

# Inference of Experimental Radial Impurity Transport on Alcator C-Mod: Bayesian Parameter Estimation and Model Selection.

F. Sciortino<sup>1</sup>‡, N.T. Howard<sup>1</sup>, E.S. Marmor<sup>1</sup>, T. Odstrcil<sup>1</sup>, N.M. Cao<sup>1</sup>, R. Dux<sup>2</sup>, A.E. Hubbard<sup>1</sup>, J.W. Hughes<sup>1</sup>, J.H. Irby<sup>1</sup>, Y. Marzouk<sup>1</sup>, L.M. Milanese<sup>1</sup>, M.L. Reinke<sup>3</sup>, J.E. Rice<sup>1</sup>, P. Rodriguez-Fernandez<sup>1</sup>

<sup>1</sup>Massachusetts Institute of Technology, Cambridge, MA, 02139, USA

<sup>2</sup>Max Planck Institute - Institute for Plasma Physics, Garching, Germany

<sup>3</sup>Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA

**Abstract.** We present a fully Bayesian approach for the inference of radial profiles of impurity transport coefficients and compare its results to neoclassical, gyrofluid and gyrokinetic modeling. Using nested sampling, the Bayesian Impurity Transport Inference (BITE) framework can handle complex parameter spaces with multiple possible solutions, offering great advantages in interpretative power and reliability with respect to previously demonstrated methods. BITE employs a forward model based on the `pySTRAHL` package, built on the success of the well-known `STRAHL` code [Dux, IPP Report, 2004], to simulate impurity transport in magnetically-confined plasmas. In this paper, we focus on calcium (Ca,  $Z=20$ ) Laser Blow-Off injections into Alcator C-Mod plasmas. Multiple Ca atomic lines are diagnosed via high-resolution X-ray Imaging Crystal Spectroscopy and Vacuum Ultra-Violet measurements. We analyze a sawtoothed I-mode discharge for which neoclassical and turbulent (quasilinear and nonlinear) predictions are also obtained. We find good agreement in diffusion across the entire radial extent, while turbulent convection and density profile peaking are estimated to be larger in experiment than suggested by theory. Efforts and challenges associated with the inference of experimental pedestal impurity transport are discussed.

## 1. Introduction

Stable high-performance tokamak operation requires complex trade-offs to maintain acceptable impurity conditions. Core-edge integration demands minimization of core radiation and dilution [1] while also sustaining a detached divertor regime, envisioned to only be possible via puffing of impurities in the edge [2]. Choices of wall materials are also known to affect plasma fueling, pedestal formation and high-confinement (H-mode) operation [3]. Developing and validating accurate models for impurity transport and atomic data are paramount to make predictions in a whole-device perspective.

‡ Email: sciortino@psfc.mit.edu

Historically, transport model validation for fusion plasmas has mostly focused on heat transport [4, 5]. Accurate estimation of experimental particle transport coefficients has proven to be more demanding, although significant progress has been made, particularly on understanding the origin of electron density profile peaking [6–10]. Detailed validation of impurity transport theory is more challenging, both from an experimental and a theoretical perspective.

Since transport codes tend to make predictions for *fluxes*, validation efforts would ideally compare *fluxes* between experiments and theoretical predictions. This, however, would be a formidable task for impurity transport, to the required level of accuracy, since knowledge of sources, atomic rates and diagnostic calibrations each present significant challenges. We thus resort to transport coefficients, diffusion ( $D$ ) and convection ( $v$ ), to separate the particle flux ( $\Gamma$ ) via  $\Gamma := -D \nabla n + v n$ .  $D$  and  $v$  are more effective metrics to compare experiment and theory without over-reliance on experimental details. In quasi-steady conditions with no impurity core sources, we may expect  $\Gamma = 0$ , which allows one to relate particle density gradients to the  $v/D$  ratio via  $v/D = \nabla n/n$ . It is only in the presence of time-dependent dynamics that  $D$  and  $v$  may be experimentally separated, offering much stronger constraints for transport model validation. For low- $Z$  ions that may be assumed to be fully-ionized, core transport coefficients may be estimated using relatively simple flux-gradient methods [11], which however do not generalize to higher- $Z$  impurities for which atomic physics plays a key role. In this work, we make use of transient impurity injections produced via the Laser Blow-Off (LBO) technique [12].

The inference of  $D$  and  $v$  relies on an iterative process in which predictions of charge state densities and emissivity profiles are computed for given sets of parameters choices, attempting to match experimental observations. For each choice of parameters, a “goodness-of-fit” metric is calculated in order to identify parameters that most closely match reality. This is referred to as an *inverse problem*, as opposed to the operation of computing a signal prediction for a given choice of model and parameters (*forward modeling*). Ideally, we would like to solve an inverse problem with the lowest possible number of free parameters (inputs to the forward model) while still reproducing experimental behavior to the highest possible degree of fidelity.

Efforts to estimate radial profiles of  $D$  and  $v$  have been reported from all major devices, including ASDEX-Upgrade [13–16], DIII-D [17, 18], JET [19, 20], Tore Supra [21, 22], TCV [23], HL-2A [24], W7-X [25], MAST [26], NSTX [27] and Alcator C-Mod [28–32]. In several studies, gas puffing or supersonic pellets were used to inject small amounts of impurities; both of these methods may incur difficulties when modeling the impurity source. On the other hand, the Laser Blow-Off (LBO) technique [12] allows the injection of non-recycling impurities via a time-resolved and non-perturbative laser ablation, whose source function can be more easily characterized.

On Alcator C-Mod, the inference of experimental impurity transport coefficients was previously attempted for L-mode discharges using a  $\chi^2$  minimization via the Levenberg-

Marquardt Algorithm [31]. This approach suffers from the risk of finding *local*  $\chi^2$  minima rather than *global* ones, depending on initial guesses of model parameters. It also lacks of rigorous uncertainty quantification and model selection, as discussed by Chilenski *et al.* [33], who adopted Bayesian techniques with synthetic data to demonstrate the critical importance of avoiding under- or over-fitting [33]. Unfortunately, synthetic data often makes it difficult or impossible to explore issues of *model inadequacy* that may result from oversimplified experimental analysis, inaccurate atomic rates, discrepancies of detector models and other important factors [34]. For example, the omission of sawtooth modeling or the time-dependence of atomic rates, as in previous C-Mod work [31–33], may prevent the optimal model complexity to emerge.

In this work, we demonstrate parameter estimation and model selection using a forward model that has been improved in both its physics fidelity and its iterative capability. We describe our approach on experimental data, thus parting from any idealizations of synthetic datasets. Our fully-Bayesian framework employs nested sampling algorithms and introduces faster forward modeling, higher physics fidelity, improved treatment of experimental measurements and nuisance parameters, as well as more rigorous combination of diagnostic instruments. We attempt to model an extensive range of experimental details to improve comparison of experiment with neoclassical, gyro-fluid and gyrokinetic models.

The structure of this paper is as follows. In section 2, we describe the typical experimental setup for impurity transport experiments at Alcator C-Mod and we introduce a high-performance I-mode discharge that will be used to demonstrate our inference methods. A detailed description of our inference framework is given in section 3. In section 4 we compare results of experimental inferences to neoclassical and turbulence modeling. The latter is analyzed in more detail in section 5. Finally, in section 6 we discuss and summarize this work. Additional details on numerical schemes and experimental data are provided in the appendices.

## 2. Experimental Setup and Measurements

The analysis methods presented in the next section are applicable to a wide range of discharges and we therefore describe the general experimental setup that can be used for impurity transport inferences. We focus on data from Alcator C-Mod [35], a compact ( $a = 0.22$  m,  $R = 0.68$  m), high field, diverted tokamak that operated up to the end of 2016 with up to 5.5 MW of ion cyclotron resonance heating (ICRH). C-Mod parameters have ranged within  $B_T = [2.0 - 8.1]$  T,  $I_p = [0.4, 2.0]$  MA,  $n_e(r = 0) = [0.2 - 6.0] \times 10^{20} \text{ m}^{-3}$  and  $T_e(0)$  up to 9 keV [36].

On Alcator C-Mod, a multi-pulse LBO system has been used for transient impurity studies for several years [31, 37–40]. Here, we consider experiments where a  $\text{CaF}_2$  film was ablated and calcium ( $Z_{Ca} = 20$ ) transport was studied by diagnosing impurity emission with a combination of spectroscopic measurements. A significant advantage of Ca over other ions is that it is non-intrinsic and non-recycling, and it is also simple to ablate from a LBO

slide [37].

An initial condition for particle transport reconstructions was provided by an optical fiber viewing the edge of the plasma, near the injection location of the LBO. A bandpass filter limited measured brightness to the  $420 \pm 10$  nm range, providing a proxy measurement of Ca-I emission [31]. This system offers high signal-to-noise ratio (SNR) and sometimes provides evidence of CaF<sub>2</sub> clusters entering the plasma; this often correlates with irregularities in signals on several other spectrometers.

For our inferences, a central role is played by C-Mod's XICS diagnostic [41], which has 3 ms readout time and typical signal integration over 3 ms intervals, for a total time resolution of 6 ms. The two crystals of this system are normally configured to observe lines from the H- and He-like charge states of argon. Doppler broadening and Doppler shift of the emission spectra provide radial profiles of ion temperature ( $T_i$ ) and toroidal rotation ( $v_\phi$ ) following tomographic inversion [41, 42].

For impurity transport studies, one of the two XICS crystals could be substituted with one capable of viewing the Ca<sup>18+</sup> (He-like) spectrum, often discretizing the detector spatial coverage into 32 independent spatial chords. In standard C-Mod discharges, the brightest Ca line in the core is the  $1s^2 \ ^1S_0 - 1s2p \ ^1P_1$  resonance (w) line at  $3177.26 \text{ m}\text{\AA}$ , and past inferences of impurity transport coefficients in Refs. [31, 32, 38] relied mostly on its measurement using existing spectral analysis routines [43]. As described in section 3.2, fitting and uncertainty quantification for the z line ( $1s^2 \ ^1S_0 - 1s2p \ ^1P_1$ ,  $3211.13 \text{ m}\text{\AA}$ ) for impurity transport studies has also recently become routine.

Finally, a single-chord VUV spectrometer (XEUS) focused on the 10-70 Å range can be used to view multiple emission lines for Li-like Ca [44]. In recent years, a companion Long-wavelength Extreme Ultraviolet Spectrometer (LoWEUS) spectrometer [45] has allowed emission in the 100-300 Å region to also be analyzed. Unfortunately, LoWEUS data are available only in a small subset of C-Mod data. While the spectral range of these spectrometers typically offers valuable observations of Ca<sup>16+</sup> and Ca<sup>17+</sup> at  $r/a > 0.7$ , their relatively low time resolution (2 ms) does not usually allow clear observation of the signal rise phase following LBO injections.

In this paper, we focus on a sawtoothed I-mode discharge run with  $B_t \approx 5.5 T$  and  $I_p \approx 1.0 MA$ , reaching  $T_e \approx 5 \text{ keV}$  and  $n_e \approx 1.7 \times 10^{20} \text{ m}^{-3}$  on axis. Using the neoclassical prediction of Sauter *et al.* [46], we estimate  $Z_{eff} \approx 2.1 \pm 0.2$  during the current flat-top and assume radial  $Z_{eff}$  profiles to be flat, consistently with past Visible Bremsstrahlung measurements on C-Mod [47]. This leads to a volume-averaged (dimensionless) collisionality, defined as  $\nu_{ei} = 0.1 Z_{eff} \langle n_e \rangle R / \langle T_e \rangle^2$  [48], of  $0.9 \pm 0.2$ , a relatively low value for standard Alcator C-Mod operation. This discharge presents a temperature pedestal, but the electron density decreases smoothly across the LCFS. This is consistent with the typical phenomenology of I-mode, which is characterized by good energy confinement, no Edge-Localized Modes (ELMs), a weakly-coherent  $\sim 100 - 300 \text{ kHz}$  mode near the LCFS, high levels of particle transport and low impurity retention [49, 50].

Fig. 1 shows w and z line brightness signals from sample XICS chords following a Ca LBO injection during the current flat-top of this discharge. Orange lines in Fig. 1 show signal fits from the impurity transport inference described in section 4. Tangency radii for these chords are shown in the plots on the right hand side and are below the midplane for low-numbered chord and above it for high-numbered ones. In high performance discharges such as the presented I-mode, He-like Ca emission is collected via XICS up to  $r/a \approx 0.8$ , outside of which XEUS still measures several Li-like Ca lines.

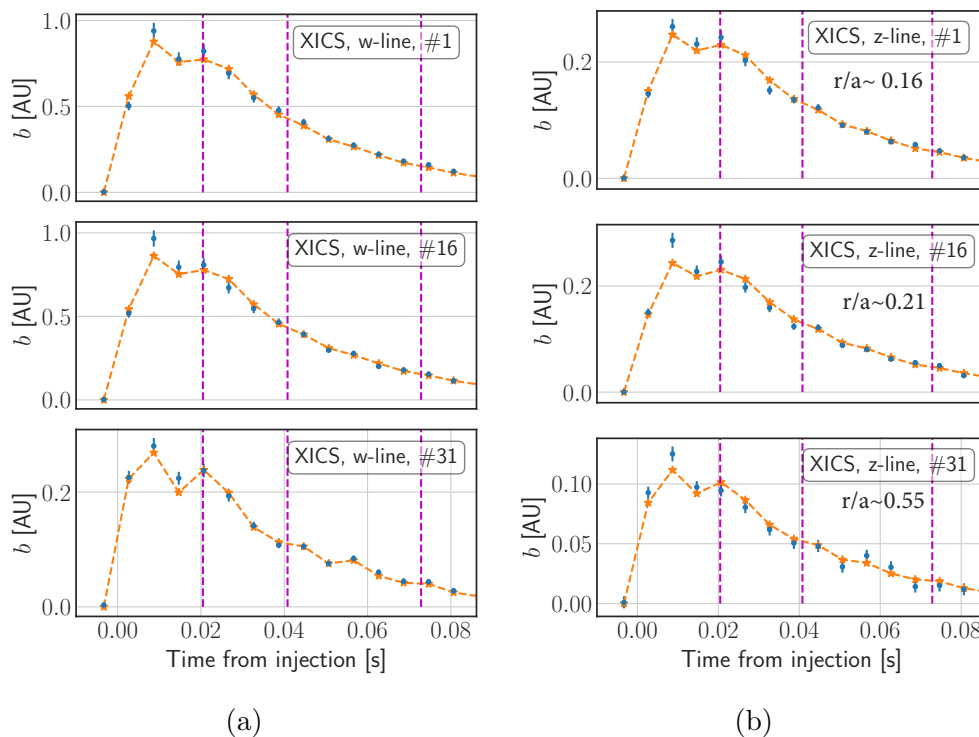


Figure 1: XICS  $\text{Ca}^{18+}$  w (left) and z line (right) signals from sample chords following a LBO injection into the I-mode discharge of interest. Tangency radii are indicated within plots on the right. Orange dashed lines show fits from the impurity transport inference described in section 4. Vertical magenta dashed lines indicate times at which sawteeth occur.

In Fig. 2, we show time traces of H-minority Ion Cyclotron Resonance Heating (ICRH) power ( $P_{RF}$ ) with a near-axis resonance, central electron temperature ( $T_{e,0}$ ), total radiated power from bolometry ( $P_{rad}$ ) and line-averaged density ( $\bar{n}_e$ ) from the same I-mode discharge. As it is often the case in C-Mod, sawteeth modulate core temperatures and they are known to also affect impurity profiles. Previous work on C-Mod focused on inferring “sawtooth-averaged” impurity transport [31, 51], which may be a sufficiently good approximation to background transport coefficients in low-performance discharges with small sawteeth. In order to examine higher- $T_e$  discharges, broadening the range of analyzed confinement states and enabling higher SNR from spectroscopic signals in the plasma edge, some form of sawtooth modeling is necessary, and will be described in section 3.1.




Figure 2: Typical time series of radio-frequency power ( $P_{RF}$ ), near-axis electron temperature ( $T_e^0$ ), radiated power ( $P_{rad}$ ) and line-averaged electron density ( $n_e$ ) for the I-mode discharge described in the main text. The LBO injection is seen not to perturb the electron density.

A requirement for our transport inferences to be meaningful for validation purposes is that the plasma is quasi-steady and mostly unperturbed by LBO injections. This is illustrated in Fig. 2, where the line averaged density (lowest panel, measured by a Two-Color Interferometer) appears unaffected by the  $\text{CaF}_2$  influx (shown by radiated power, measured via bolometry).

We adopt different kinetic profile fitting methods for impurity transport forward modeling and for theoretical (neoclassical, gyro-fluid and gyrokinetic) predictions. For forward modeling, we make use of a parametric time-dependent Radial Basis Functions (RBF), taking care to avoid smoothing of sawteeth and to enforce pedestal structure. Use of fast Electron Cyclotron Emission data allows us to fit very well the  $T_e$  modulation, while for  $n_e$  we rely on Thomson Scattering data at lower time resolution.

For theoretical modeling, we adopt Gaussian Process Regression (GPR), which minimizes user bias by inferring hyperparameters corresponding to fits of highest probability [52]. In Fig. 3, we show kinetic profile fits ( $n_e$ ,  $T_e$  and  $T_i$ ) and corresponding normalized inverse gradient scale lengths ( $a/L_{n_e}$ ,  $a/L_{T_e}$  and  $a/L_{T_i}$ , with  $L_X \equiv -X/(dX/dr)$ ), averaging experimental data over 150 ms near the LBO injection of interest. In the absence of experimental  $T_i$  measurements in the pedestal, we assume high collisional coupling between electrons and ions and take  $T_i \approx T_e$  in this region, as experimentally observed in past work [53]. We quantify GPR uncertainties by propagating diagnostic uncertainties and data scattering in time via the Law of Total Variance and MCMC sampling [54], using a Gibbs kernel with tanh length scale, as suggested in Ref. [52].




Figure 3: Time-averaged kinetic profiles for the I-mode discharge, fitted via Gaussian Process Regression. In (b), we focus on gradient scale lengths between  $r/a = 0.2$  and  $r/a = 0.8$ .

### 3. Methods

In this section, we describe the Bayesian Impurity Transport Inference (BITE) framework developed in this work and broadly represented in Fig. 4. At the highest level, BITE uses the nested sampling algorithm to iteratively test predictions from possible parameter choices. In doing so, it repeatedly runs a forward model for impurity transport, represented in the right hand side box of Fig. 4. In what follows, we describe in greater depth each aspect of this diagram.

#### 3.1. Impurity Transport Forward Modeling

To forward-model the radial transport of an impurity “I”, we make use of the continuity equations for each impurity charge state “Z”,

$$\frac{\partial n_{I,Z}}{\partial t} = -\frac{1}{r} \frac{\partial}{\partial r} \left( r \bar{\Gamma}_{I,Z} \right) + Q_{I,Z}, \quad (1)$$

with the radial particle flux  $\Gamma_{I,Z}$  being

$$\Gamma_{I,Z} = -D \frac{\partial n_{I,Z}}{\partial r} + v n_{I,Z}. \quad (2)$$

The source/sink term  $Q_{I,Z}$  couples the transport equation of each ionization stage with the neighboring ones via ionization and recombination processes, for which we use rates



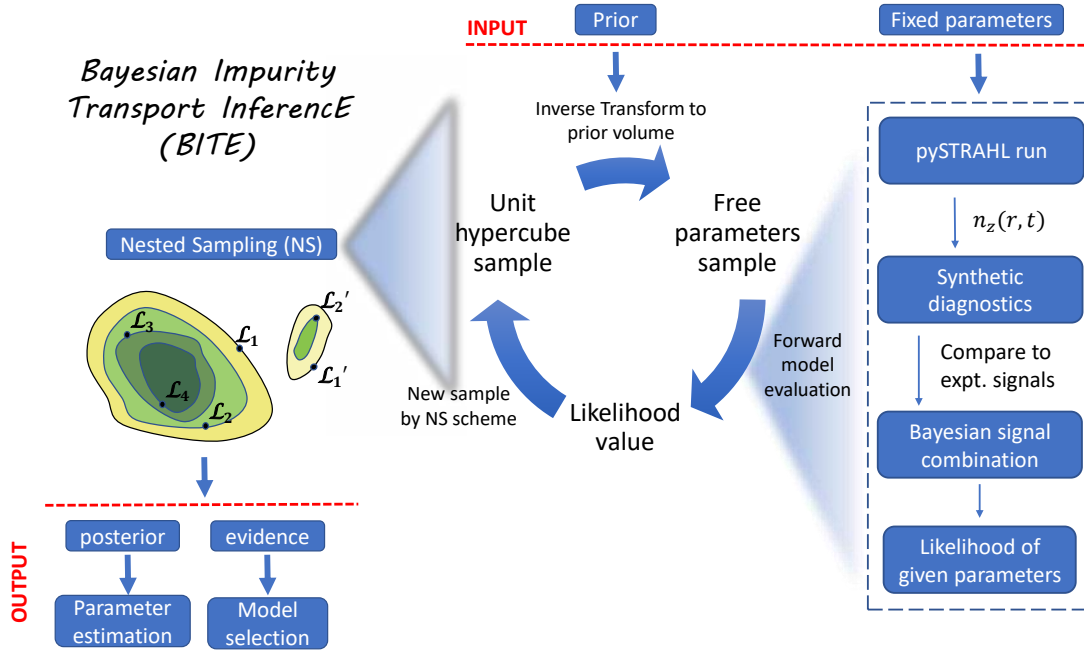


Figure 4: Overview of the components of the BITE framework to infer radial profiles of impurity transport coefficients.

from the Atomic Database Analysis System (ADAS) [55]§. We consider flux-averaged transport coefficients and label flux surfaces based on the volume they enclose, using  $r = \sqrt{V/(2\pi^2 R_{axis})}$ .

Theoretical transport models can often be used to predict coefficients  $D_0$  and  $v_0$  with the respect to the low-field-side (LFS) density  $n_0$ . To compare these results, on a minor radius coordinate  $r_m$ , to those from forward modeling, on the  $r$  coordinate defined above, one can perform the following transformations [9]:

$$D = D_0 \frac{n_0}{\langle n \rangle} \left( \frac{\partial r}{\partial r_m} \right)^2 \quad \text{and} \quad v = v_0 \frac{n_0}{\langle n \rangle} \frac{\partial r}{\partial r_m} + D \frac{\partial}{\partial r_m} \left( \ln \frac{n_0}{\langle n \rangle} \right) \quad (3)$$

where factors of  $n_0/\langle n \rangle$  account for poloidal density asymmetries, expected to be significant for heavy ions [56, 57]. The  $n_0/\langle n \rangle$  ratio may be estimated using a neoclassical code or, more simply, an analytical prediction of centrifugal effects [57]. A detailed description of this correction procedure, which can modify  $D$  and  $v$  estimates by up to 50% for Ca ions in C-Mod plasmas, can be found in Appendix A of Ref. [9].

Eqs. 1 and 2 represent a system of coupled differential equations for all charge states. We solve this using a new **pySTRAHL** package, which builds on the well-known **STRAHL** code [58], with which it has been thoroughly benchmarked. **pySTRAHL** is written in Fortran 90 and

§ We take thermal charge exchange processes to be negligible, assuming that neutrals do not penetrate much further than the LCFS in the high-density C-Mod conditions.



couples to Python via the `f2py` package, maintaining modeling capabilities of the original STRAHL while avoiding the redundant input/output (I/O) normally required to iterate over STRAHL inputs. At present, pySTRAHL iterations can either be initialized directly in Python or by reading the output of a first STRAHL run which is written to a NETCDF file and can later be read into computer memory. Following this initialization stage, pySTRAHL iterations only require modification of input variables in Python.

With either initialization method, pySTRAHL requires similar inputs to STRAHL: plasma geometry, atomic rates, kinetic profiles, impurity sources, recycling and scrape-off-layer (SOL) parameters and transport coefficients (diffusion and convection as a function of radius). The latter are the quantities that we wish to infer, and are therefore iterated over according to a chosen algorithm. At all iterations, the values of  $D$  and  $v$  at a specified number of knots are interpolated using Piecewise Cubic Hermite Interpolating Polynomial (PCHIP) splines [59], which enforce monotonicity between knots and avoid over-shooting. In BITE, these profiles are typically summed with an additional Gaussian feature in the  $v$  edge profile, representing the expected particle pinch in the pedestal region. This Gaussian is typically centered at the LCFS and its width is allowed to be a free parameter. This parameterization enables one to represent sharp changes in the  $v$  profile without necessitating a large number of additional free spline knots or detailed experimental data in the pedestal.

In order to compare transport models for impurity propagation through the pedestal, we allow  $v/D$  to scale with the  $Z$  value of each charge state outside of a specified radial region. Such scaling is expected by neoclassical theory [29, 60] and has been observed in previous experimental work [15]. The radial location outside of which  $v$  is taken to scale with  $Z$  is determined by a parameter that we normally fix for simplicity to be  $r/a = 0.9$  (approximately the pedestal top). One can then test how results compare to those obtained by standard runs with no  $v$  charge-state dependence.

In sawtooth discharges, the application of a sawtooth model can be important. STRAHL allows one to flatten all charge state density profiles inside of the sawtooth mixing radius at given crash times; this phenomenological picture was suggested by experimental observations by Seguin *et al.* [61]. A drawback of this simple sawtooth model is that it creates sharp density gradients at each sawtooth crash, thus demanding high temporal resolution to remain numerically stable. By implementing a smoother (neither more nor less physical) sawtooth crash model, such requirements have been alleviated, effectively halving the runtime of pySTRAHL simulations. Crashes were implemented using complementary error functions as

$$n_z(r) \leftarrow \frac{\bar{n}_z}{2} \operatorname{erfc}\left(\frac{r - r_{mix}}{w}\right) + \frac{n_z(r)}{2} \operatorname{erfc}\left(\frac{r_{mix} - r}{w}\right) \quad (4)$$

where the left-pointing arrow shows how  $n_z(r)$  changes across the crash. Here,  $n_z$  is the charge state density,  $\bar{n}_z$  is its average evaluated between the magnetic axis and the mixing

|| A weaker neoclassical dependence of  $D$  on  $Z$  is also expected, but we have not attempted to model this effect.

radius before sawtooth crashes and  $w \approx 1$  cm is a “crash region width”, which sets the spatial smoothness of the sawtooth crash.

The output of pySTRAHL consists of time-dependent charge state densities for all ionization stages of an impurity ion in each of the particle reservoirs. By summing densities in the plasma, wall and divertor regions, we have ensured that particle conservation is maintained in our simulations. At all iterations, charge state distributions are combined with atomic rates to compute local emissivities. These are then line-integrated to compare to spectroscopic diagnostic brightness. Radial profiles of transport coefficients are then modified to better match experimental data (based on algorithmic schemes) and are given as inputs to pySTRAHL, allowing for a new comparison with experimental signals.

pySTRAHL offers major advantages over standard STRAHL operation for iterative work. First of all, by caching spatio-temporal grids, atomic rates, geometry and other fixed parameters at an initialization stage, we obtain a speed improvement of approximately 70%, bringing our runtime down to  $\approx 150$  ms per simulation for the I-mode case discussed in this work¶. The speed-up can be even more significant for runs where input/output (I/O) operations would normally take a larger fraction of the STRAHL runtime. Moreover, by avoiding I/O at every iteration we are able to efficiently parallelize our inference and make use of hundreds of CPUs. Without this capability, reaching convergence with the advanced algorithms described in section 3.3, requiring tens of millions of iterations at times, would be unfeasible.

### 3.2. Spectroscopic analysis

The accurate analysis and interpretation of atomic spectra, particularly for high-resolution data including overlapping lines, is paramount for transport inferences. The development of the Bayesian Spectral Fitting Code (BSFC) has recently offered new opportunities to analyze XICS data on C-Mod where traditional spectral fitting tools would find it difficult to effectively subtract overlapping satellite lines from a primary line of interest [62]. BSFC addresses this problem with a decomposition of each lineshape into a Hermite polynomial series and using Bayesian sampling techniques, analogous to those used in BITE. This allows robust and rigorous truncation, as well as uncertainty quantification. In Fig. 5 we show spectral fits obtained for the I-mode discharge of interest at  $t \approx 1.21$  s, following the Ca injection at  $t = 1.2$  s. Fig. 5a shows  $n \geq 4$  satellite lines blending with the  $1s^2-1s2p$  resonance w line (rest wavelength of 3177.3 mÅ). In Fig. 5b, we show the forbidden z line (rest wavelength of 3211.1 mÅ) blending with the j line, whose amplitude, width and Doppler shift are physically constrained by the nearby k line. The spectra in Fig. 5a and 5b are measured on two adjacent detector submodules, whose accurate cross-calibration has been demonstrated in previous work (e.g. Ref. [39]) where w and z line spectra could be simultaneously matched by atomic modeling. Therefore, we can make use of the relative amplitude of these lines without running into cross-calibration issues. Since the diagnostic

¶ These times refer to cluster nodes using Intel(R) Xeon(R) CPU E5-2683 v4, 2.1 GHz.

is not absolutely calibrated, we normalize both the w and z line signals to the maximum signal observed across all spatial chords for the two lines together.

Measuring both the w and z line emission offers a powerful way to constrain impurity transport, since these lines arise from different physical processes: the upper level of the w transition is populated via collisional excitation of the ground state of the He-like ion; for the z transition, the upper state is instead populated via radiative recombination of the H-like ion and inner shell ionization of Li-like one. As one might expect, these processes have different  $T_e$  dependence, with z line emissivity becoming stronger at lower  $T_e$  and signals extending to the pedestal [30,42]. In order to compute atomic rates, we make use of formulae and coefficients provided by Mewe and Schrijver [63], which have previously been shown to match C-Mod XICS spectra with great accuracy [40]. This calculation requires detailed knowledge of the H-, He- and Li-like ion populations, which are output by pySTRAHL after each run; therefore, w and z line rates are computed at every iteration in BITE. The ratio of w and z line intensities effectively constrains whether the impurity of interest is recombining, ionizing or is in steady-state at the local temperature [64], providing a strong constraint on plasma transport properties. Optimizing the calculation of these rates in Fortran 90, we have been able to reduce the runtime to 75 ms, making our entire forward model execution (including a pySTRAHL run and the application of synthetic diagnostics) shorter than 300 ms.

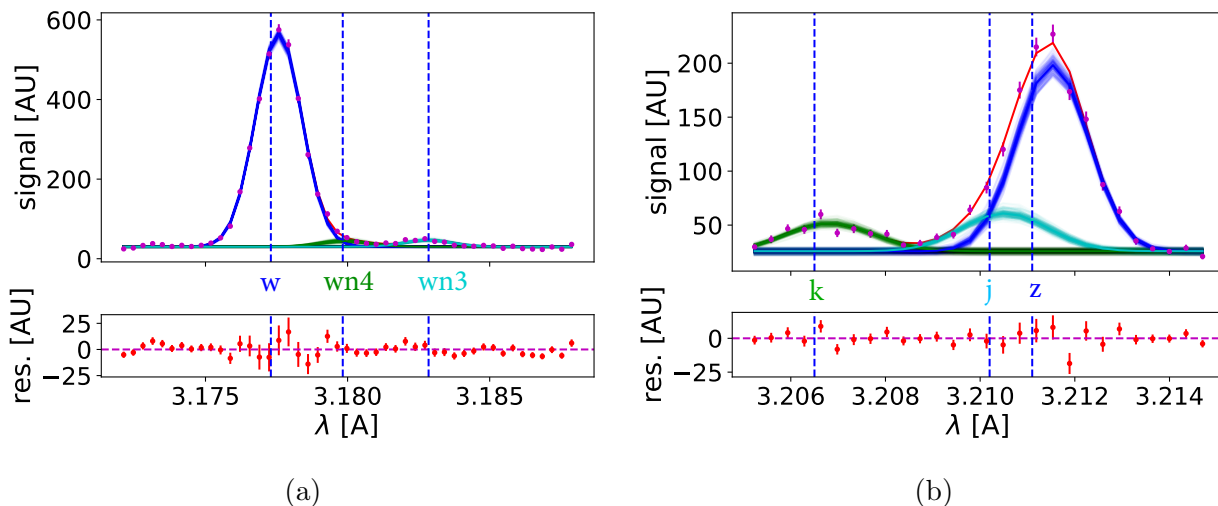


Figure 5: Typical fits of the XICS He-like Ca spectra. In (a), the brightest line is the w resonance; in (b), it is the z forbidden line. These fits correspond to  $\sim 1$  ms after the LBO injection in the I-mode plasma discussed in the main text. Bottom panels show spectral fitting residuals.

For VUV spectrometers, the lower spectral resolution does not permit advanced fitting procedures and we resort to standard binning of known atomic lines. We combine the brightness of adjacent lines for which ADAS photon emissivity coefficients are available, so as to increase the signal-to-noise ratio. For  $\text{CaF}_2$  LBO injections, we found the (1.8683,1.8727)

nm, (1.9775, 1.9632, 1.9790) nm and (2.0122, 2.0289) nm Li-like Ca groupings to be clearly visible on the XEUS spectrometer [44].

To account for possible systematic errors, we apply a minimum relative uncertainty of  $\pm 5\%$  to all brightness values as in past analysis [31, 51]. For VUV signals, uncertainties are most effectively estimated from signal scatter before LBO injections.

### 3.3. Bayesian Inference

Bayesian analysis methods are commonly applied in several branches of Physics and Engineering, but are relatively less known in fusion research. For this reason, we introduce some key Bayesian concepts in this section. A comprehensive review can be found in Ref. [65].

Consider a framework to infer parameters  $\theta$  (vector symbols are omitted in the following) given experimental data  $\mathcal{D}$  and a model  $\mathcal{M}_i$ . Bayes' formula then reads

$$p(\theta, \gamma | \mathcal{D}, \mathcal{M}_i) = \frac{p(\mathcal{D} | \theta, \gamma, \mathcal{M}_i) p(\theta, \gamma | \mathcal{M}_i)}{\mathcal{Z}(\mathcal{D} | \mathcal{M}_i)}, \quad (5)$$

where

$$\mathcal{Z}(\mathcal{D} | \mathcal{M}_i) = \int p(\theta, \gamma | \mathcal{M}_i) p(\mathcal{D} | \theta, \gamma, \mathcal{M}_i) d\theta d\gamma \quad (6)$$

is the *evidence* (or *marginal likelihood*) of model  $\mathcal{M}_i$ . Here,  $\theta$  is a vector containing the parameters that we wish to infer and  $\gamma$  is another vector containing nuisance parameters, i.e. parameters that are needed by our model evaluation, but whose values are not of physical interest (see section 3.4). Any “choice” of parameters  $\{\theta, \nu\}$  may be referred to as a parameter *sample*. The term  $p(\theta, \gamma | \mathcal{M}_i) \equiv p(\theta, \gamma)$  (a probability density for the parameters, given a specific model) is generally referred to as the prior (see section 3.6), while  $p(\mathcal{D} | \theta, \gamma, \mathcal{M}_i) \equiv \mathcal{L}(\mathcal{D} | \theta)$  is the likelihood. The left hand side,  $p(\theta, \gamma | \mathcal{D}, \mathcal{M}_i) \equiv p(\theta | \mathcal{D})$ , is the posterior. Note that in order to find the parameters  $\theta, \gamma$  that maximize the posterior density function (the “Maximum A-Posteriori”, or MAP, estimate), we can ignore  $\mathcal{Z}$ , since this is not a function of these parameters. However, while the integration required to compute  $\mathcal{Z}$  tends to be a computationally expensive task, its value (or approximations thereof) is the fundamental *quantitative* Bayesian metric to compare how much the data support one mathematical model versus another. The comparison of  $\mathcal{Z}$  values (or their equivalent normalization) is referred to as *model selection*. A complementary task to model selection is *parameter estimation*, which attempts to find the best parameters for a given model. Both model selection and parameter estimation are important for *model validation*, although the former is often not treated in detail due to its inherent complexity.

In the absence of good estimates of the Bayesian evidence, one may resort to some of its approximations, e.g. the Akaike Information Criterion (AIC) or the Bayesian Information Criterion (BIC), which however are derived by making questionable assumptions on the form of the (unknown) posterior distribution [66]. Note that the  $\chi^2$  is a metric describing how accurate signal matching is and does not provide a judgement on the model; in other

words, it corresponds to a likelihood function, rather than an approximation of the evidence. Its normalization by the number of “degrees of freedom”, referred to as *reduced- $\chi^2$* , is a deceptively simple and often unjustifiable metric for model selection [67].

### 3.4. Signal modeling and combination

Every diagnostic has finite spatio-temporal resolution, thus limiting the accuracy and precision of measurements. In attempting to match `pySTRAHL` dynamics and experimental data, we integrate simulation results over time bins representing each instrument’s photon collection time. Moreover, our inferences include a few useful *nuisance* parameters, i.e. free variables that are not of interest *per se*, but which are important in order to reproduce the physical dynamics of impurity injections. For example, we infer corrections to the estimated location of the sawtooth mixing radius and diagnostic time bases (which may not be exactly synchronized otherwise). We also allow the locations of spline knots for  $D$ ,  $v$  profiles to be freely inferred, forcing a small minimum distance between them to prevent exploration of extreme gradients of transport coefficients, which can lead to numerical instabilities. No evidence is observed of Ca recycling and therefore we avoid free parameters corresponding to this effect.

Inspired by the astrophysical literature (e.g. Refs. [66, 68–70]), we adopt a Gaussian likelihood over measurement errors and reconsider the common practice [18, 31] of weighting  $\chi^2$  from different diagnostics with fixed factors as

$$\chi^2 = \sum_k \alpha_k \chi_k^2 \quad (7)$$

where  $k$  is an index identifying an experimental signal to be matched. The  $\alpha_k$  weight is unknown, and commonly fixed by an experimentalist to match expectations on which signal matching should be prioritized. We consider this in a Bayesian light and look for an appropriate prior over  $\alpha_k$ . By a maximum-entropy argument (see Appendix A for details), one may find that an appropriate prior with unit expectation is the exponential  $P(\alpha_k) = \exp(-\alpha_k)$ . In practice, it is often preferable to limit the prior to have a finite width, in order to avoid unreasonable values. While no analytic solution to this constrained entropy maximization is available [69], the gamma distribution (of which the exponential is a special case) fulfills our objectives, while not strictly being a solution to the above problem. As with the exponential prior, we are able to integrate a Gaussian likelihood over such a gamma prior for  $\alpha_k$  analytically (see Appendix A). This allows us to effectively substitute the simple  $\chi^2$  metric with

$$\ln P(D|\theta) = \sum_{k=1} \left[ \ln \Gamma \left( \frac{n_k}{2} + \nu \right) - \left( \frac{n_k}{2} + \nu \right) \ln \left( \frac{\chi_k^2}{2} + \nu \right) - \ln \Gamma(\nu) + \nu \ln \nu \right]. \quad (8)$$

where we defined  $\nu := a = 1/b$  to fix the gamma prior mean  $\langle \alpha_k \rangle = 1$  and we dropped constant factors that do not depend on the inference parameters. This likelihood is analogous

to the typical  $\chi^2$  one, but sums over diagnostic weights that follow a gamma probability distribution with  $\nu := a = 1/b$ . By setting  $\nu$  to different values, we allow  $\alpha_k$  to be more or less free to weight diagnostics differently. Appendix A shows some possible choices of  $\nu$  that we considered; the inference shown in section 4 used  $\nu = 25$ . Note that there is one value of  $\alpha_k$  for each diagnostic  $k$ , but since we analytically marginalized the likelihood over the  $\alpha_k$  prior there is no direct sampling of  $\alpha_k$  to be done: inclusion of these parameters in an inference does not incur into any new free parameters (and thus adds no additional computational cost).

By assigning a prior to the “weight factors”, we allow our algorithm to infer appropriate values based on observed under-estimation of uncertainties, inaccuracies of atomic data, correlated signals, and other uncontrollable issues. Setting  $\nu = 25$  in Eq. 8, as we do in this work, conservatively constraints weights not be different from their first estimates by more than  $\approx 50\%$ .

### 3.5. Nested Sampling

The BITE framework makes use of the Nested Sampling (NS) Monte Carlo method [71] for model selection and parameter estimation. The primary objective of NS is to evaluate the evidence integral in Eq. 6, which can be re-written as

$$\mathcal{Z} = \int \mathcal{L} dX \quad (9)$$

where  $X$  is the prior survival function (referred to as “prior volume”) above a given value of likelihood  $\mathcal{L}$ :

$$X(\lambda) = \int_{\{\theta: \mathcal{L} > \lambda\}} \pi(\theta) d\theta. \quad (10)$$

$X(\lambda)$  may be thought of as the fraction of the prior hyper-volume in which  $\mathcal{L} > \lambda$ . By writing the evidence integral in the form of Eq. 9, one reduces a complex multi-dimensional integral to a 1D integral, which may be computed via standard numerical methods once samples of  $\mathcal{L}$  and  $X$  have been collected.

NS sets out to do so by first obtaining pseudo-random samples within a unit-hypercube having as many dimensions as the inference problem of interest. These samples are passed to a user-defined function that must transform them to the prior space of interest, usually via the inverse Cumulative Distribution Function (CDF) of the specific prior used for each dimension. Parameters are then passed to user-defined routines that run the forward model (in our case, pySTRAHL and synthetic diagnostics) and return a likelihood value. The NS algorithm thus collects the prior and likelihood probabilities for each unit-hypercube sample and orchestrates where future samples should be obtained in order to explore the (unknown) posterior distribution. NS makes use of multiple *live points* that can explore the posterior in parallel. At any iteration, the live point with the lowest likelihood is eliminated from the set (becoming “inactive”), and replaced with a new sample, which is



accepted (becoming “active”) if its likelihood is higher than the one of the eliminated point, or else is rejected. The algorithm therefore progressively explores nested shells of likelihood, reducing the prior volume that is sampled to find the best fitting parameters. Following this sampling procedure, the algorithm “ranks” the obtained values of likelihood based on their prior volume,  $X$ . It then becomes possible to evaluate the 1D integral in Eq. 9, once a tolerance condition for the  $\mathcal{Z}$  estimate is reached. As an important by-product of this process, NS allows one to use the set of live points from all iterations to reconstruct the explored posterior, possibly recovering any multimodal structure (i.e. separate peaks of the posterior distribution). In the “Importance Nested Sampling” (INS) variant, live points that were discarded at every iteration are also used to further improve the estimate of  $\mathcal{Z}$  [72].

The main difficulty in implementing the NS algorithm is the necessity to efficiently draw unbiased samples within iso-likelihood contours in the prior space. The `MultiNest` algorithm [72, 73] does so by fitting (potentially overlapping) ellipsoids to the set of live points and sampling from within their union. Details of the algorithm and `MultiNest` parameters that we choose for transport inferences are given in Appendix B.

### 3.6. Choice of priors

As mentioned above, nested sampling is an inherently Bayesian method. The use of priors, on top of likelihoods, allows us to provide valuable information to the algorithm and avoid unphysical regions, while also demanding care to avoid inappropriate bias. *We adopt the philosophy that all information available to us and not already encapsulated in the likelihood should be carefully put to use via priors.* Our choices are described in this section.

Whenever possible, we avoid the use of uniform priors over a parameter value. These would imply that we have no idea of what is the most likely region of parameter space where true physical solutions lie - something that is often not true. For example, we apply a Gaussian prior, rather than a uniform one, over the sawtooth inversion radius, estimated with  $\sim 1$  cm spatial accuracy via ECE or SXR diagnostics. Adopting a uniform prior over  $D$  would also be a mistake, given that we are interested in exploring diffusion that may go from neoclassical expectations near the magnetic axis (of the order of  $10^{-2}$   $m^2/s$ ) to much larger values, possibly approaching 100  $m^2/s$  at mid-radius, where turbulent transport is dominant. *Uninformative* sampling in this case requires one to apply uniform sampling to  $\log(D)$ ; in the Bayesian literature, this is referred to as a *Jeffreys prior* [74, 75].

Assuming that neoclassical and turbulent processes act in such a way as to eliminate fine structure from radial profiles, we expect  $D$  and  $v$  profiles to vary smoothly as a function of radius. In order to encapsulate this expectation into our priors we make use of *Gaussian copulae*, through which we set correlations between the values of  $D$  and  $v$  (separately) at adjacent spline knots. A Gaussian copula “couples” different sampled parameters,  $\vec{u} = [u_1, u_2, \dots, u_d]$ , from a  $d$ -dimensional unit hypercube. It is defined via the expression

$$C_R^{Gauss}(\vec{u}) := \Phi_R(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d)) \quad (11)$$



where  $R$  is a correlation matrix,  $\Phi_R$  is the joint cumulative distribution function of a multivariate normal distribution with zero mean vector and covariance matrix  $R$ , and  $\Phi^{-1}$  is the inverse cumulative distribution function of a standard normal. The result of the transformations in Eq. 11,  $C_R^{Gauss}(\vec{u})$ , is a set of parameters that are individually (i.e. marginally) uniformly distributed, but not *jointly* uniformly distributed in the unit hypercube. Rather, they present correlations indicated by the correlation matrix  $R$ . We choose  $R$  to be tridiagonal, with diagonal values of 1.0 and off-diagonal entries of 0.5. This makes our expectation of “smoothness” in  $D$  and  $v$  profiles more explicit without necessarily forcing personal bias on the posterior [76].

We choose to apply prior constraints over  $D$ ,  $v/D$  rather than  $D$ ,  $v$ . We find this to be more conservative since  $v/D$  is related to density peaking and is unlikely to take very large values; on the other hand,  $v$  may vary widely so long as a physical  $v/D$  is maintained (see section 5.1). In devices with direct impurity density measurements (e.g. via CER), values of  $v/D$  may be strongly constrained; however, on C-Mod only line-integrated brightness is available and this appears to constrain  $D$  more strongly than  $v/D$ . We therefore impose a relatively restrictive prior on  $v/D$  [ $m^{-1}$ ] at midradius, using a Gaussian distribution with mean of 0 and standard deviation of  $2 m^{-1}$  ( $\mathcal{N}(0, 2)$ ). This allows enough freedom to explore positive and negative peaking within a range that spans measurements on other devices [18, 22, 77, 78] and encompassing observed values of  $1/L_{ne} \sim 1$  at midradius in the I-mode discharge described in this paper. Near-axis, we allow significantly more variation with  $v/D \sim \mathcal{N}(0, 10)$ . In the pedestal, where we have only weak experimental constraints but one may expect large inward pinch convection, we set weak priors on the Gaussian pinch amplitude ( $\sim \mathcal{N}(-100, 50)$ ) and its width ( $\sim \mathcal{N}(0.03, 0.03)$ , in  $r/a$  units)<sup>+</sup>. We also impose that a “non-negligible” fraction of LBO-injected particles should enter the confined plasma in pySTRAHL simulations. Previous work (e.g. Ref. [79]) reported LBO penetration fractions equal or greater than 5-10%, suggesting that one may conservatively set a minimum penetration fraction of 1-2% in BITE. This simple condition makes a large  $D$  ( $\gtrsim 5 m^2/s$ ) and a small inward edge  $v$  ( $|v| \lesssim 50 m/s$ ) at the LCFS likely unrealistic. Finally, in the absence of more detailed knowledge, time base synchronization of experimental signals, being determined by random triggering events in each detector, is set via uniform priors with appropriate bounds for each detector.

#### 4. Experimental Results

The inference of  $D$ ,  $v$  profiles and their uncertainty is strongly dependent on the flexibility of the  $D$  and  $v$  “models” that one applies. In this section, we illustrate in turn the results of model selection and parameter estimation for the I-mode case of interest.

<sup>+</sup> Truncated normal priors are used when appropriate to prevent sampling of unphysical parameters, e.g. negative radial “widths”.

#### 4.1. Model Selection

As described in section 3.3, the Bayesian evidence is the central metric in Bayesian inference to assess what level of model complexity is best supported by data. The Bayesian evidence may be interpreted as a *relative* metric, since the ratio of values of evidence have a rigorous statistical interpretation for model selection (see Ref. [66] for a detailed discussion). In Fig. 6, we show the base-10 logarithm of Bayes Factors (BF), defined as ratios of Bayesian evidence values, for a range of model complexities (number of free parameters). Each BF is relative to the case with highest log-evidence. Cases on the left of the BF peak are under-fitting, i.e. their models are insufficient to represent experimental data adequately; cases on the right are over-fitting, i.e. they require more parameters than experimental uncertainties suggest is reasonable to use. In all cases, we set  $D$  and  $v$  radial knots to be the same for simplicity.

We compare the BF obtained when we set  $v$  to scale linearly with  $Z$  in the pedestal (taken as  $r/a > 0.9$ ) in pySTRAHL (blue) and when we set  $v$  to be independent of  $Z$  (green). From this plot, it is clear that the  $v \sim Z$  model is best supported by data, although the strength of statistical evidence favoring the model-selected case with respect to the next case with the highest BF is relatively weak. However, the difference in predicted posterior distributions between these cases is also small, thus making the distinction between them not particularly significant. On the other hand, the difference between posteriors of cases differing by more than 3 units of  $\log_{10}(BF)$  can be dramatic. None of the cases with high BF were found to be multimodal, i.e. having separate statistical modes of comparable evidence, although the posterior distributions often exhibit non-Gaussian features.

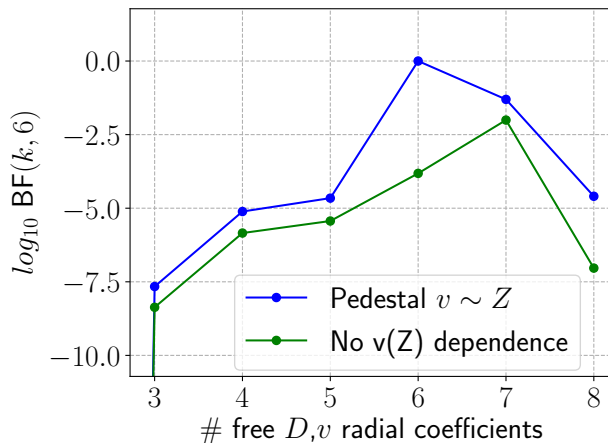


Figure 6: Scaling of the log of the Bayes Factor (BF) with number of free  $D$ ,  $v$  radial coefficients in the analyzed I-mode discharge. The model with 6 free  $D$  and  $v$  radial coefficients and with  $v$  scaling with  $Z$  in the pedestal is consistently, although weakly, favored by experimental data at all dimensionalities.

Obviously, uncertainties on experimental data strongly affect the dimensionality achieving the highest log-evidence (or, equivalently, BF). The importance of carefully

quantifying such uncertainties for challenging model selection problems such as ours cannot be overstated; our detailed approach was described in previous sections.

#### 4.2. Parameter Estimation

Fig. 7 shows the inferred  $D$ ,  $v$  and  $v/D$  profiles from inferences obtained with the three levels of complexity that resulted in the highest BF (cf. Fig. 6), with 5 (blue), 6 (green) and 7 (red) radial coefficients for  $D$  and  $v$ . All of these cases used  $v(r/a > 0.9) \sim Z$ , since this choice is found to give higher BF. In Fig. 7 we display profiles up to the LCFS, showing that a large inward pinch (plotted for fully-stripped Ca) is inferred in the pedestal. The maximum magnitude of such pinch ( $\approx -150$  m/s) is not shown in order to visualize both core and edge profiles on the same figure; its value is not very well constrained by our data and is found to depend on details of the applied model. The apparent existence of a pedestal impurity pinch does not preclude a short impurity confinement time: in this discharge,  $\tau_{imp} \approx 30 \pm 2$  ms, while the energy confinement time  $\tau_e \approx 23 \pm 1$  ms.




Figure 7: Inferred profiles of  $D$  and  $v$ , together with the  $Rv/D$  ratio, for inferences with 5, 6 and 7 radial  $D$  and  $v/D$  coefficients (with  $v$  calculated from their product). The 6-coefficient inference has higher log-evidence and is therefore *model selected*; 5- and 7-coefficient cases are only shown for reference.

The vertical magenta dashed line show the approximate location of the sawteeth inversion radius. Sample signal fits resulting from these transport coefficients are shown in Fig. 1 and Appendix B. We remark that by setting the w and z line signals to be relatively calibrated, we make their fitting much more challenging: virtually no discrepancy between

simulated and experimental signals would be seen if we allowed lines to be independently normalized\*. The uncertainties shown in Fig. 7 are not assumed to be Gaussian and represent the 1-99, 10-90 and 25-75 quantiles of the posterior function. The Bayesian weight factors described in section 3.4 (see also Appendix A) are found to effectively weight XICS signals (for both the w and z lines) more heavily than VUV lines. All signals are found to have effective weights of less than 1, indicating that diagnostic uncertainties are likely too small to have full consistency with the pySTRAHL model. The use of Bayesian weight factors results in an increase of error bars for the profiles in Fig. 7 by approximately 30% at  $\rho_\phi = 0.6$ , compared to the results obtained with fixed diagnostic weights of 1. Setting  $v \sim Z$  in the pedestal is found to lower the absolute magnitude of edge values of  $D$  and  $v$  required to match data, without significantly affecting inferred profiles in the core.

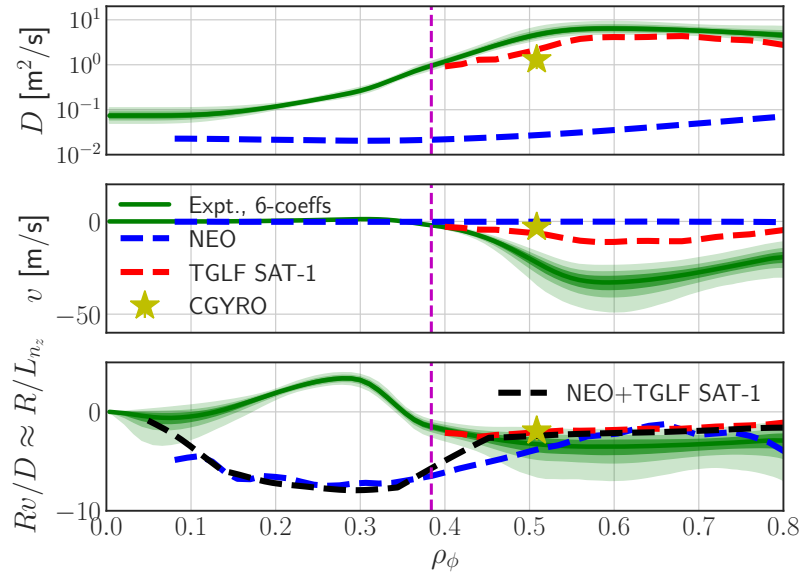


Figure 8: Core-focused view of inferred profiles of  $D$  (on a logarithmic scale) and  $v$ , together with the  $Rv/D$  ratio, for the model-selected inference (also shown in Fig. 7). Neoclassical, gyro-fluid and gyrokinetic modeling are overlaid. The black line in the lowest plot shows an interpolated profile of combined NEO and TGLF predictions for the total  $R v/D$ .

On top of the experimentally-inferred radial profiles of impurity transport coefficients for the model with the highest BF, Fig. 8 shows the results of selected transport models. In blue, we show predictions from the NEO [80, 81] neoclassical transport code and in red from the Trapped Gyro-Landau Fluid (TGLF) reduced quasilinear turbulence model [82, 83]. In the bottom panel, we also show the prediction for the total  $Rv/D$  value (black), taken as an additive measure of neoclassical and turbulent transport (i.e.  $R v/D_{tot} =$

\* This effect is **not** due to inaccurate atomic rates, but to true physical complexity of experimental data.

$R(v_{neo} + v_{turb})/(D_{neo} + D_{turb})$ ). Scales have been adapted in each panel to complement the visualization of Fig. 7 and allow better comparison to transport models.

NEO is a multi-species drift-kinetic code that solves the first-order (in the drift-ordering parameter  $\rho^* = \rho_i/a \ll 1$ ) drift-kinetic-Poisson equations [80, 81]. Since neoclassical predictions are formally limited by near-axis potato orbits and by ion orbit losses near the LCFS [81, 84], we limit our NEO simulations to the range  $0.1 \lesssim \rho_\phi \lesssim 0.95$ . Fig. 8 shows that inferred diffusion values near the magnetic axis are larger, but of the same order of magnitude, as the NEO predictions. The inferred values are, of course, to some extent dependent on the applied sawtooth model in pySTRAHL simulations, but we have found that  $D$  and  $v$  profiles of neoclassical character are always inferred inside the sawtooth inversion radius. We note that this is only an approximate comparison, given that NEO was run with time-averaged kinetic profiles, as opposed to the time-dependent pySTRAHL predictions including multiple impurity charge states and sawtooth modeling. The NEO profiles displayed here were obtained for  $\text{Ca}^{18+}$  (the dominant charge state in this discharge); future work will present more detailed predictions accounting for the presence of multiple charge states.

At mid-radius, where the effect of sawteeth is reduced, TGLF SAT-1 compares favorably with absolute magnitudes of turbulent  $D$ , but differs by more than a factor of 2 from  $v$  (and  $v/D$ ) experimental estimates. Fig. 7 also shows a single-point prediction for  $D$ ,  $v$  and their ratio  $v/D$  from a nonlinear, ion-scale,  $Q_i$ -matched CGYRO [85] simulation at  $r/a = 0.6$ . Both TGLF and CGYRO modeling are described in greater detail in the next section.

In Fig. 9a we show the self-similar impurity density profiles for the highest ionization stages corresponding to the inferred transport coefficients in Fig. 8, normalized such that the median of the total density is 1 on axis. By “self-similar” here we refer to the property of maintaining profile shape while decreasing in magnitude over time after the LBO injection. Density uncertainties were computed from the inferred  $D$ ,  $v$  uncertainties, i.e. by propagating  $D$ ,  $v$  uncertainties using chains and weights from MultiNest. This gives a clear visualization of how uncertainties in Fig. 8 translate to  $n_z$  profiles, retaining all degrees of correlation between parameters.

In Fig. 9b we compare the median fractional abundances for the highest charge states, shown with continuous lines, with those given by only balancing effective ionization and recombination rates, i.e. without any transport. This figure shows that the presence of transport significantly modifies the localization of each charge state. In particular, we note that it allows these highly-ionized charge states to live in the pedestal region, where atomic physics by itself would predict that only lower charge states should exist. This excess of highly-ionized states in the pedestal is not present during the rise phase of the LBO injection, since impurity ions must first penetrate into the core to undergo nearly-complete ionization and then be transported outwards; only at this stage, the upper level of x-ray transitions can be significantly populated via recombination, as previously described by Rice *et al.* [86]. Such recombination effect contributes to the He-like Ca z line emissivity measured via XICS, but not to the corresponding w line, as described in section 3.2, thus providing a non-trivial

constraint to our inference.

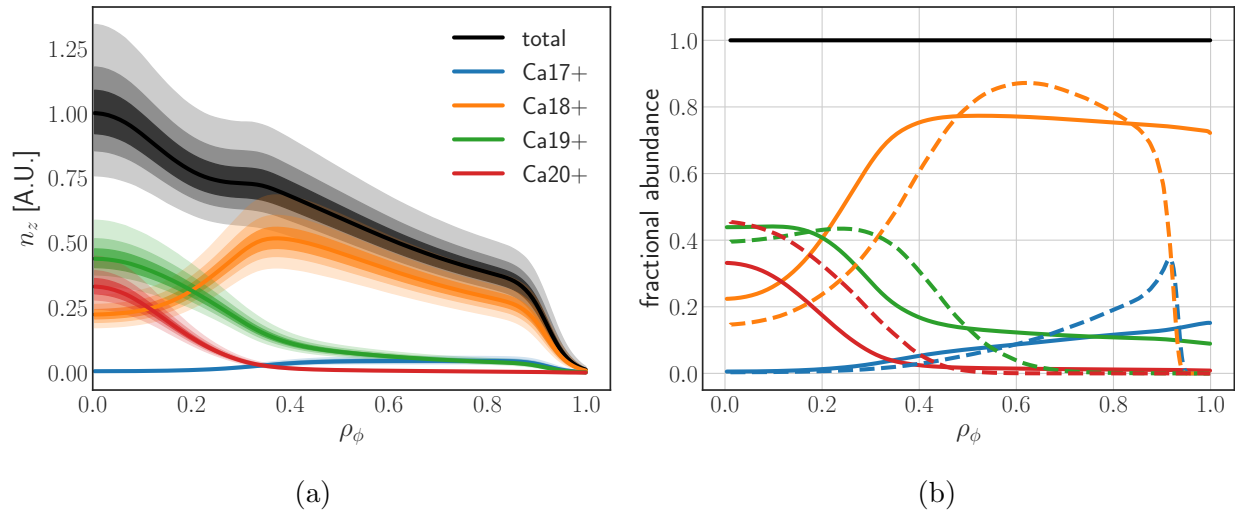


Figure 9: (a) Experimentally-inferred self-decaying profile shapes for the highest ionization stages of Ca after an LBO injection into the I-mode discharge. In (b), we compare the median fractional abundances of (a), shown with continuous lines, with those predicted by ionization equilibrium only (no transport), shown with dashed lines.

## 5. Turbulence Modeling

Having described our attempts at quantifying transport coefficients for the I-mode case of interest, we now proceed to describe in greater depth the turbulence modeling to which these results can be compared. The turbulence modeling in this section is not in any way informed by the experimental results presented above, thus offering the means for an independent validation effort.

In this section, we describe both nonlinear gyrokinetic simulations and quasilinear modeling. The latter is commonly used for integrated modeling, where fast iterations are required to interpret experimental observations or predict plasma behavior. In particular, we make use of TGLF [82, 83], which solves the Trapped Gyro-Landau Fluid equations to approximate the linear eigenmodes of gyrokinetic drift-wave instabilities: trapped ion and electron modes (TIM, TEM), ion and electron temperature gradient modes (ITG, ETG), and kinetic ballooning modes (KBM). Using these, TGLF computes quasilinear estimates to turbulent fields [82, 87], making use of a parametric saturation rule whose coefficients are set to match a database of fully-nonlinear gyrokinetic simulations.

By adding trace Ca impurities to TGLF or CGYRO runs, one may obtain estimates for turbulent diffusion and convection components. This is a valid procedure because *in the trace limit* a species has such a low concentration as to become negligible in the Poisson equation and in Ampère’s Law. Consequently, the turbulent electromagnetic potential

becomes independent of the distribution function of the trace particles and the particle flux becomes *linear* in the thermodynamic gradients [88]:

$$\frac{R\Gamma_z}{n_z} = D_z \frac{R}{L_{n,z}} + D_{T,z} \frac{R}{L_{T,z}} + Rv_{p,z} \quad (12)$$

Here we separated the effect of density and temperature gradients on the impurity flux, and defined a particle diffusion coefficient  $D_z$  (identical in definition to what we obtained in our experimental analysis), a thermodiffusion term  $D_{T,z}$  and a pure convection term  $v_{p,z}$  ( $z$  subscripts will be dropped henceforth). For clarity, we also define  $v_{T,z} := D_{T,z}/L_{T,z}$ , which has units of m/s and is more properly identified as a thermal convection term. Some authors also define a term that is linearly proportional to toroidal rotation gradients (“roto-diffusion”), e.g. Refs. [89, 90]. We shall not consider this since having species with different toroidal rotation is not possible in CGYRO and this prevents us from rigorously studying this term. Roto-diffusion is thus “grouped” in this analysis with other convection terms.

In order to get  $D$ ,  $v_T$  and  $v_p$  estimates from TGLF or CGYRO, we introduce 3 trace species. These may either be all included in the same run, as we do for CGYRO, or in separate runs; the latter is a better option for codes like TGLF whose computational cost scales with the square of the number of species. Relative to the main-ion density, we slightly vary<sup>‡</sup> the input density gradient of the second and third species and, similarly, the input temperature gradient of the third species only. We then collect a set of linear equations for each species’ flux into matrix form

$$\frac{R}{n_z} \begin{pmatrix} \Gamma_{z,1} \\ \Gamma_{z,2} \\ \Gamma_{z,3} \end{pmatrix} = \begin{pmatrix} R/L_{n,z,1} & R/L_{T,z,1} & 1 \\ R/L_{n,z,2} & R/L_{T,z,2} & 1 \\ R/L_{n,z,3} & R/L_{T,z,3} & 1 \end{pmatrix} \begin{pmatrix} D_z \\ v_{T,z} \\ v_{p,z} \end{pmatrix} \quad (13)$$

By inverting the  $3 \times 3$  matrix, one can then solve for the local transport coefficients. An analogous method was also used to obtain the NEO and TGLF  $D$  and  $v$  radial profiles in Fig. 8. In this case, we omitted the temperature perturbation in Eq. 13 and thus only obtained a total convection term (labeled  $v$ ), since thermal convection cannot be separated experimentally by our inference framework. It is nonetheless interesting to distinguish  $v_T$  from  $v_p$  in theoretical modeling to compare to predictions for thermodiffusion in different turbulent regimes [91].

### 5.1. Quasilinear TGLF modeling

In this section, we explore the sensitivity of the radial profiles of impurity transport coefficients predicted by TGLF, shown in Fig. 8. We make use of the SAT-1 “multiscale” saturation rule [92, 93] since this is found to better reproduce kinetic profiles while matching experimental heat fluxes from TRANSP [94] within the TGYRO framework [95].

<sup>‡</sup> The actual percentage variation is inconsequential for trace-level concentrations.



Analysis of the TGLF linear spectrum shows that ion-scale modes with  $k_y \rho_s$  up to  $\approx 0.9$  have real frequencies indicating propagation in the ion diamagnetic drift direction. Here,  $k_y = nq/r$ ,  $n, q$  and  $r$  are the toroidal mode number, safety factor and outboard midplane minor radius, respectively, and  $\rho_s$  is the ion sound speed gyroradius. Together with a strong sensitivity to scans in  $a/L_{Ti}$ , and weaker sensitivity to other typical drift wave drives, linear spectra suggest that ITG modes are strongly dominant in this plasma, in agreement with past I-mode simulations [96,97]. The observation of positive thermal convection,  $v_T > 0$ , also confirms the identification of ITG being dominant [91]. While TGLF and CGYRO spectra are different, particularly at scales where ITG and TEM modes co-exist, they agree in the identification of the most intense turbulent mode being at  $k_y \rho_s \approx 0.4$ .

In Fig. 10 we show independent TGLF scans of  $a/L_{Ti}$ ,  $a\nu_{ei}/c_s$ ,  $a/L_n$  and  $T_i/T_e$  within 2 standard deviations, as estimated from experimental data (cf. Fig. 3). Here,  $c_s$  is the local sound speed and  $\nu_{ei}$  is the electron-ion collision frequency, for which we considered an uncertainty of 20% following Ref. [98]. We use  $2\sigma$  rather than the more conventional  $1\sigma$  in order to encompass  $\approx 95\%$  of possible outcomes, taking all uncertainties to be Gaussian. These scans suggest that  $D$  and the total  $v$  have an overall uncertainty of at least 30% of their




Figure 10: Diffusion and convection coefficients from TGLF  $2\sigma$  scans around experimental I-mode parameters. Uncertainties are estimated locally at each radial location for each parameter.

magnitude. We note, however, that a complete assessment of uncertainties and comparison to the BITE results in Fig. 8 would ideally require varying input parameters simultaneously; correlations across TGLF inputs could be important, but are not available from the 1D scans of Fig. 10. Interestingly,  $v/D$  is almost fixed within the radial range that we explore; in other words, changes of  $D$  and  $v$  compensate such that the expected impurity peaking

indicated by  $v/D$  remains constant. The effect of parallel velocity shear on these results has been found to be negligible. Transport coefficients were found to generally have weaker dependence on  $a/L_{T_e}$  than on  $a/L_{T_i}$  (not shown). However, greater sensitivity to  $a/L_{T_e}$  at outer radii suggests that electron-scale modes grow stronger (and ion-scale modes weaker) as one approaches the top of the pedestal; this is found to correlate with smaller diffusion and convection magnitudes.

## 5.2. Gyrokinetic modeling

We next describe results from nonlinear ion-scale gyrokinetic simulations with CGYRO [85]. We focused our runs on  $r/a = 0.6$ , where experimental profiles are well determined and the effect of sawteeth is expected to be small. We used a domain size with  $L_x/\rho_s \approx 100$  and  $L_y/\rho_s \approx 100$ , radial grid spacing of  $\Delta x/\rho_s = 0.061$  and  $\Delta y/\rho_s = 0.065$ , with 344 radial modes and 22 toroidal modes, giving  $\max(k_x \rho_s) \approx 10.5$  and  $\max(k_y \rho_s) \approx 1.4$ . In velocity space, we use a grid with 16 pitch angles, 8 energy and 24 poloidal points. We adopt experimental profile inputs, Miller geometry, electromagnetic ( $\tilde{\phi}$  and  $\tilde{A}_{||}$ ) effects, as well as gyrokinetic electrons. We apply the sonic rotation scheme recently added to CGYRO [99], whose effects with respect to the linear ordering in Mach number are known to account for significant differences of transport for heavy impurities [100]. Averages and uncertainties on simulation outputs are estimated using a moving average over windows of length equal to 3 times the estimated correlation time.

Sensitivity scans for each of the parameters mentioned above, as well as for numerical dissipation values, were run by increasing each parameter by 50% in turn and checking for significant variations in output. The I-mode discharge of interest has a value of dimensionless electron collision frequency (defined as in Ref. [85]) of  $\bar{\nu}_{ee} \approx 0.1$  at the examined location - a relatively *low*-collisionality case for C-Mod. Linear scans of collisionality did not display strong dependencies in any outputs, but did highlight that this parameter affects particle transport more strongly than heat transport. This may not be surprising, in view of the well-known role of collisionality on electron transport and density peaking [6–8]. Application of the Sugama operator [85, 101, 102], as opposed to pitch-angle scattering via the Lorentz operator, has also been found to strongly affect particle transport predictions.

Nonlinear scans of  $a/L_{T_i}$  were used to find the value at which CGYRO could match the experimental turbulent ion heat flux,  $Q_i$ , as determined by power balance via TRANSP [103]. This was found to be at  $a/L_{T_i} \approx 1.65$ , approximately 25% below the experimental  $a/L_{T_i}$  estimate, for which CGYRO predicts a value of  $Q_i$  that is 6 times higher. As in previous GYRO simulations of I-mode discharges [96], we have found this case to be particularly stiff and close to marginality. This makes an exact  $Q_i$  matching procedure more difficult, since  $Q_i$  varies significantly over time. Consistently with past work [104, 105], we find a clear underestimation of  $Q_e$  at all scanned values of  $a/L_{T_i}$ . Changes in  $a/L_{T_e}$  (by 10%) or in  $E \times B$  shear (by 50%) saw no significant variation of  $Q_e/Q_i$  (less than 10%), which remains less than 1/2 of the value predicted by TRANSP. This is consistent with expectations from linear

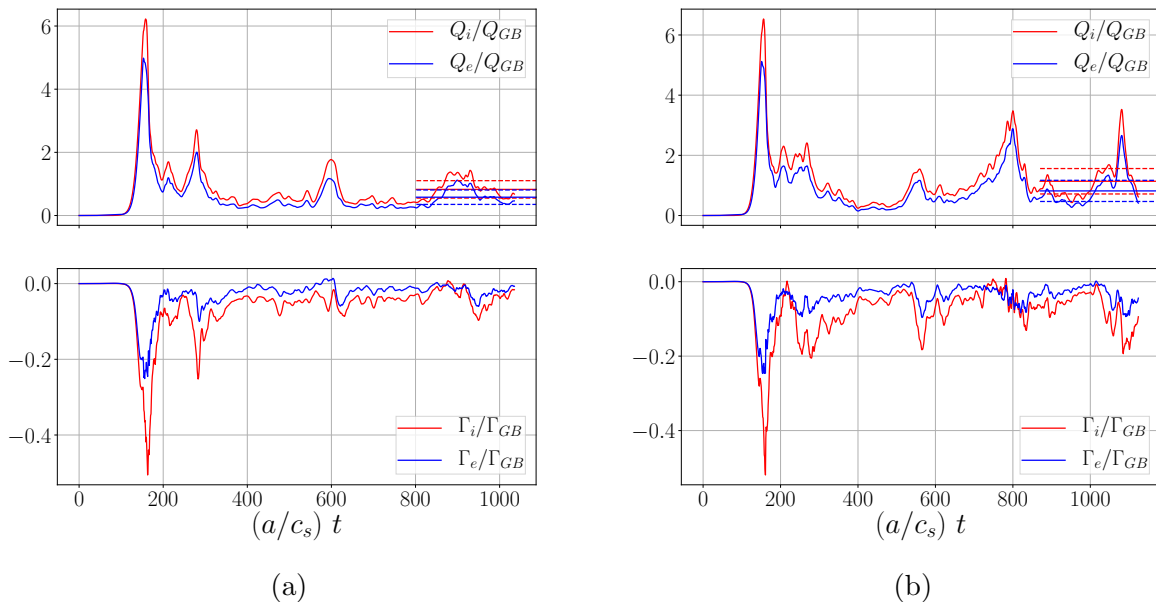


Figure 11: Time evolution of total ion and electron fluxes in nearly  $Q_i$ -matched CGYRO simulations, differing in  $a/L_{T_i}$  by only 0.5%. The case with larger  $a/L_{T_i}$  (right) displays significantly more intermittency. We use a moving average of heat fluxes over the last 25% of each simulation (dashed lines) to compare to experimental values.

CGYRO scans, which show little sensitivity to typical TEM or ETG drives, suggesting that multiscale simulations [93] may still be unable to resolve the observed discrepancies in  $Q_e$ . Nonetheless, given the predominant role of ion-scale fluctuations in determining turbulent impurity transport, these results do not preclude a useful comparison of impurity transport predictions with experiment.

In Fig. 11 we show the time history of heat and particle fluxes for two simulations that differed by only 0.5% in their  $a/L_{T_i}$  values. As a result of changing the temperature gradient by such small amount, turbulent intermittency is seen to vary significantly. This may be expected given the marginality and low heat fluxes in this regime. Horizontal lines in the heat flux plots (top) display the moving average over the last 25% of the simulation, giving  $Q_i^{sim}/Q_{gB} \approx 0.83 \pm 0.27$  and  $1.14 \pm 0.47$  for the simulations with lower  $a/L_{T_i}$  (left) and higher  $a/L_{T_i}$  (right), respectively, where  $Q_{gB} := n_e T_e c_s (\rho_s/a)^2$  is the local gyro-Bohm unit of heat flux. Both of these values are reasonably close to the experimental estimate of  $Q_i^{expt}/Q_{gB} \approx 1.09 \pm 0.3$ , slightly under- or over-predicting it within uncertainty.

In Fig. 12 we show the time evolution of Ca transport coefficients from the simulation in Fig. 11a, again separating pure and thermal convection as in the TGLF simulations shown in section 5.1. Continuous lines are the quantities obtained by the sum of all toroidal modes; dashed lines show results obtained from only the strongest mode, found via linear

simulations to be at  $k_y \rho_s \approx 0.4$ . While the dashed line obviously does not match the total  $D$  in magnitude (continuous line),  $v/D$  is closely matched for each convection component separately; all discrepancies between single-mode and overall traces are due to transient growth of other toroidal modes. Furthermore, crosses at the end of the simulation time domain in the second panel indicate  $v/D$  results obtained from a purely linear simulation at  $k_y \rho_s = 0.4$ . The proximity of these results suggests that quasilinear estimates (such as those obtained via TGLF, see Fig. 8) should be able to closely match fully-nonlinear gyrokinetic predictions for  $v/D$ . On the other hand, these results also show that matching  $D$  and  $v$  components separately between experiment and theory offers a much more stringent test than matching  $v/D$  alone. We note that  $v_T/D$ , the thermal convection peaking component, is positive (as expected in the ITG regime, in agreement with TGLF) and it is smaller than the pure pinch; since  $v_T$  scales with  $1/Z$  [91], this may not be surprising for Ca impurities.




Figure 12: Time evolution of diffusive and convective transport components of trace impurity species over time for the lower- $a/L_{T_i}$  simulation of Fig. 11a. Dashed lines indicate how transport coefficients computed only from  $k_y \rho_s \approx 0.4$  track the total coefficients. Crosses on the right hand side of the lower panel show  $v/D$  predictions from a linear CGYRO simulation at  $k_y \rho_s = 0.4$ .

In Fig. 13 we show spectra of heat and particle fluxes, as well as diffusion and convective components of trace Ca impurities, again for the simulation corresponding to Fig. 11a. Here,  $\Gamma_{gB} := n_e c_s (\rho_s/a)^2$  is the gyro-Bohm unit of particle flux. Interestingly, the inversion of main-ion ( $\Gamma_i$ ) and electron particle flux ( $\Gamma_e$ ) at  $k_y \rho_s \approx 0.35$  is not observed with the Lorentz collision operator, which predicts negative deuterium flux at all ion scales and larger positive electron particle flux at  $k_y \rho_s < 0.4$ . We note a clear similarity between heat flux spectrum

and the impurity diffusion spectrum, consistently with the fact that heat transport is mostly of diffusive nature. Impurity convection terms do not seem to be significant for  $k_y \rho_s \gtrsim 0.5$ , i.e. at electron scales all transport appears to be of diffusive nature.




Figure 13: Wavenumber spectra of heat and particle fluxes for both electrons and (total) ions, as well as of trace impurity diffusion ( $D$ ) and convective components (pure,  $v_p$ , and thermal convection,  $v_T$ ) over the non-linearly simulated range of  $k_y \rho_s \lesssim 1.4$ . These spectra correspond to the lower- $a/L_{T_i}$  case of Fig. 11a.

## 6. Discussion and Summary

In previous sections, we presented the results of experimental inferences of radial profiles of Ca transport coefficients and compared them to predictions from neoclassical and turbulent transport models. The nested sampling algorithm has allowed us to reliably estimate the Bayesian evidence and do model selection for appropriate and flexible parameterizations of transport coefficients and nuisance parameters. Doing so with an inference framework

based on a minimization algorithm, such as the Levenberg-Marquart one, would have required questionable approximations and compromises. Ultimately, a comparison between algorithms would be valuable, but requires the construction of a database of inferences, which will be part of future work.

The complexity of our modeling, both on the experimental and theoretical sides, leaves room for multiple potential sources of error in the validation process. From an experimental perspective, one may particularly expect inaccuracies due to diagnostic time bases synchronization or MHD activity in the plasma core. The use of a 1D model and the omission of thermal charge exchange in the plasma edge may also be problematic. While we cannot easily quantify uncertainties in ADAS effective ionization and recombination rates, we have thoroughly explored the sensitivity of our results to photon emissivity rates. Having obtained similar results with and without the inclusion of the  $\text{Ca}^{18+}$  z line, which depends on different atomic physics from the w line, and having tested our inferences with and without time-dependent atomic rates (modulated by sawteeth in our experiment), we can confidently state that our results are robust to such issue. Since Ca is non-intrinsic and non-recycling, we have been able to avoid significant modeling complications related to SOL transport, thus making our results more robust. Inferred uncertainties have been found not to be strongly sensitive to prior shapes and hyperparameters, except for  $v/D$ , which is only weakly constrained by our data and is therefore effectively “regularized” by our prior. Increasing the  $v/D$  prior width to larger values can indeed lead to inferences with higher peaking. While such solutions cannot be excluded due to our limited measurements, we note that our prior is consistent with expectations of  $R/L_{n_z}$  not being dramatically different from  $R/L_{n_e}$ .

We highlight that our experimental  $D$ ,  $v_{tot}$  uncertainties are constrained by the “model” that we apply, i.e. the number of spline knots, the choice of  $D$ ,  $v_{tot}$  parameterization, the sawtooth crash model, etc.. More physical (e.g., 2D) or more flexible models (not necessarily with more free parameters) would improve uncertainty quantification by overcoming the apparent model inadequacy, particularly in the pedestal. Still, even taking inferred uncertainties at face values, we note that TGLF scans (Fig. 10) have shown sensitivity of modeled  $D$ ,  $v$  values to several inputs, making it plausible to have a closer match between experimental and modeling predictions than at first apparent from Fig. 8.

Interestingly, the application of a Z-dependence to the particle convection in the pedestal region was found to match experimental data better than a model with no Z-dependence for charge states, based on the Bayesian evidence. This observation has been corroborated with inferences on discharges that have not been described in this paper. In all cases, our experimental signals only offer weak constraints in the pedestal, therefore suggesting that even our core-focused measurements may constrain pedestal transport to some extent. Our inferences suggest that the I-mode discharge discussed in this paper has a significant edge pinch, analogously to H-mode discharges. This is compatible with the observation of relatively large edge gradients in  $n_e$  (see Fig. 3), even in the absence of a well-defined

“density pedestal”. As described in section 4, the impurity confinement time is found to be low (comparable to  $\tau_e$ ) because of the simultaneous effect of high diffusion.

Modeling of the I-mode condition of interest with TGLF has shown that experimentally-inferred diffusion can be approximately matched within uncertainties, whereas convection at mid-radius is predicted to be smaller than experimentally inferred. Non-linear ion-scale  $Q_i$ -matched CGYRO simulations predicted smaller  $D, v_{tot}$  values at  $r/a = 0.6$ , thus increasing discrepancies with experiment with respect to TGLF. Although these simulations could not match experimental results quantitatively, we used them to elucidate the role of pure and thermal convection in determining peaking for our I-mode condition. Density peaking, being related to the  $v/D$  ratio, has been found to be well described by the toroidal mode at  $k_y \rho_s = 0.4$ , both in linear and nonlinear simulations. However, separate comparison of  $D$  and  $v_{tot}$ , rather than their ratio, clearly provides a much stronger validation constraint. In section 5.2, we have remarked that different gyrokinetic collision operators can affect particle fluxes much more strongly than heat fluxes. This observation will be the subject of future work.

In summary, this paper showed how BITE addresses many of the issues identified in previous work, particularly bridging the gap between experimental and synthetic data. BITE’s advances range from improved X-ray spectral fitting to inclusion of sawtooth temperature and density modulation, from faster iterative speed using `pySTRAHL` on a computing cluster to improved pedestal modeling. Bayesian model selection was used to avoid under- and over-fitting of experimental signals, allowing us to rigorously estimate diffusion and convection radial profiles. We compared our results to neoclassical NEO and quasilinear turbulent TGLF simulations across the radial extent, as well as to a single CGYRO nonlinear prediction at midradius, showing qualitative (and in some cases quantitative) agreement. The presented methods offer significant advances for data analysis on C-Mod and other devices, but are obviously no substitute for high quality experimental data. Future work will aim at further constraining transport inferences across the plasma, particularly focusing on expanding diagnostic capabilities and physical model fidelity.

## Acknowledgments

We are thankful to the Alcator C-Mod team for outstanding work over the years and to J. Candy, E. Belli, O. Meneghini and S. Smith for helpful conversations on transport modeling. Part of our analysis was performed using the OMFIT integrated modeling framework [106]. This work was directly supported by US DoE grants DE-SC0014264 and made use of the MIT-PSFC partition of the *engaging* cluster at the MGHPCC facility, funded by DoE award No. DE-FG02-91-ER54109, as well as the National Energy Research Scientific Computing Center (NERSC)—a US DoE Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.



# Appendices

## A. Bayesian Signal Combination

The technique described in the main text to combine experimental datasets, e.g. from different plasma diagnostics, is an extension of methods that have been adopted in the astrophysical literature for many years. As in Ref. [69], we look for the distribution of  $\alpha_k$  of maximum entropy

$$S[P(\alpha)] = - \int_0^\infty P(\alpha_k) \ln P(\alpha_k) d\alpha_k$$

under the constraint of normalization  $\int_0^\infty P(\alpha_k) d\alpha_k = 1$  and unit mean  $E[\alpha_k] = \int_0^\infty \alpha_k P(\alpha_k) d\alpha_k = 1$ . This sets our expectation that the appropriate weight  $\alpha_k$  for each diagnostic is 1. Such constrained optimization can be solved analytically using Lagrange multipliers and gives the exponential solution  $P(\alpha_k) = \exp(-\alpha_k)$ , which may be taken to appropriately represent our state of knowledge about  $\alpha_k$ .

As pointed out in Ref. [69], if we were interested in limiting the range of  $\alpha_k$  values to be explored, one could follow the same maximum entropy recipe described above, but now adding an additional constraint on the variance of  $\alpha_k$ . This optimization cannot be analytically solved unless we extend the domain of  $\alpha_k$  to  $[-\infty, +\infty]$ , thus allowing for non-physical weights, to find the familiar Gaussian distribution. This domain extension can be problematic, and we address it by exploiting the fact that the exponential solution found above is a special case of the gamma distribution

$$p(x|a, b) = \frac{1}{\Gamma(a)b^a} x^{a-1} e^{-x/b} \quad (14)$$

for  $a = 1, b = 1$ . Generalizing the maximum-entropy argument above, we therefore adopt a gamma distribution that resembles a “skewed Gaussian” and still obeys the constraints above.

Integrating the product of the prior in Eq. 14 with a Gaussian likelihood, one has

$$P(D|\theta) = \prod_{k=1}^N \frac{1}{(2\pi)^{n_k/2} |V_k|^{1/2} \Gamma(a) b^a} \int_0^\infty \alpha_k^{n_k/2} e^{-\frac{1}{2}\alpha_k \chi_k^2} e^{-\alpha/b} d\alpha \quad (15)$$

where  $V_k$  are the measurement covariance matrices, and  $a, b$  are the gamma distribution parameters of Eq. 14. After some algebra, and dropping constant factors that also appear in the standard Gaussian log-likelihood, this integration directly leads to Eq. 8. This allows one to impose finite variance for the prior over signal weights. Although the gamma distribution is not exactly a maximum-entropy solution, it is a logical extension to achieve our objectives. In order to fix the expectation value of the gamma prior to 1, we set  $\nu := a = 1/b$ . Fig. 14 shows the resulting distribution for a number of choices of  $\nu$ .

One slight drawback of marginalizing over the weights prior is that no single value of the weights is directly inferred; rather, the entire prior distribution is indirectly explored.




Figure 14: Gamma priors over diagnostic weights for some values of the  $\nu$  parameter. All of these choices have a distribution mean of 1, indicated by the vertical black line. Quantiles of each distribution are shown in different colors: 25-75 (red), 10-90 (blue), 1-99 (green).  $\nu = 25$  was used in the inferences described in the main text.

However, one may obtain an estimate for which weights were found to be most likely by differentiating Eq. 15 with respect to  $\alpha_k$  and setting the result to zero. This gives an “effective” weight for each diagnostic

$$\alpha_k^{eff} = \frac{(n_k + 2a - 2)b}{b\chi^2 + 2}. \quad (16)$$

The corresponding effective weights that would be obtained with an exponential prior are

$$\alpha_k^{eff}(\theta) = \frac{n_k}{\chi_k^2(\theta)} \quad (17)$$

Using Eq. 16 in our inferences with  $a = 1/b = 25$ , we usually infer values in the interval  $0.8 < \alpha_k < 1$  for XICS and  $0.6 < \alpha_k < 0.8$  for XEUS. This suggests that wider experimental uncertainties are needed to obtain a Gaussian distribution of data points around the inferred solution, especially for the latter diagnostic, for which background subtraction and signal binning is expected to be less accurate.

We note that, even without marginalizing over a weights prior, one could allow such weights to be inferred as free parameters. This requires weights to be included within the likelihood model, e.g. in the denominator of a Gaussian likelihood, and not only in the exponential. Of course, this increases the problem dimensionality and therefore also the runtime, making the analytic marginalization over signal weights described above more attractive.

## B. Application of MultiNest

As described in section 3.5, nested sampling offers the means to quantify both the Bayesian evidence and the posterior distribution, offering an excellent alternative to Markov Chain

Monte Carlo (MCMC) methods. **MultiNest** is an implementation of NS that fits ellipsoids to the set of live points and samples from within their union. We refer the reader to the Refs. [72, 73] for details of the numerical methods. In this appendix, we describe only key features and choices of important parameters that we made for the present work.

One attractive feature of **MultiNest** is its speed and effectiveness in exploring parameter spaces with up to approximately 30-50 dimensions. While other algorithms, such as **PolyChord** [107], perform better at higher dimensionality, **MultiNest** is the ideal tool for our transport inferences extending to a maximum of 30 free parameters. For standard (non-dynamic) nested sampling, the larger the number of live points ( $n_{live}$ ) used, the lower the chance of missing important parameter space. In our inferences, we vary the number of live points based on the inference dimensionality ( $D$ ) in order to keep the evidence error constant [107], using  $n_{live} = 200 + 25 \times D$ .

**MultiNest** defines a *target efficiency*, or “inverse enlargement factor”,  $f$ , which expands the ellipsoids’ hyper-volume ( $0 < f \leq 1$ ) to avoid over-shrinking at any iteration, at the cost of slower convergence. In our inferences, we conservatively set 0.1 when using vanilla NS and 0.01 for INS. Increasing  $n_{live}$  or decreasing  $f$  has been observed not to affect our results, while obviously incurring higher computational cost.

In its INS variant, **MultiNest** can reach significantly faster convergence, particularly in its “constant-efficiency” mode. This makes INS a convenient choice for dimensionality scans, although **MultiNest**’s ability to isolate posterior modes is only available in vanilla NS. We make use of the latter to identify which experimental measurements and priors are better suited to exclude multi-modality from our impurity transport inferences. Such data-driven approach allows us to understand ambiguities of our data that may lead to unphysical local minima in the  $D$  and  $v$  posterior distributions.

### C. XEUS signal fits

For the sake of completeness, we show the XEUS spectroscopic signal matches from the I-mode case that was described in the main text. Sample signal fits for the XICS Ca w and z lines are shown in Fig. 1. Signal fits for the two VUV groups of lines measured by XEUS (viewing Ca<sup>17+</sup> emission) are shown in Fig. 15.




Figure 15: Signal fits for the XEUS groups of lines for the I-mode case discussed in the main body.

- [1] Pütterich T, Fable E, Dux R, O'Mullane M, Neu R, Siccino M. Determination of the tolerable impurity concentrations in a fusion reactor using a consistent set of cooling factors. *Nuclear Fusion*. 2019;59(5):ab0384. Available from: <https://doi.org/10.1088/1741-4326/ab0384>.
- [2] Reinke ML, Meigs A, Delabie E, Mumgaard R, Reimold F, Potzel S, et al. Expanding the role of impurity spectroscopy for investigating the physics of high-Z dissipative divertors. *Nuclear Materials and Energy*. 2017 8;12:91–99.
- [3] Dunne MG. Impact of impurity seeding and divertor conditions on transitions, pedestal structure and ELMs. *Nuclear Fusion*. 2017;57(2):25002. Available from: <http://dx.doi.org/10.1088/0029-5515/57/2/025002>.
- [4] Holland C. Validation metrics for turbulent plasma transport. *Phys Plasmas*. 2016;23(060901).
- [5] Greenwald M. Verification and validation for magnetic fusion. In: *Physics of Plasmas*. vol. 17; 2010. Available from: <https://doi.org/10.1063/1.3298884>.
- [6] Angioni C, Peeters AG, Pereverzev GV, Ryter F, Tardini G. Density Peaking, Anomalous Pinch, and Collisionality in Tokamak Plasmas. *Physical Review Letters*. 2003 5;90(20):4. Available from: <https://link.aps.org/doi/10.1103/PhysRevLett.90.205003>.
- [7] Angioni C, Fable E, Greenwald M, Maslov M, Peeters AG, Takenaga H, et al. Particle transport in tokamak plasmas, theory and experiment. *Plasma Physics and Controlled Fusion*. 2009;51(12).
- [8] Greenwald M, Angioni C, Hughes JW, Terry J, Weisen H. Density profile peaking in low collisionality H-modes: Comparison of Alcator C-Mod data to ASDEX Upgrade/JET scalings. *Nuclear Fusion*. 2007;47(9).
- [9] Angioni C. Tungsten transport in JET H-mode plasmas in hybrid scenario , experimental observations and modeling. *Nuclear Fusion*. 2014;54(083028).
- [10] Tala T, Nordman H, Salmi A, Bourdelle C, Citrin J, Czarnecka A, et al. Density peaking in JET—determined by fuelling or transport? *Nuclear Fusion*. 2019;59(12):126030. Available from: <https://doi.org/10.1088/1741-4326/ab4248>.
- [11] Wade MR, Hillis DL, Hogan JT, Maingi R, Menon MM, Mahdavi MA, et al. Helium transport and exhaust studies in enhanced confinement regimes in DIII-D. *Physics of Plasmas*. 1995;2(6):2357–2365. Available from: <https://doi.org/10.1063/1.871489>.
- [12] Marmar ES, Cecchi JL, Cohen SA. System for rapid injection of metal atoms into plasmas. *Review of Scientific Instruments*. 1975;46(9):1149–1154.
- [13] Dux R. Impurity transport in ASDEX upgrade. *Fusion Science and Technology*. 2003;44(3):708–715.
- [14] Dux R, Neu R, Peeters AG, Pereverzev G, Mück A, Ryter F, et al. Influence of the heating profile on impurity transport in ASDEX Upgrade. *Plasma Physics and Controlled Fusion*. 2003;45(9):1815–1825.
- [15] Pütterich T, Dux R, Janzer MA, McDermott RM. ELM flushing and impurity transport in the H-mode edge barrier in ASDEX Upgrade. In: *Journal of Nuclear Materials*. vol. 415; 2011. p. 334–339.
- [16] Bruhn C, McDermott RM, Angioni C, Ameres J, Bobkov V, Cavedon M, et al. Erratum: A novel method of studying the core boron transport at ASDEX Upgrade (*Plasma Physics and Controlled Fusion* (2018) 60 (085011) DOI: 10.1088/1361-6587/aac870); 2020. Available from: <https://doi.org/10.1088/1361-6587/ab72e4>.
- [17] Grierson BA, Burrell KH, Nazikian RM, Solomon WM, Garofalo AM, Belli EA, et al. Impurity confinement and transport in high confinement regimes without edge localized modes on DIII-D. *Physics of Plasmas*. 2015;22(5).
- [18] Odstrcil T, Howard N, Sciortino F, Chrystal C, Wilks T, Thome K, et al. The Dependence of the Impurity Transport on the Dominant Turbulent Regime in ELM-y H-mode Discharges. Submitted for publication. 2020;.
- [19] Dux R, Giroud C, Zastrow KD. Impurity transport in internal transport barrier discharges on JET. *Nuclear Fusion*. 2004;44(2):260–264.
- [20] Giroud C, Barnsley R, Buratti P, Coffey IH, Von Hellermann M, Jupén C, et al. Method for experimental determination of Z dependence of impurity transport on JET. *Nuclear Fusion*.

- 2007 4;47(4):313–330. Available from: <http://stacks.iop.org/0029-5515/47/i=4/a=010?key=crossref.6a7ddbf23950c7f219caaf7138d12c69>.
- [21] Parisot T, Guirlet R, Bourdelle C, Garbet X, Dubuit N, Imbeaux F, et al. Experimental impurity transport and theoretical interpretation in a Tore Supra lower-hybrid heated plasma. *Plasma Physics and Controlled Fusion*. 2008;50(5):1–18.
- [22] Villegas D, Guirlet R, Bourdelle C, Garbet X, Hoang GT, Sabot R, et al. Experimental and theoretical study of nickel transport dependence on gradients in Tore Supra. *Nuclear Fusion*. 2014;54:073011.
- [23] Scavino E, Bakos JS, Dux R, Weisen H. Effects of plasma shape on laser blow-off injected impurity transport in TCV. *Plasma Physics and Controlled Fusion*. 2003;45(11):1961–1974.
- [24] Cui ZY, Zhang K, Morita S, Ji XQ, Ding XT, Xu Y, et al. Study of impurity transport in HL-2A ECRH L-mode plasmas with radially different ECRH power depositions. *Nuclear Fusion*. 2018;58(5).
- [25] Geiger B, Wegner T, Beidler CD, Burhenn R, Buttenschön B, Dux R, et al. Observation of anomalous impurity transport during low-density experiments in W7-X with laser blow-off injections of iron. *Nuclear Fusion*. 2019;59(4).
- [26] Henderson SS, Garzotti L, Casson FJ, Dickinson D, O’Mullane M, Patel A, et al. Charge dependence of neoclassical and turbulent transport of light impurities on MAST. *Plasma Physics and Controlled Fusion*. 2015;57(9):1–17.
- [27] Delgado-Aparicio L, Stutman D, Tritz K, Finkenthal M, Kaye S, Bell R, et al. Impurity transport studies in NSTX neutral beam heated H-mode plasmas. *Nuclear Fusion*. 2009;49(8).
- [28] Rice JE, Goetz JA, Granetz RS, Greenwald MJ, Hubbard AE, Hutchinson IH, et al. Impurity toroidal rotation and transport in Alcator C-Mod ohmic high confinement mode plasmas. *Physics of Plasmas*. 2000 5;7(5):1825–1830. Available from: <http://aip.scitation.org/doi/10.1063/1.874004>.
- [29] Pedersen TS, Granetz RS, Hubbard AE, Hutchinson IH, Marmor ES, Rice JE, et al. Radial impurity transport in the H mode transport barrier region in Alcator C-Mod. *Nuclear Fusion*. 2000;40(10):1795–1804.
- [30] Rice JE, Terry JL, Marmor ES, Granetz RS, Greenwald MJ, Hubbard AE, et al. Impurity transport in alcator C-Mod plasmas. *Fusion Science and Technology*. 2007;51(3):357–368.
- [31] Howard NT, Greenwald M, Mikkelsen DR, Reinke ML, White AE, Ernst D, et al. Quantitative comparison of experimental impurity transport with nonlinear gyrokinetic simulation in an Alcator C-Mod L-mode plasma. *Nuclear Fusion*. 2012;52(6).
- [32] Chilenski MA, Greenwald MJ, Marzouk YM, Rice JE, White AE. Efficient design and verification of diagnostics for impurity transport experiments. *Review of Scientific Instruments*. 2018;89(1).
- [33] Chilenski M, Greenwald MJ, Marzouk Y, Rice JE, White AE. On the importance of model selection when inferring impurity transport coefficient profiles. *Plasma Physics and Controlled Fusion*. 2019;61(12):125012. Available from: <http://dx.doi.org/10.1088/1361-6587/ab4e69>.
- [34] Morrison RE, Oliver TA, Moser RD. Representing model inadequacy: A stochastic operator approach. *SIAM-ASA Journal on Uncertainty Quantification*. 2018 4;6(2):457–496.
- [35] Hutchinson IH, Boivin R, Bombarda F, Bonoli P, Fairfax S, Fiore C, et al. First results from Alcator-C-MOD. *Physics of Plasmas*. 1994;1(5):1511–1518.
- [36] Greenwald M, Bader A, Baek S, Bakhtiari M, Barnard H, Beck W, et al. 20 years of research on the Alcator C-Mod tokamak. *Physics of Plasmas*. 2014;21(11).
- [37] Howard NT, Greenwald M, Rice JE. Characterization of impurity confinement on Alcator C-Mod using a multi-pulse laser blow-off system. *Review of Scientific Instruments*. 2011;82(3):0–5.
- [38] Howard NT, White AE, Greenwald M, Holland C, Candy J, Rice JE. Impurity transport, turbulence transitions and intrinsic rotation in Alcator C-Mod plasmas. *Plasma Physics and Controlled Fusion*. 2014;56(12).
- [39] Rice JE, Reinke ML, Ashbourn JMA, Gao C, Victora MM, Chilenski MA, et al. X-ray observations of Ca19 +, Ca18 + and satellites from Alcator C-Mod tokamak plasmas. *Journal of Physics B: Atomic, Molecular and Optical Physics*. 2014;47(7).
- [40] Rice JE, Reinke ML, Ashbourn JMA, Gao C, Bitter M, Delgado-Aparicio L, et al. X-ray observations

- of medium Z H- and He-like ions with satellites from C-Mod tokamak plasmas. *Journal of Physics B: Atomic, Molecular and Optical Physics*. 2015;48(14):1–8. Available from: <http://dx.doi.org/10.1088/0953-4075/48/14/144013>.
- [41] Reinke ML, Podpaly YA, Bitter M, Hutchinson IH, Rice JE, Delgado-Aparicio L, et al. X-ray imaging crystal spectroscopy for use in plasma transport research. *Review of Scientific Instruments*. 2012;83(113504).
- [42] Rice JE, Gao C, Reinke ML, Diamond PH, Howard NT, Sun HJ, et al. Non-local heat transport, rotation reversals and up/down impurity density asymmetries in Alcator C-Mod ohmic L-mode plasmas. *Nuclear Fusion*. 2013;53(3).
- [43] Reinke ML, Podpaly Y, Gao C, Science P. Operation and Validation of The HIREXSR Analysis COde MIT-Plasma Science and Fusion Center Alcator C-Mod; 2013.
- [44] Reinke ML, Beiersdorfer P, Howard NT, Magee EW, Podpaly Y, Rice JE, et al. Vacuum ultraviolet impurity spectroscopy on the Alcator C-Mod tokamak. In: *Review of Scientific Instruments*. vol. 81; 2010. .
- [45] Lepson JK, Beiersdorfer P, Kaita R, Majeski R, Boyle D. Responsivity calibration of the LoWEUS spectrometer. In: *Review of Scientific Instruments*. vol. 87. AIP Publishing LLC; 2016. Available from: <http://aip.scitation.org/doi/10.1063/1.4960729>.
- [46] Sauter O, Angioni C, Lin-Liu YR. Erratum: Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime (*Physics of Plasmas* (1999) 6 (2834)); 2002.
- [47] Fiore CL, Ernst DR, Rice JE, Zhurovich K, Basse N, Bonoli PT, et al. Internal transport barriers in Alcator C-Mod. *Fusion Science and Technology*. 2006;51(3):303–316.
- [48] Angioni C, Weisen H, Kardaun OJWF, Maslov M, Zabolotsky A, Fuchs C, et al. Scaling of density peaking in H-mode plasmas based on a combined database of AUG and JET observations. *Nuclear Fusion*. 2007;47(9):1326–1335.
- [49] Whyte DG, Hubbard AE, Hughes JW, Lipschultz B, Rice JE, Marmor ES, et al. I-mode: An H-mode energy confinement regime with L-mode particle transport in Alcator C-Mod. *Nuclear Fusion*. 2010;50(10).
- [50] Rice JE, Reinke ML, Gao C, Howard NT, Chilenski MA, Delgado-Aparicio L, et al. Core impurity transport in Alcator C-Mod L-, I- and H-mode plasmas. *Nuclear Fusion*. 2015;55(3).
- [51] Chilenski MA. *Experimental Data Analysis Techniques for Validation of Tokamak Impurity Transport Simulations*. MIT; 2017.
- [52] Chilenski MA, Greenwald M, Marzouk Y, Howard NT, White AE, Rice JE, et al. Improved profile fitting and quantification of uncertainty in experimental measurements of impurity transport coefficients using Gaussian process regression. *Nuclear Fusion*. 2015;55(2):23012. Available from: <http://dx.doi.org/10.1088/0029-5515/55/2/023012>.
- [53] Churchill RM, Theiler C, Lipschultz B, Hutchinson IH, Reinke ML, Whyte D, et al. Poloidal asymmetries in edge transport barriers. *Physics of Plasmas*. 2015;22(5).
- [54] Bishop CM. *Pattern Recognition and Machine Learning*; 2006.
- [55] Summers HP, O’Mullane MG. Atomic data and modelling for fusion: The ADAS Project. In: *AIP Conference Proceedings*. vol. 1344; 2011. p. 179–187.
- [56] Reinke ML, Hutchinson IH, Rice JE, Howard NT, Bader A, Wukitch S, et al. Poloidal variation of high-Z impurity density due to hydrogen minority ion cyclotron resonance heating on Alcator C-Mod. *Plasma Physics and Controlled Fusion*. 2012;54(4).
- [57] Odstrcil T, Pütterich T, Angioni C, Bilato R, Gude A, Odstrcil M. The physics of W transport illuminated by recent progress in W density diagnostics at ASDEX Upgrade. *Plasma Physics and Controlled Fusion*. 2018;60(1).
- [58] Dux R. *STRAHL User Manual*; 2006.
- [59] Fritsch FN, Carlson RE. Monotone Piecewise Cubic Interpolation. *SIAM Journal on Numerical Analysis*. 1980 4;17(2):238–246. Available from: <http://epubs.siam.org/doi/10.1137/0717021>.



- [60] Helander P. Neoclassical transport in a rotating impure plasma. *Physics of Plasmas*. 1998;5(4):1209–1211.
- [61] Seguin FH, Petrasso R, Marmor ES. Effects of Internal Disruptions on Impurity Transport in Tokamaks. *Physical Review Letters*. 1983 8;51(6):455–458. Available from: <https://link.aps.org/doi/10.1103/PhysRevLett.51.455>.
- [62] Cao NM, Sciortino F. Bayesian Spectral Moment Estimation and Uncertainty Quantification. *IEEE Transactions on Plasma Science*. 2019;p. 1–9.
- [63] Mewe R, Schrijer J. Helium-like Ion Line Intensities. *Astronomy and Astrophysics*. 1978;65:99–114.
- [64] Gabriel AH. Dielectronic Satellite Spectra for Highly-Charged Helium-Like Ion Lines. *Monthly Notices of the Royal Astronomical Society*. 1972;160(1):99–119. Available from: <https://academic.oup.com/mnras/article/160/1/99/2604664>.
- [65] Von Toussaint U. Bayesian inference in physics. *Reviews of Modern Physics*. 2011 9;83(3):943–999. Available from: <https://journals.aps.org/rmp/abstract/10.1103/RevModPhys.83.943>.
- [66] Trotta R. Bayes in the sky: Bayesian inference and model selection in cosmology. vol. 49. Taylor & Francis; 2008.
- [67] Andrae R, Schulze-Hartung T, Melchior P. Dos and don'ts of reduced chi-squared. arXiv preprint. 2010 12; Available from: <http://arxiv.org/abs/1012.3754>.
- [68] Lahav O, Bridle SL, Hobson MP, Lasenby AN, Sodr e L. Bayesian 'hyper-parameters' approach to joint estimation: The Hubble constant from CMB measurements. *Monthly Notices of the Royal Astronomical Society*. 2000;315(4):45–49.
- [69] Hobson MP, Bridle SL, Lahav O. Combining cosmological data sets: Hyperparameters and Bayesian evidence. *Monthly Notices of the Royal Astronomical Society*. 2002;335(2):377–388.
- [70] Marshall P, Rajguru N, Slosar A. Bayesian evidence as a tool for comparing datasets. *Physical Review D - Particles, Fields, Gravitation and Cosmology*. 2006;73(6):1–4.
- [71] Skilling J. Nested Sampling. *AIP Conference Proceedings*. 2004;395(2004):395–405.
- [72] Feroz F, Hobson MP, Cameron E, Pettitt AN. Importance Nested Sampling and the MultiNest Algorithm. arXiv. 2013;p. 1–28. Available from: <http://arxiv.org/abs/1306.2144>.
- [73] Feroz F, Hobson MP. Multimodal nested sampling: An efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses. *Monthly Notices of the Royal Astronomical Society*. 2008;384(2):449–463.
- [74] Jeffreys H. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London Series A, Mathematical and physical sciences*. 1946 9;186(1007):453–461. Available from: <http://www.royalsocietypublishing.org/doi/10.1098/rspa.1946.0056>.
- [75] Jaynes ET. Prior Probabilities. *IEEE Transactions on Systems Science and Cybernetics*. 1968;4(3):227–241.
- [76] Nelsen RB. Introduction. In: *An Introduction to Copulas*. 2nd ed. Springer New York; 2007. p. 1–5.
- [77] Angioni C, McDermott RM, Fable E, Fischer R, P utterich T, Rytter F, et al. Gyrokinetic modelling of electron and boron density profiles of H-mode plasmas in ASDEX Upgrade. *Nuclear Fusion*. 2011;51(2).
- [78] Manas P, Camenen Y, Benkadda S, Weisen H, Angioni C, Casson FJ, et al. Gyrokinetic modeling of impurity peaking in JET H-mode plasmas. *Physics of Plasmas*. 2017;24(6). Available from: <http://dx.doi.org/10.1063/1.4985330>.
- [79] Rice JE, Terry JL, Marmor ES, Bombarda F. X ray observations of up-down impurity density asymmetries in Alcator C-Mod plasmas. *Nuclear Fusion*. 1997 2;37(2):241–249. Available from: <http://stacks.iop.org/0029-5515/37/i=2/a=I09?key=crossref.bd024016ef25402e8136f710c804303b>.
- [80] Belli EA, Candy J. Kinetic calculation of neoclassical transport including self-consistent electron and impurity dynamics. *Plasma Physics and Controlled Fusion*. 2008;50(9):0–32.
- [81] Belli EA, Candy J. Full linearized Fokker-Planck collisions in neoclassical transport simulations. *Plasma Physics and Controlled Fusion*. 2012 1;54(1):015015. Available from: <http://stacks.iop>.

- [org/0741-3335/54/i=1/a=015015?key=crossref.824bcf2d11d8b9b8849cb24daddca98e](https://doi.org/10.1088/1741-4326/54/i=1/a=015015?key=crossref.824bcf2d11d8b9b8849cb24daddca98e).
- [82] Staebler GM, Kinsey JE, Waltz RE. A theory-based transport model with comprehensive physics. In: *Physics of Plasmas*. vol. 14; 2007. .
- [83] Staebler GM, Howard NT, Candy J, Holland C. A model of the saturation of coupled electron and ion scale gyrokinetic turbulence. *Nuclear Fusion*. 2017;57(6).
- [84] Helander P, Sigmar DJ. *Collisional transport in magnetized plasmas*. Cambridge University Press; 2010. Available from: [https://books.google.com/books/about/Collisional\\_Transport\\_in\\_Magnetized\\_Plas.html?id=nm-V01E0H2MC](https://books.google.com/books/about/Collisional_Transport_in_Magnetized_Plas.html?id=nm-V01E0H2MC).
- [85] Candy J, Belli EA, Bravenec RV. A high-accuracy Eulerian gyrokinetic solver for collisional plasmas. *Journal of Computational Physics*. 2016;324:73–93.
- [86] Rice JE, Graf MA, Terry JL, Marmor ES, Giesing K, Bombarda F. X-ray observations of helium-like scandium from the alcator c-mod tokamak. *Journal of Physics B: Atomic, Molecular and Optical Physics*. 1995;28(5):893–905.
- [87] Staebler GM, Kinsey JE, Waltz RE. Gyro-Landau fluid equations for trapped and passing particles. *Physics of Plasmas*. 2005;12(10):1–24.
- [88] Angioni C, Peeters AG, Pereverzev GV, Bottino A, Candy J, Dux R, et al. Gyrokinetic simulations of impurity, He ash and  $\alpha$  particle transport and consequences on ITER transport modelling. *Nuclear Fusion*. 2009 5;49(5):055013. Available from: <http://stacks.iop.org/0029-5515/49/i=5/a=055013?key=crossref.90ec31be31bc8d5c2dcf353f6c196869>.
- [89] Angioni C, Sertoli M, Bilato R, Bobkov V, Loarte A, Ochoukov R, et al. A comparison of the impact of central ECRH and central ICRH on the tungsten behaviour in ASDEX Upgrade H-mode plasmas. *Nuclear Fusion*. 2017;57(5):aa6453. Available from: <https://doi.org/10.1088/1741-4326/aa6453>.
- [90] Casson FJ, Peeters AG, Angioni C, Camenen Y, Hornsby WA, Snodin AP, et al. Gyrokinetic simulations including the centrifugal force in a rotating tokamak plasma. *Physics of Plasmas*. 2010;17(10).
- [91] Angioni C, Camenen Y, Casson FJ, Fable E, McDermott RM, Peeters AG, et al. Off-diagonal particle and toroidal momentum transport: A survey of experimental, theoretical and modelling aspects; 2012.
- [92] Staebler GM, Candy J, Howard NT, Holland C. The role of zonal flows in the saturation of multi-scale gyrokinetic turbulence. *Physics of Plasmas*. 2016;23(6).
- [93] Howard NT, Holland C, White AE, Greenwald M, Candy J, Creely AJ. Multi-scale gyrokinetic simulations: Comparison with experiment and implications for predicting turbulence and transport. *Physics of Plasmas*. 2016;23(5).
- [94] Breslau J, Gorelenkova M, Poli F, Sachdev J, Yuan X, Science UOo. TRANSP; 2018. Available from: <https://www.osti.gov/biblio/1489900>.
- [95] Candy J, Holland C, Waltz RE, Fahey MR, Belli E. Tokamak profile prediction using direct gyrokinetic and neoclassical simulation. *Physics of Plasmas*. 2009;16(6). Available from: <https://doi.org/10.1063/1.3167820>.
- [96] White AE, Howard NT, Creely AJ, Chilenski MA, Greenwald M, Hubbard AE, et al. Nonlinear gyrokinetic simulations of the I-mode high confinement regime and comparisons with experiment. *Physics of Plasmas*. 2015;22(5).
- [97] Fu XR, Horton W, Bepamyatnov IO, Rowan WL, Benkadda S, Fiore CL, et al. Turbulent impurity transport modeling for Alcator C-Mod. *Journal of Plasma Physics*. 2013;79(5):837–846.
- [98] Sung C, White AE, Mikkelsen DR, Greenwald M, Holland C, Howard NT, et al. Quantitative comparison of electron temperature fluctuations to nonlinear gyrokinetic simulations in C-Mod Ohmic L-mode discharges. *Physics of Plasmas*. 2016;23(4). Available from: <https://doi.org/10.1063/1.4945620>.
- [99] Belli EA, Candy J. Impact of centrifugal drifts on ion turbulent transport. *Phys Plasmas*. 2018;032301(December 2017). Available from: <http://dx.doi.org/10.1063/1.5020298>.

- [100] Casson FJ, Angioni C, Belli EA, Bilato R, Mantica P, Odstrcil T, et al. Theoretical description of heavy impurity transport and its application to the modelling of tungsten in JET and ASDEX upgrade. *Plasma Physics and Controlled Fusion*. 2015;57(1):14031. Available from: <http://dx.doi.org/10.1088/0741-3335/57/1/014031>.
- [101] Sugama H, Watanabe TH, Nunami M. Linearized model collision operators for multiple ion species plasmas and gyrokinetic entropy balance equations. *Physics of Plasmas*. 2009;16(11):112503.
- [102] Belli EA, Candy J. Implications of advanced collision operators for gyrokinetic simulation. *Plasma Physics and Controlled Fusion*. 2017;59(4).
- [103] Hawryluk RJ. An empirical approach to tokamak transport. In: *Physics of Plasmas Close to Thermonuclear Conditions*. Pergamon; 1981. p. 19–46. Available from: <https://www.sciencedirect.com/science/article/pii/B9781483283852500091>.
- [104] Mikkelsen DR, Howard NT, White AE, Creely AJ. Verification of gene and gyro with l-mode and i-mode plasmas in alcator c-mod. *Physics of Plasmas*. 2018 4;25(4).
- [105] Creely AJ, Howard NT, Rodriguez-Fernandez P, Cao N, Hubbard AE, Hughes JW, et al. Validation of nonlinear gyrokinetic simulations of L- and I-mode plasmas on Alcator C-Mod. *Physics of Plasmas*. 2017;24(5).
- [106] Meneghini O, Smith SP, Lao LL, Izacard O, Ren Q, Park JM, et al. Integrated modeling applications for tokamak experiments with OMFIT. *Nuclear Fusion*. 2015 8;55(8):083008. Available from: <http://stacks.iop.org/0029-5515/55/i=8/a=083008?key=crossref.5f4846d96e96da6689b641740716977c>.
- [107] Handley WJ, Hobson MP, Lasenby AN. POLYCHORD: Nested sampling for cosmology. *Monthly Notices of the Royal Astronomical Society: Letters*. 2015;450(1):L61–L65.