumptions for estimating thermal ions* | | |
| $T_i/T_e$ | 1.5 | Ratio of ion to electron temperature |
| $n_{e,C}/n_e$ | 0.2 | Fraction of electron density due to carbon |
| $C_{\sigma,ion}$ | 4.0 | Uncert. multiplier for est. therm. ions |
| $C_{\mathrm{edge},p_i}$ | 1.5 | Multiplier for $p_i$ at edge |
| $\psi_{N,\mathrm{edge},p_i}$ | 0.93 | Definition of "edge" for $C_{\mathrm{edge},p_i}$ |
| *Assumptions for estimating fast ions* | | |
| $T_{fast}$ | 75.0 keV | Fast ion temperature |
| $\sigma_{T_{\mathrm{fast}}}$ | 6.0 keV | Uncertainty in fast ion temperature |
| $p_{fast}/p_e$ | 0.3 | Fast ion pressure as fraction of $p_e$ |
| $\sigma_{p_{\mathrm{fast}}}/p_e$ | 0.3 | Fast ion pressure uncert. as fraction of $p_e$ |
| $C_{fast,NBI}$ | 0.1 | $p_{fast}$ enhancement factor for NBI |

expand profile analysis abilities through the use of OMFITprofiles [17], as well as applying the results to pedestal stability analysis via ELITE [2,43]. Future work is aimed bring these capabilities together into an automatic plasma analysis suite.

The prime candidate for CAKE-based kinetic reconstructions for stability analysis are those with MSE data available, due to the importance of core current and $q$ profile to the proper calculation of MHD stability. However, above mentioned contingencies for cases where MSE is unavailable are made for the purpose of a robust fall back that is still able to resolve the edge and pedestal well. In the case where MSE measurements are missing, the equilibira produced focuses on the reconstruction of the pedestal region, and the core pressure and toroidal current profiles are peaked and otherwise featureless. This could be sufficient for certain edge focused diagnostic mapping or analysis, but in general MSE data are recommended for high quality outputs. It is assumed that experimentalists who are interested in good core reconstructions will design the experiment around the availability of MSE measurements. However, for planned database analysis and the establishment of a consistent reference point, two standard equilibrium sets are being run for the current experimental campaign at DIII-D, those with and without MSE constraint applied.

The automation work done for CAKE also provides guidance for ongoing efforts to produce real time kinetic equilibrium. Ongoing effort to apply these strategies to real time equilibrium reconstruction will provide inputs to the STRIDE [44,45] code that performs real-time MHD stability calculations. For example, recent work has used machine learning strategies for physics facilitated tearing prediction [46] aimed at control purposes, and CAKE is well suited to for generating the needed database of kinetic equilibria for such machine learning efforts.

## 6. Conclusion

The Consistent Automatic Kinetic Equilibria (CAKE) workflow produces useful reconstructions of DIII-D [4,3] plasma pressure and current distributions without assuming any input or guidance from the user other than shot numbers and timing requests. It is able to robustly produce kinetic equilibria constrained by TS, CER, and MSE measurements as well as boot-strap current and fast ion calculations, and do so to a high numerical accuracy. Automatic profile analysis typically produces reasonable pressure and current constraints with default settings, supporting the goal of consistent treatment while also making the code easier and faster to use. The primary uses of these results are envisioned as being (1) input to statistical analysis of many plasma shots and timeslices, (2) proof of concept of a fully automated procedure that could be applied to real time and used in a plasma control system with a view toward disruption avoidance, (3) convenient access to kinetic equilibrium reconstructions for analyses that are not sensitive to the errors or uncertainties introduced by assumptions used to make CAKE reliable, for machine learning training sets for example, and (4) first iteration input to a manual kinetic reconstruction workflow. CAKE results are compared against DIII-D reconstructions produced by various manual workflows and found to be mostly consistent. Further, good ideal MHD stability calculations were demonstrated. CAKE will be upgraded to further improve the handling of sparse data, the generation of fast ion profiles for constraint calculations.

## Disclaimer



## CREDIT statement

**Z.A. Xing**: Writing – Original Draft, Writing – Review & Editing, Data Curation, Visualization, Investigation, Software

**D. Eldon**: Writing – Original Draft, Visualization, Data Curation, Investigation, Software

**A.O. Nelson**: Software

**M.A. Roelofs**: Software

**W.J. Eggert**: Software

**O. Izacard**: Software

**A.S. Glasser**: Software

**N.C. Logan**: Software, Data Curation

**O. Meneghini**: Software

**S.P. Smith**: Software

**R. Nazikian**: Conceptualization, Supervision

**E. Kolemen**: Conceptualization, Project administration, Funding acquisition, Supervision

## Conflict of interest

The authors declare no conflict of interest.

## Declaration of Competing Interest

The authors report no declarations of interest.

## Acknowledgement

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 Nos. DE-FC02-04ER54698, DE-FOA-0001386, DE-SC0015878, and DE-AC02-09CH11466.

Development of this project was done in the OMFIT scientific modeling framework [26]. Profile analysis depends on the efforts of several diagnostic groups at DIII-D, Thomson scattering, Charge Exchange Recombination spectroscopy, Motional Spark Effect polarimetry, and of course magnetics.

The authors are grateful to Steve Sabbagh, Tom Osborne, Richard Groebner, Brian Grierson, and the OMFITprofiles team for discussions on general kinetic EFIT calculations and workflows.

## Appendix A. Tuning of settings in CAKE

Proper functioning of the CAKE module requires several settings to be tuned well. The user is not expected to do this as the default settings and logic should provide reasonable performance for all cases. The ability to adjust settings is provided for developers to optimize the initial default tuning and test potential improvements later. These settings include thresholds for outlier rejection, factors for increasing uncertainty when missing data have been replaced with estimates, and the number of knots in the EFIT stream functions ($p'$ and $FF'$). Key tuning settings are listed in Tables A.1–A.4. In order to allow easier inspection of CAKE setup and to ensure that the documented settings stay up to date, these tables of settings and assumptions are generated by the CAKE module as LaTeX files from internal setting variables.

## Appendix B. Profile fit functions for CAKE

CAKE uses several fits based on the LMFIT package in Python as well as wrappers around spline fits from SciPy. For the SciPy splines, error analysis accomplished via Monte Carlo trials. For the LMFIT-based fits, functions are defined in CAKE for use in minimization. For the key functions, partial derivatives are computed analytically and used to find uncertainty in the fit according to $\Sigma^f = \vec{J}\overleftrightarrow{\Sigma}^x\vec{J}^T$, where $\vec{J}$ is the Jacobian and $\overleftrightarrow{\Sigma}^x$ is the covariance matrix returned by LMFIT. CAKE also has the ability to calculate partial derivatives numerically and these are used for checking analytic partial derivatives and for testing new functions which do not have analytic partial derivatives available yet.

### *B.1 Modified hyperbolic tangent fit*

The modified tanh fit model used in CAKE (given by Eq. (B.1)) is based on the MTANH function defined by Groebner et al. [30], but is recast in terms of pedestal height $h$ and SOL offset $f$ and it allows for higher order polynomial $p_i$ terms in the modification.

$$y = f + \frac{h-f}{2}\left[\left(1 - tanh\left(\frac{x-s}{w}\right)\right) + \sum_i p_i Z^i\right] \tag{B.1}$$

where

$$Z = \frac{ze^{-z}}{e^z + e^{-z}}, \quad z = \frac{x-s}{w} \tag{B.2}$$

and $s$ and $w$ are the tanh symmetry point and half-width respectively.

As an alternative to the polynomial modification in the core, the edge ($\psi_N > 0.8$) can be fit to a tanh ($p_i = 0$) and the residual then fit with a spline. Error analysis for the spline fit is done by Monte Carlo trials. These fit methods are designated within CAKE as `lmfit_mtanh2` with the polynomial modification and `lmfit_tanh_then_mod` with the spline to the residual.

### *B.2 Asymmetric tanh fit*

The asymmetric tanh fit is implemented to handle cases where the upper knee of the pedestal is sharper than the lower knee of the pedestal, as in shot 158,103 at 2300 ms (see figure). The model is given by Eq. (B.3):

$$y = \begin{cases} \frac{w_2}{w_1+w_2}\left(1 - \frac{w_1}{w_2}tanh(\frac{x-s}{w_1})\right)\left(h-f\right)+f & x \le s \\ \frac{w_2}{w_1+w_2}\left(1 - tanh(\frac{x-s}{w_2})\right)\left(h-f\right)+f & x > s \end{cases} \tag{B.3}$$

where $x$ is the independent variable (normally $\psi_N$ in this application), $h$ is the pedestal top value or pedestal height, $f$ is the offset or SOL value, $s$ is the symmetry point of both tanh functions as well as the changeover point, and $w_1$ and $w_2$ are the half-widths of the two tanh functions. The function is continuous in its value and first derivative at the changeover point $s$ and the fit parameters $p$ and $f$ retain the same physical meaning as in other tanh fits

such as defined by Groebner et al. [30]. The total width $w_1 + w_2$ should have the same meaning as the total width ($2w$) in standard tanh fits. The symmetry point may have a different interpretation in this function.

This function can be modified to include a polynomial in the core by adding in $y_m$ from Eq. (B.4), but because the amplitude of the inner tanh term now depends on width, the function becomes significantly nonlinear in the fit parameters and linear error propagation fails when the modification is added.[4] Therefore, this fit is used on only the edge region ($\psi_N > 0.8$) and the residual of the fit function evaluated over the entire profile is then fit with a cubic spline. Uncertainty in the spline fit to the residual is determined by the Monte Carlo method. The final result is the sum of the asymmetric tanh and the spline.

$$y_m = (h-f)\sum_i p_i Z^i$$

$$Z = \frac{z_1 e^{-z_1}}{e^{z_1} + e^{-z_1}}$$

$$z_1 = \frac{x-s}{w_1}$$

$Z$ becomes small at $x > s$, so this modification is only significant in the core.

Defining $y_1$ and $y_2$ as the top and bottom cases of Eq. (B.3), $z_i = (x-s)/w_i$, $T_i = \tanh(z_i)$ and $S_i^2 = \mathrm{sech}(z_i)^2$, the partial derivatives of $y_1$ and $y_2$ are:

$$\frac{\partial y_1}{\partial h} = \frac{w_2}{w_1+w_2}\left(1 - \frac{w_1}{w_2}T_1\right)$$

$$\frac{\partial y_1}{\partial s} = (h-f)\frac{S_1^2}{w_1+w_2}$$

$$\frac{\partial y_1}{\partial w_1} = (h-f)\left(\frac{z_1 S_1^2 - T_1}{w_1+w_2} + \frac{w_1 T_1 - w_2}{(w_1+w_2)^2}\right)$$

$$\frac{\partial y_1}{\partial w_2} = (h-f)\left(1 - \frac{w_2 - w_1 T_1}{w_1+w_2}\right)\frac{1}{w_1+w_2}$$

$$\frac{\partial y_1}{\partial f} = 1 - \frac{\partial y_1}{\partial h}$$

$$\frac{\partial y_2}{\partial h} = \frac{w_2}{w_1+w_2}(1-T_2)$$

$$\frac{\partial y_2}{\partial s} = (h-f)\frac{S_2^2}{w_1+w_2}$$

$$\frac{\partial y_2}{\partial w_1} = (h-f)\frac{-w_2}{(w_1+w_2)^2}(1-T_2)$$

$$\frac{\partial y_2}{\partial w_2} = (h-f)\left(\left(1-T_2\right)\frac{w_1}{w_1+w_2} + z_2 S_2^2\right)\frac{1}{w_1+w_2}$$

$$\frac{\partial y_2}{\partial f} = 1 - \frac{\partial y_2}{\partial h}$$

The partial derivatives of the modification $y_m$ are simply added to $\partial y$ if the modification is used:

$$\frac{\partial y_m}{\partial h} = \sum_i p_i Z^i$$

$$\frac{\partial y_m}{\partial s} = -(h-f)\frac{1}{w_1}\left(\frac{\mathrm{d}Z}{\mathrm{d}z_1}\right)\sum_i i p_i Z^{i-1}$$

$$\frac{\partial y_m}{\partial w_1} = -(h-f)\frac{x-s}{w_1^2}\left(\frac{\mathrm{d}Z}{\mathrm{d}z_1}\right)\sum_i i p_i Z^{i-1}$$

$$\frac{\partial y_m}{\partial w_2} = 0$$

$$\frac{\partial y_m}{\partial f} = -\sum_i p_i Z^i$$

$$\frac{\partial y_m}{\partial p_i} = (h-f)Z^i$$

$$\frac{\mathrm{d}Z}{\mathrm{d}z_1} = Z\left(\frac{1}{z_1} - 1 - T_1\right) \tag{B.23}$$

This fit method is designated as `lmfit_asym_tanh_then_mod` within CAKE.

[4] The nonlinearity of the modified fit model can be confirmed by comparing analytic vs. numerical calculations of the partial derivatives of the model vs. fit parameters: without the modification, numerical and analytic results agree very well. With the modification, $\partial y/\partial w_2$ no longer matches when d$x$ in the numerical derivative is set by uncertainty in the fit parameter. Shrinking d$x$ improves the agreement of numerical vs. analytic partial derivatives somewhat, but the result of linear error propagation is still unreasonable. Including higher order polynomial terms in the modification worsens the situation. Note that difficulty in calculating numerical partial derivatives is a symptom of nonlinearity, but not a cause of trouble with error propagation because analytic partial derivatives are used.

### B.3 Splines

CAKE's default spline method uses UnivariateSpline in SciPy with Monte-Carlo error analysis and is internally labeled as `basic_monte_carlo_spline`.

### B.4 Polynomial fit

For basic test purposes, a polynomial fit function is provided. The model is

$$y = \sum_i p_i x^i \tag{B.24}$$

The polynomial is set to 5th order by default and the method is designated as `lmfit_poly` within CAKE. It is not the default fit for any parameter, but is provided for developers to use in testing.

## Appendix C. Supplementary data

Supplementary data associated with this article can be found, in the online version, at https://doi.org/10.1016/j.fusengdes.2020.112163.

## References


[1] A.H. Glasser, The direct criterion of newcomb for the ideal MHD stability of an axisymmetric toroidal plasma, Phys. Plasmas 23 (2016) 072505.
[2] P.B. Snyder, H.R. Wilson, J.R. Ferron, L.L. Lao, A.W. Leonard, T.H. Osborne, A. D. Turnbull, D. Mossessian, M. Murakami, X.Q. Xu, Edge localized modes and the pedestal: a model based on coupled peeling-ballooning modes, Phys. Plasmas 9 (2002) 2037.
[3] J.L. Luxon, A design retrospective of the DIII-D tokamak, Nucl. Fusion 42 (2002) 614.
[4] D.N. Hill, theJDIII-D, Team, DIII-D research towards resolving key issues for ITER and steady-state tokamaks, Nucl. Fusion 53 (2013) 104001.
[5] L.L. Lao, H. St John, R.D. Stambauh, A.G. Kellman, W. Pfeiffer, Reconstruction of current profile parameters and plasma shapes in tokamaks, Nucl. Fusion 25 (1985) 1611.
[6] J.L. Luxon, B.B. Brown, Magnetic analysis of non-circular cross-section tokamaks, Nucl. Fusion 22 (6) (1982) 813.
[7] J. Blum, E. Lazzaro, J. O'Rourke, B. Keegan, Y. Stephan, Problems and methods of self-consistent reconstruction of tokamak equilibrium profiles from magnetic and polarimetric measurements, Nucl. Fusion 30 (8) (1990) 1475.
[8] V.D. Shafranov, Determination of the parameters $\beta_i$ and $l_i$ in a tokamak for arbitrary shape of plasma pinch cross-section, Plasma Phys. 13 (1971) 757.
[9] R. Fischer, A. Bock, M. Dunne, J.C. Fuchs, L. Giannone, K. Lackner, P.J. McCarthy, E. Poli, R. Preuss, M. Rampp, M. Schubert, J. Stober, W. Suttrop, G. Tardini, M. Weiland, Coupling of the flux diffusion equation with the equilibrium reconstruction at ASDEX Upgrade, Fusion Sci. Technol. 69 (April (2)) (2016) 526–536.
[10] R. Fischer, A. Bock, A. Burckhart, O.P. Ford, L. Giannone, V. Igochine, M. Weiland, M. Willensdorfer, the ASDEX Upgrade Team, Sawtooth induced q-profile evolution at ASDEX Upgrade, Nucl. Fusion 59 (March (5)) (2019) 056010.
[11] L.L. Lao, H. St. John, R.D. Stambaugh, W. Pfeiffer, Separation of $\beta_p$ and $l_i$ in tokamaks of non-circular cross-section, Nucl. Fusion 25 (1985) 1421.
[12] P.J. McCarthy, The ASDEX Upgrade team, Identification of edge-localized moments of the current density profile in a tokamak equilibrium from external magnetic measurements, Plasma Phys. Control. Fusion 54 (2012) 015010.
[13] L.L. Lao, H.E. St. John, Q. Peng, J.R. Ferron, E.J. Strait, T.S. Taylor, W.H. Meyer, C. Zhang, K.I. You, Mhd equilibrium reconstruction in the diii-d tokamak, Fusion Sci. Technol. 48 (2) (2005) 968–977.
[14] S.A. Sabbagh, R.E. Bell, M.G. Bell, J. Bialek, A.H. Glasser, B. Leblanc, J.E. Menard, F. Paoletti, D. Stutman, E. Fredrickson, A.M. Garofalo, D. Gates, S.M. Kaye, L. L. Lao, R. Maingi, D. Mueller, G. Navratil, M. Ono, M. Peng, E. Synakowski, W. Zhu, Beta-limiting instabilities and global mode stabilization in the National Spherical Torus Experiment, Phys. Plasmas 9 (May (5)) (2002) 2085–2092.
[15] S.A. Sabbagh, J.M. Bialek, R.E. Bell, A.H. Glasser, B.P. LeBlanc, J.E. Menard, F. Paoletti, M.G. Bell, R. Fitzpatrick, E.D. Fredrickson, A.M. Garofalo, D.A. Gates, S. M. Kaye, L.L. Lao, R. Maingi, D. Mueller, G.A. Navratil, D. Stutman, W. Zhu, N.S.T. X. the, Research Team, The resistive wall mode and feedback control physics design in NSTX, Nucl. Fusion 44 (2004) 560.
[16] Y.S. Park, S.A. Sabbagh, W.H. Ko, J.G. Bak, J.W. Berkery, J.M. Bialek, M.J. Choi, S. H. Hahn, Y.K. In, S.C. Jardin, Y.M. Jeon, J. Kim, J.G. Kwak, S.G. Lee, Y.K. Oh, H. K. Park, S.W. Yoon, G.S. Yun, Investigation of instabilities and rotation alteration in high beta kstar plasmas, Phys. Plasmas 24 (2017) 012512.
[17] N.C. Logan, B.A. Grierson, S.R. Haskey, S.P. Smith, O. Meneghini, D. Eldon, OMFIT tokamak profile data fitting and physics analysis, Fusion Sci. Technol. 74 (1–2) (2018) 125–134.
[18] G.Q. Li, Q.L. Ren, J.P. Qian, L.L. Lao, S.Y. Ding, Y.J. Chen, Z.X. Liu, B. Lu, Q. Zang, Kinetic equilibrium reconstruction on EAST tokamak, Plasma Phys. Control. Fusion 55 (2013) 125008.
[19] D. Eldon, B.D. Bray, T.M. Deterly, C. Liu, M. Watkins, R.J. Groebner, A.W. Leonard, T.H. Osborne, P.B. Snyder, R.L. Boivin, G.R. Tynan, Initial results of the high resolution edge Thomson scattering upgrade at DIII-D, Rev. Sci. Instrum. 83 (2012) 10E343.
[20] T.N. Carlstrom, J.C. DeBoo, R. Evanko, C.M. Greenfield, C.L. Hsieh, R.T. Snider, P. Trost, A compact, low cost, seven channel polychromator for Thomson scattering measurements, Rev. Sci. Instrum. 61 (1990) 2858.
[21] K.H. Burrell, P. Gohil, R.J. Groebner, D.H. Kaplan, J.I. Robinson, W.M. Solomon, Improved charge-coupled device detectors for high-speed, charge exchange spectrscopy studies on the DIII-D tokamak, Rev. Sci. Instrum. 75 (2004) 3455.
[22] A.W. Leonard, W.H. Meyer, B. Geer, D.M. Behne, D.N. Hill, 2d tomography with bolometry in DIII-D, Rev. Sci. Instrum. 66 (1995) 1201.
[23] R.J. Hawryluk, An empirical approach to tokamak transport, in: Physics of Plasmas Close to Thermonuclear Conditions, vol. 1, 1980, p. 19. http://w3.pppl.gov/transp/papers/Hawryluk.pdf.
[24] H. St. John, T.S. Taylor, Y.R. Lin-Liu, A.D. Turnbull, Transport simulation of negative magnetic shear discharges, in: Plasma Physics and Controlled Nuclear Fusion Research 1994, vol. 3, International Atomic Energy Agency, 1994, pp. 603–614.
[25] D. Eldon, R.L. Boivin, R.J. Groebner, T.H. Osborne, P.B. Snyder, A.D. Turnbull, J. A. Boedo, K.H. Burrell, E. Kolemen, L. Schmitz, G.R. Tynan, H.R. Wilson, Investigation of peeling-ballooning stability prior to transient outbursts accompanying transitions out of H-mode in DIII-D, Phys. Plasmas 22 (2015) 052109.
[26] O. Meneghini, S.P. Smith, L.L. Lao, O. Izacard, Q. Ren, J.M. Park, J. Candy, Z. Wang, C.J. Luna, V.A. Izzo, Integrated modeling applications for tokamak experiments with OMFIT, Nucl. Fusion 55 (2015) 083008.
[27] D. Eldon, E. Kolemen, J.L. Barton, A.R. Briesemeister, D.A. Humphreys, A. W. Leonard, R. Maingi, M.A. Makowski, A.G. McLean, A.L. Moser, P.C. Stangeby, Controlling marginally detached divertor plasmas, Nucl. Fusion 57 (2017) 066039.
[28] R.J. Colchin, D.L. Hillis, R. Maingi, C.C. Klepper, N.H. Brooks, The filterscope, Rev. Sci. Instrum. 74 (2003) 2068.
[29] R. Nazikian, C. Paz-Soldan, J. deGrassie, D. Eldon, T.E. Evans, N.M. Ferraro, B. A. Grierson, R.J. Groebner, S. Haskey, J.D. King, N. Logan, G.R. McKee, R. A. Moyer, M. Okabayashi, D.M. Orlov, T.H. Osborne, J.-K. Park, T.L. Rhodes, M. W. Shaefer, S.P. Smith, P.B. Snyder, E.J. Strait, M.R. Wade, Pedestal bifurcations and resonant field penetration near the threshold of ELM suppression in the DIII-D tokamak, Phys. Rev. Lett. 114 (2015) 105002.
[30] R.J. Groebner, D.R. Baker, K.H. Burrell, T.N. Carlstrom, J.R. Ferron, P. Gohil, L. L. Lao, T.H. Osborne, D.M. Thomas, W.P. West, J.A. Boedo, R.A. Moyer, G. R. McKee, R.D. Deranian, E.J. Doyle, C.L. Rettig, T.L. Rhodes, J.C. Rost, Progress in quantifying the edge physics of the H mode regime in DIII-D, Nucl. Fusion 41 (2001) 1789.
[31] O. Sauter, C. Angioni, Y.R. Lin-Liu, Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime, Phys. Plasmas 6 (1999) 2834, https://doi.org/10.1063/1.873240. See Sauter et al. [32] for erratum.
[32] O. Sauter, C. Angioni, Y.R. Lin-Liu, Erratum: 'neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime' [phys. plasmas 6, 2834 (1999)], Phys. Plasmas 9 (2002) 5140, https://doi.org/10.1063/1.1517052. Erratum for Sauter et al. [31].
[33] W.M. Solomon, K.H. Burrell, P. Gohil, R. Groebner, D. Kaplan, Cross-calibrating spatial positions of light-viewing diagnostics using plasma edge sweeps in DIII-D, Rev. Sci. Instrum. 74 (2003) 5084.
[34] R.S. Wilcox, M.W. Shafer, N.M. Ferraro, G.R. McKee, L. Zeng, T.L. Rhodes, J. M. Canik, C. Paz-Soldan, R. Nazikian, E.A. Unterberg, Evidence of toroidally localized turbulence with applied 3d fields in the DIII-D tokamak, Phys. Rev. Lett. 117 (2016) 135001.

[35] P. Stangeby, J. Canik, J. Elder, C. Lasnier, A.W. Leonard, D. Eldon, M. Makowski, T. Osborne, J. Watkins, Identifying the location of the OMP separatrix in DIII-D using power accounting, Nucl. Fusion (accepted) 55 (2015) 093014.
[36] D. Wróblewski, L.L. Lao, Polarimetry of motional stark effect and determination of current profiles in DIII-D, Rev. Sci. Instrum. 63 (1992) 5140.
[37] T.C. Luce, C.C. Petty, W.H. Meyer, C.T. Holcomb, K.H. Burrell, L.J. Bergsten, Method for correction of measured polarization angles from montional stark effect spectroscopy for the effects of electric fields, Plasma Phys. Control. Fusion 58 (2016).
[38] S. Koh, C.S. Chang, S. Ku, J.E. Menard, H. Weitzner, W. Choe, Bootstrap current for the edge pedestal plasma in a diverted tokamak geometry, Phys. Plasmas 19 (2012) 072505.
[39] E.A. Belli, J. Candy, O. Meneghini, T.H. Osborne, Limitations of bootstrap current models, Plasma Phys. Control. Fusion 56 (2014) 045006.
[40] E.A. Belli, J. Candy, Kinetic calculation of neoclassical transport including self-consistent electron and impurity dynamics, Plasma Phys. Control. Fusion 50 (2008) 095010.
[41] R.J. Goldston, D.C. McCune, H.H. Towner, S.L. Davis, R.J. Hawryluk, G.L. Schmidt, New techniques for calculating heat and particle source rates due to neutral beam injection in axisymmetric tokamaks, J. Comput. Phys. 41 (1981) 61–78.
[42] T.D. Rognlien, J.L. Milovich, M.E. Rensink, G.D. Porter, A fully implicit, time dependent 2-D fluid code for modeling tokamak edge plasmas, J. Nucl. Mater. 196–198 (1992) 347–351.
[43] H.R. Wilson, P.B. Snyder, G.T.A. Huysmans, R.L. Miller, Numerical studies of edge localized instabilities in tokamaks, Phys. Plasmas 9 (2002) 1277.
[44] A.S. Glasser, E. Kolemen, A.H. Glasser, A Riccati solution for the ideal MHD plasma response with applications to real-time stability control, Phys. Plasmas 25 (2018) 032507.
[45] A.S. Glasser, E. Kolemen, A robust solution for the resistive mhd toroidal $\Delta'$ matrix in near real-time, Phys. Plasmas 25 (2018) 082502.
[46] Y. Fu, D. Eldon, K. Erickson, K. Kleijwegt, L. Lupin-Jimenez, D. Mark, N. Boyer, N. Eidietis, O. Barbour, E. Izacard, Kolemen, Machine learning control for disruption and tearing mode avoidance, Phys. Plasmas 27 (February (2)) (2020) 022501.