

**Experimental Inference of Particle Transport in Tokamak
Plasmas**

by

Francesco Sciortino

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2021

© Massachusetts Institute of Technology 2021. All rights reserved.

Author
Department of Physics
August 13, 2021

Certified by
Earl S. Marmor
Senior Research Scientist
Thesis Supervisor

Accepted by
Depto Chakrabarty
Associate Department Head of the Physics Department

Experimental Inference of Particle Transport in Tokamak Plasmas

by

Francesco Sciortino

Submitted to the Department of Physics
on August 13, 2021, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

The achievement of sustainable operation in tokamaks depends crucially on accurate understanding of fuel, impurity, and neutral particle dynamics. In this work, radial profiles of experimental particle transport coefficients have been inferred following laser blow-off (LBO) impurity injections into both Alcator C-Mod and DIII-D plasmas. Development of the Aurora modeling package has supported the creation of Bayesian frameworks that leverage a wide range of spectroscopic diagnostics. This investigation spans regimes without Edge-Localized Modes including Enhanced D-Alpha H-mode and I-mode on C-Mod, and diverted negative triangularity on DIII-D. On C-Mod, a novel forward model for the entire Ca K_α spectrum has been combined with Extreme Ultra-Violet (EUV) spectroscopy of multiple charge states. On DIII-D, analogous EUV measurements complement Soft X-Ray and Charge Exchange Recombination spectroscopy. While the impact of Charge eXchange (CX) between impurities and background neutrals from heating beams is found to be relatively small, edge neutrals are shown to be extremely important for ionization balance and radiation in the pedestal region. This conclusion is supported by SOLPS-ITER simulations, shown to compare favorably to a database of Ly $_\alpha$ measurements near the C-Mod midplane. We find neoclassical, gyrofluid, and nonlinear gyrokinetic modeling to be in relatively good agreement with experimental estimates of diffusion, whereas significant discrepancies in convection are evident in several cases. In particular, experimental observations of hollow impurity profiles often cannot be reproduced by microturbulence models within uncertainties, suggesting that current transport codes may be missing critical physics for impurity peaking predictions of future devices. As a whole, this work provides one of the highest-fidelity assessments of cross-field impurity transport in tokamaks, offering the means to extend comparisons between theory and experiments in the particle transport channel.

Thesis Supervisor: Earl S. Marmor
Title: Senior Research Scientist

Acknowledgments

I have been truly privileged to spend five years at MIT and experience what an incredible place it is. During this time, I have enjoyed tremendous support and freedom, for which I want first to thank my supervisors, Earl Marmor and Nathan Howard. Their advice and encouragement have kept me motivated to contribute to the quest for fusion energy, which first inspired me during my undergraduate studies at Imperial College and EPFL. Particularly over the past year, during the COVID-19 pandemic, I have deeply appreciated my Zoom conversations with Earl and Nathan, sometimes touching on research and bears in Vermont, telescopes in New York, and our attempts to keep optimistic during difficult times. Even though the MIT Physics Department rules prevent the name of a scientist who is not “senior” from being at the front of a Ph.D. thesis, I want to acknowledge Nathan’s direct supervision and mentorship over my years at MIT.

I also want to thank John Rice for his guidance and for sharing his contagious enthusiasm for spectroscopy. John is certainly to blame for me going deeper and deeper into XICS analysis over time. I have also learnt a lot from Tomáš Odstrčil. In his role as a postdoc and later as GA Staff Scientist, Tomáš has taught me a ridiculous number of Python tricks and also kept me humble, showing me the complexities of experimental analysis. He has been a key ally in the development of Aurora and the ImpRad module, neither of which would have been half as impressive without his help.

Orso Meneghini and Sterling Smith immensely helped me (and many others) along the journey – their philosophy of open software and collaborative science has not gone lost with me. Youssef Marzouk has been an excellent teacher on Bayesian inference (for the record, I regret not taking more classes taught by him). I have also been fortunate to work with Adam Foster, without whom my latest work on K_α spectra would not have been possible. With Norman Cao, I found a friend and colleague of incredible sharpness. I am proud of our joint efforts in setting up the successful “Turbulence & Transport” journal club that has been meeting for the past few years at the the MIT PSFC.

Over time, I have been greatly helped by Shaun Haskey, who allowed me to use his FIDASIM tools, and by Jeremy Lore, who contributed to getting me started with SOLPS-ITER and shared some of his simulations with me. Colin Chrystal and Kathreen Thome have also been central to enabling the CER analysis on which I based my DIII-D work.

Igor Bykov provided fast camera data, without which my DIII-D modeling would have been more challenging and less accurate. I am also grateful to Alessandro Marinoni, Max Austin, and Carlos Paz-Soldan for leading the recent negative triangularity campaigns at DIII-D and letting me be part of the experimental team.

Over the past year, I have had the pleasure of joining the Friday meetings of the “MIT-W&M Neutrals Lounge”. For these, I want to thank Richard Reksoatmodjo, Jerry Hughes, Saskia Mordijck, Aaron Rosenthal, Abhi Mathews, and Andrés Miller. I have also greatly appreciated Jim Terry’s encouragement to delve into edge physics, and Matt Reinke’s help and inspiration to work on spectroscopy. Martin O’Mullane helped me navigate the intricacies of ADAS, and I am grateful for his advice, especially since my questions must have displayed great ignorance at times. For years, I have been looking up to the work of Thomas Pütterich, Ralph Dux, and Rachel McDermott, whose spectroscopy and modeling expertise I have benefited from at several points of my academic journey. I am looking forward to working more with them at MPI-IPP.

Pablo Rodriguez-Fernandez and Lucio Milanese have helped me grow scientifically, but most importantly they have been great friends. I have many more to thank at MIT, and I particularly want to mention Julian Picard, Kevin Montes, Bryan Linehan, Libby Tolman, Alex Tinguely, Alex Creely, and Will McCarthy. I will always be impressed at the intellectual curiosity and sense of friendship of Matteo Sbroscia, who reviewed and attempted to make sense of this thesis. Daniel Rodan Legrain and Nicolás Riva have helped me to keep up to the very last sprint.

During the past five years, I have greatly enjoyed doing outreach at the PSFC and volunteering for the Cambridge Community Learning Center. In the good and the bad times, it is Beth who has kept me going with her incredible patience and love. While practicing her *active listening* with me, she inadvertently got me to tell her way too much about tokamaks . . . multiple times. She even read this entire thesis, in a display of support (and masochism?) that I am truly amazed by. I am looking forward to moving to Europe together and getting ever closer to the wonderful Hadleys.

It has been almost two years since I last saw my family in Italy. My parents have supported and encouraged me in my journeys ever since I first left my hometown at age 17. My sister, Anna, has been the first and most inspirational figure of my life. Naturally, my thesis is dedicated to her.

Contents

1	Introduction	21
1.1	The Pursuit of Nuclear Fusion	21
1.2	Magnetic Confinement Fusion	24
1.2.1	The Alcator C-Mod Tokamak	26
1.2.2	The DIII-D Tokamak	28
1.2.3	Plasma Transport in Tokamaks	30
1.3	Summary of Primary Results	33
1.4	Organization of This Thesis	34
2	Particle Transport in Tokamaks	37
2.1	Kinetic Theory	38
2.2	Classical Collisional Transport	39
2.3	Neoclassical Transport	41
2.4	Turbulent Transport	44
2.4.1	The Gyrokinetic Equation	44
2.4.2	Gyro-Fluid Theory	47
2.4.3	The Quasilinear Approximation	48
2.4.4	Characterization of Turbulent Modes in Tokamaks	52
2.5	Magneto-Hydrodynamic Modes	56
2.6	Scrape-off-Layer Transport	57
2.7	Experimental Determination of Particle Transport	59
3	Principles of Plasma Spectroscopy	63
3.1	Fundamental Atomic Processes in Plasmas	65
3.2	Ionization Equilibrium	68

3.3	Generalized Collisional Radiative Coefficients	70
3.4	Radiated Power	72
3.5	X-Ray Spectroscopy	75
3.5.1	K_α Spectra	77
4	Particle Transport Diagnostics	83
4.1	Electron Density and Temperature Diagnostics	85
4.2	Ultra-Violet Spectroscopy	86
4.3	Charge Exchange Recombination Spectroscopy	89
4.4	X-Ray Crystal Spectroscopy	91
4.5	Soft X-Ray Arrays	93
5	Bayesian Inference of Particle Transport	95
5.1	Key Concepts of Bayesian Inference	97
5.2	Nested Sampling	100
5.3	Likelihood Model	102
5.3.1	The Aurora package	104
5.3.2	Treatment of the Plasma Background	108
5.3.3	Automatic Rescaling of Synthetic Signals	110
5.4	Prior Distributions	111
5.4.1	Free spline knots	113
5.4.2	Physically-correlated D - v - v/D priors	115
5.4.3	Combination of Constraints from Multiple Diagnostics	117
5.5	On The Use of Manufactured Solutions	119
6	Neutral Particle Studies	121
6.1	Introduction	121
6.2	Comparison of D Neutral Measurements and SOLPS-ITER Predictions with EIRENE	123
6.3	C-Mod Deuterium Ly_α Database	129
6.3.1	Database Development and Methodology	129
6.3.2	Neutral Density Regressions	133
6.3.3	Neutral Penetration Depth	137

6.3.4	Separatrix Electron density	138
6.3.5	Cross-Field Deuterium Pedestal Transport	139
6.3.6	Effective Deuterium Transport Coefficients in the Pedestal	143
6.3.7	Summary of Ly_α database study	145
6.4	Impact of Charge Exchange of Neutrals with Heavy Ions	146
6.5	Charge Exchange with NBI Neutrals	149
6.6	Discussion and Summary	153
7	Inference of Impurity Transport on Alcator C-Mod	155
7.1	Particle Transport Across Confinement Regimes	156
7.2	Inferences of Impurity Transport on C-Mod	159
7.2.1	Review of Previous Efforts	159
7.2.2	Experimental Setup	160
7.2.3	Inferences Using Individual Line Brightnesses	166
7.2.4	Constraints from Line Ratios	173
7.2.5	Inferences using the full K_α spectrum	176
7.3	Comparison of Experimental and Theoretical Modeling	189
7.4	Discussion and Summary	201
8	Inference of Impurity Transport on DIII-D	203
8.1	The OMFIT ImpRad Module	205
8.2	Experimental Data and Analysis Methods	206
8.3	Impurity Transport in DIII-D Discharges with RMPs	215
8.3.1	Inferences of Impurity Transport	217
8.3.2	Comparison of Experimental and Theoretical Modeling	219
8.4	Diverted Negative Triangularity	223
8.4.1	DIII-D Experiments in Diverted Negative Triangularity	225
8.4.2	Impurity Transport Inferences	229
8.4.3	Comparison of Experimental and Theoretical Modeling	235
8.5	Discussion and Summary	242
9	Summary and Conclusions	245
9.1	Considerations For Effective Particle Transport Inferences	247

9.2	Future Work	248
A	The Aurora package	251
A.1	Atomic Data and Spectral Predictions	252
A.2	Aurora Recycling Model	253
A.3	Superstaging	256
A.4	Charge Exchange in the ITER Pedestal	263
B	The Bayesian Spectral Fitting Code	267
B.1	Hermite Polynomial Decomposition	268
B.2	Sampling Methods	271
B.3	Applications to XICS Spectra	272
B.4	Summary	275
C	Bayesian combination of experimental datasets	277
C.1	Bayesian Priors for Diagnostic Weights	279
C.2	Generalization of Constrained Weights Prior	283
C.3	Summary	286
D	Application of Nested Sampling	287
E	Neutral dynamics in Alcator C-Mod	289
E.1	Penetration of Edge Neutrals	289
E.2	Experimental Analysis of D Ly α Signals	290
E.3	High-n CX X-Ray Measurements	293
F	Kinetic Neutral 1D Modeling	299

List of Figures

1-1	Maxwellian velocity averages of some fusion cross sections.	23
1-2	Layouts of illustrative magnetic flux surfaces for Alcator C-Mod, DIII-D, SPARC and ITER with respective vacuum vessels.	27
1-3	Engineering rendering of the Alcator C-Mod tokamak.	28
1-4	Photographs of the interior and exterior of the DIII-D tokamak.	29
1-5	Concept of laser blow-off injections into a tokamak plasma.	31
2-1	Idealized picture of density-driven drift waves in a plasma slab.	53
2-2	Sketch of major turbulent modes leading to transport in tokamak core plasmas.	54
3-1	Relaxation times for a number of ions as a function of electron temperature at fixed electron density.	70
3-2	Cooling factors over electron temperature for a variety of ions of interest for fusion modeling.	75
3-3	Grotrian diagram showing the energy levels (not to scale) and transitions leading to K_α spectra.	79
3-4	Local He-like Ca spectrum at $n_e = 10^{20} \text{ m}^{-3}$ and (a) $T_e = T_i = 1 \text{ keV}$, and (b) $T_e = T_i = 3.5 \text{ keV}$	80
3-5	Local He-like Ca spectrum at $n_e = 10^{20} \text{ m}^{-3}$ and $T_e = T_i = 1 \text{ keV}$ on a log scale, showing many satellite lines contributions.	81
4-1	Electromagnetic spectrum, illustrating regions of interest for various tokamak diagnostics discussed in this chapter.	84
4-2	Lines of sight of the Alcator C-Mod XEUS and LoWEUS spectrometers, as well as for the spatially-discretized 32 chords of the XICS crystal used for Ca measurements.	87

4-3	VUV spectrum from the C-Mod LoWEUS diagnostic.	88
4-4	Magnetic equilibrium reconstruction of DIII-D negative triangularity discharge 180526 showing the location of localized measurements by the CER and Thomson Scattering diagnostics.	91
4-5	Tomographic reconstruction of SXR radiation from the DIII-D diverted negative triangularity discharge #180526.	93
5-1	Overview of the nested sampling approach taken for the inference of impurity transport for both Alcator C-Mod and DIII-D plasmas.	99
5-2	Overview of Aurora’s 1.5D ion transport forward model.	106
5-3	Physical sampling of diffusion (D), convection (v) and their ratio (v/D) for the case of $\chi = 1$, $\mathcal{P}_{\chi v/D} = \mathcal{N}(0, 5)$ and $\mathcal{P}_r = \mathcal{N}(0, 10)$	116
6-1	Rates for hydrogen (a) ionization and (b) recombination, and Ly_α photon emissivity coefficients for (c) ionization-driven and (d) recombination-driven processes.	122
6-2	Spatial distribution of atomic D neutral density from SOLPS-ITER modeling of the Alcator C-Mod discharge #1080416025.	126
6-3	Comparison of SOLPS-ITER (EIRENE) predictions for neutral midplane radial profiles with Ly_α spectroscopy in C-Mod L-mode, EDA H-mode, and I-mode shots.	127
6-4	Radial profile of the normalized effective ionization rate ν_{ion}^*	129
6-5	Histograms of some key parameters within the C-Mod Ly_α database.	131
6-6	Regression for the atomic D neutral pressure measured at the C-Mod midplane in terms of midplane, lower-divertor, and upper-divertor wall pressures.	134
6-7	Model-selected regression for the atomic D neutral pressure measured at the C-Mod midplane in terms of midplane wall pressure, $p_{D2,\text{mid}}$, in L-mode discharges without cryopump operation.	136
6-8	Model selected regression for the neutral penetration depth, L_{n_n} , in terms of $T_{e,\text{sep}}$, the electron temperature at the separatrix, in L-mode plasmas.	137

6-9	Model-selected regressions for the separatrix electron density, $n_{e,sep}$, for (a) L-mode, and (b) H-mode discharges. In (a), dependencies on midplane wall pressure, $p_{D2,mid}$, B_t value, and plasma current, I_p , are identified. In (b), a dependency on B_t is omitted because there is little variation of this variable within the current H-mode database.	140
6-10	Ionization per photon ratio for the Ly_α transition as a function of electron temperature for multiple values of electron density.	141
6-11	Regression of C-Mod deuterium flux through the LCFS for USN (unfavorable ion ∇B drift direction, open divertor) discharges.	142
6-12	Regression of C-Mod deuterium flux through the LCFS for LSN (favorable ion ∇B drift direction, closed divertor) discharges.	142
6-13	Pedestal effective deuterium diffusivity as a function of electron density gradient scale length.	143
6-14	Effective pedestal transport coefficients for deuterium, inferred from Ly_α measurements at the outboard midplane, separating observations in L- and H-mode.	144
6-15	Illustration of the impact of edge neutrals on Alcator C-Mod for identical Aurora simulations in an I-mode discharge (#1080416025) with a constant Ar source and realistic D and v profiles.	147
6-16	Charge state densities and line radiated power for Aurora simulations with and without CX in the DIII-D L-mode discharge (#180533).	147
6-17	Radial profiles of (a) atomic D neutral density and (b) its normalization by n_e in the DIII-D diverted negative triangularity discharge #180533.	148
6-18	Time history of NBI power for the DIII-D discharge #180526.	150
6-19	Density of the full-energy fast neutral beam component incident from the right hand (low-field) side, displayed over the poloidal cross section of the DIII-D negative triangularity discharge #180526.	151
6-20	Flux surface averaged (FSA) density components from FIDASIM for the DIII-D discharge #180526, separating full-, half- and third-energy components of fast neutrals and thermal (halo) contributions. Terms arising from the $n = 1$ (ground state) and excited states $n = 2$ and $n = 3$ are separated.	152

6-21	Radial profiles of recombination for F^{9+} ions due to CX with high energy and thermal (halo) neutrals from the DIII-D 30L beam in shot #180526.	152
6-22	Illustration of the total impact of CX between NBI neutrals and low-Z ions. .	153
7-1	Charge exchange recombination measurements from D gas puffing experiments on Alcator C-Mod, showing (intrinsic) boron density and temperature across an I-L transition.	157
7-2	Scatter plot of electron density peaking $n_e(\psi_n = 0.2)/\langle n_e \rangle$ vs. the peaking prediction by the Angioni scaling for I-, L- and H-mode shots on Alcator C-Mod.	158
7-3	Time histories of radio-frequency power, near-axis electron temperature, radiated power, and line-averaged electron density for the C-Mod L-mode, EDA H-mode and I-mode discharges.	162
7-4	Time-dependent kinetic profile fits for the investigated C-Mod I-mode discharge.	163
7-5	Kinetic profiles and normalized gradient scale lengths for the 3 Alcator C-Mod discharges where impurity transport was inferred.	164
7-6	BSFC fit of the Ca Ly_α lines and overlapping satellites.	167
7-7	Scaling of the log of the Bayes Factor with number of free D and v radial coefficients in the analyzed I-mode discharge.	168
7-8	Inferred profiles of D and v , together with the Rv/D ratio, for I-mode inferences with 5, 6, and 7 radial D and v/D coefficients.	170
7-9	Core-focused view of inferred profiles of D (on a logarithmic scale) and v , together with the Rv/D ratio, for the model-selected I-mode inference. Neoclassical, gyro-fluid and gyrokinetic modeling are overlaid.	171
7-10	Experimentally-inferred self-decaying profile shapes and comparison to fractional abundances at ionization equilibrium for the highest ionization stages of Ca after an LBO injection into the I-mode discharge.	172
7-11	Radial profiles of the w, z, x, y, and k+j relative line emissivities for the I-mode discharge.	175
7-12	Comparison of inferences obtained on synthetic data created from model D and v profiles.	179

7-13	Verification of close match of synthetic XICS signals and a reconstruction via the full-spectrum inference framework for the H-mode discharge.	180
7-14	Inference of Ca particle transport coefficients in the C-Mod I-mode discharge #1101014030, showing the distribution of free knots that were sampled with highest posterior probability via MultiNest.	183
7-15	Comparison of experimental Ca transport coefficients inferred using the full K_α spectrum in the C-Mod I-mode discharge.	184
7-16	Comparison of experimental Ca transport coefficients inferred using the full K_α spectrum in the C-Mod L-mode discharge.	185
7-17	Comparison of experimental Ca transport coefficients inferred using the full K_α spectrum in the C-Mod EDA H-mode discharge.	186
7-18	Steady-state radial density profiles of the five highest Ca charge states from the experimental EDA H-mode inference of transport coefficients.	187
7-19	188
7-20	Scan of Z in NEO simulations for the C-Mod EDA H-mode discharge at $\rho_p = 0.97$	191
7-21	Experimental fractional uncertainties of normalized kinetic gradient scale lengths (for n_e , T_e , and T_i) and the ratio of T_i and T_e for the 3 discharges of interest.	192
7-22	Diffusion and convection coefficients from TGLF 2σ scans around experimental I-mode parameters.	193
7-23	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%.	195
7-24	CGYRO spectra for L-mode (left), EDA H-mode (center), and I-mode (right) shots.	197
7-25	Time histories of heat and particle fluxes (for D ions and electrons) and transport coefficients (D and v/D components) from ion-scale heat-flux matched nonlinear CGYRO simulations for the L-mode (left), EDA H-mode (center), and I-mode (right) shots.	199
7-26	Estimates of Rv/D components from nonlinear CGYRO simulations (pure convection, thermodiffusion and total) and from experimental inferences for the Alcator C-Mod L-mode, EDA H-mode, and I-mode discharges.	200

8-1	ImpRad tab for the selection of background ion species and neutral densities.	205
8-2	Effective source functions for LBO injections of F (for shot #180526) and Al (for shots #180520 and #180530) from the spatial integration of fast camera unfiltered signals.	207
8-3	CER signals from F^{9+} following the LBO injection into shot #180526.	208
8-4	Comparison of SPRED measurements at $\lambda \approx 230 \text{ \AA}$ with the CER signal at the inner-most chord, T01.	211
8-5	Inferred Al and C transport coefficients for DIII-D shot #175674.	217
8-6	Comparison of Al transport coefficients from an experimental inference and from theoretical modeling for DIII-D shot #175674.	218
8-7	Match of experimental and modelled CER signals in the DIII-D RMP ELM-suppressed discharge #175674 for time-dependent Al injected via LBO and intrinsic and time-independent C.	220
8-8	Summary of TGYRO results for DIII-D shot #175674.	221
8-9	Heuristic picture of trapped particle motion in negative and positive triangularity.	224
8-10	Evolution and comparison of the 3 investigated diverted negative triangularity DIII-D discharges discussed.	226
8-11	Kinetic profiles for the 3 diverted negative triangularity discharges.	228
8-12	Signals from the CER T06 chord in #180520 (Al) and 180526 (F), showing exponential fits in the decay phase to quantify impurity confinement times.	229
8-13	Inferred C and F transport coefficients for DIII-D shot #180526.	230
8-14	Inferred F transport coefficients for DIII-D shot #180526.	231
8-15	Inferred Al transport coefficients for DIII-D shot #180520.	232
8-16	Inferred Al transport coefficients for DIII-D shot #180530.	233
8-17	Match of experimental and modelled CER signals in the DIII-D RMP ELM-suppressed discharge #180526 for time-dependent Al injected via LBO and intrinsic and time-independent C.	234
8-18	Match of experimental and modelled signals for SPRED F lines in the diverted negative triangularity discharge #180526.	235
8-19	Match of experimental and modelled F signals for SXR chords in the diverted negative triangularity discharge #180526.	236

8-20	Summaries of TGYRO runs for the three negative triangularity discharges examined in Chapter 8.	237
8-21	TGLF spectra at $\rho_p = 0.6$ and $\rho_p = 0.85$ from scans of major turbulence drives in the diverted negative triangularity discharge #180520.	238
8-22	TGLF spectra at $\rho_p = 0.6$ and $\rho_p = 0.85$ from scans of major turbulence drives in the diverted negative triangularity discharge #180530.	239
8-23	Diffusion and convection coefficients from TGLF 2σ scans around experimental parameters for the DIII-D negative triangularity discharge #180526 at $t = 2.75$ s (low-power L-mode phase).	240
A-1	Full ionization equilibria of (top) W and (bottom) Ca over electron temperature, T_e , and comparison to an example charge states bundling. . . .	260
A-2	Comparison of Aurora simulations for Fe transport in the ITER baseline scenario and an Alcator C-Mod discharge.	261
A-3	Application of superstaging within an Aurora simulation of W in the ITER baseline scenario.	262
A-4	Radial profiles of atomic neutral D density, normalized by electron density, in the ITER baseline from a SOLPS-ITER simulation.	264
B-1	Gauss-Hermite functions of order $n = 0 - 5$	269
B-2	BSFC fit of the He-like Ca z line, overlapping with the j satellite line, which is related to the nearby k line.	272
B-3	Comparison between true synthetic line-integrated measurements and measurements inferred from BSFC.	273
B-4	Comparison between line-integrated measurements inferred from BSFC and from THACO for (a) a subset of XICS channels following Ar injection into an L-mode plasma, and (b) a Ca injection into an I-mode plasma with hollow brightness profiles.	274
C-1	Gamma priors over diagnostic weights for some values of the ν parameter. . .	285
E-1	Fractional abundance of the neutral stage of D over the local electron density in typical Alcator C-Mod (a) L-mode and (b) H-mode discharges.	290

E-2	Ly $_{\alpha}$ analysis of Alcator C-Mod discharges #1100308004, #1100305023, and #1080416025.	292
E-3	Spectra between $n = 11$ and the ionization limit for plasmas with $\langle n_e \rangle = 0.8 \times 10^{20} \text{ m}^{-3}$ and $T_e(r = 0) = 3.2 \text{ keV}$, and for plasmas with $\langle n_e \rangle = 1.4 \times 10^{20} \text{ m}^{-3}$ and $T_e(r = 0) = 1.2 \text{ keV}$	294
E-4	Distribution of atomic D neutrals across the poloidal cross section in the L-mode shot that was compared to Ar high-n Rydberg series measurements.	296
E-5	Poloidal distributions of neutral density at multiple radial locations inside of the separatrix.	297
F-1	Comparison of KN1D results for DIII-D discharge 175674 using two values of edge D $_2$ pressure.	301

List of Tables

3.1	Electronic configurations for transitions in He-like spectra.	78
3.2	Wavelengths and energies of the resonance, intercombination, and forbidden lines for Ar and Ca as measured on Alcator C-Mod.	78
7.1	VUV line identification used for the analysis of LBO injections of Ca on Alcator C-Mod.	176
8.1	VUV line identification used for the analysis of LBO injections of F on DIII-D.	212
8.2	VUV line identification used for the analysis of LBO injections of Al on DIII-D.	213
E.1	Up/down neutral (n_n) and impurity density (n_I) asymmetries relative to X– point location and ion $\mathbf{B} \times \nabla \mathbf{B}$ drift direction.	295

1

Introduction

“We don’t want to be complacent about how audacious this is, but we have to be audacious if we’re going to meet the need.”

– M. Greenwald, February 2021 [1]


1.1 The Pursuit of Nuclear Fusion

Few scientific and technological challenges that have captured human imagination as much as nuclear fusion. The prospect of virtually infinite, safe, and sustainable energy could power much of human ambition for the 21st century and beyond. Fusion, the energy source of the stars, could be central to our reduction of fossil fuel consumption, adaptation to climate change, and a fairer distribution of resources. It may be essential to support increasing energy consumption in developing countries and permit deep-space travel. It is hard to overstate the impact that fusion could have on our society.

The first efforts to harness the power of the stars began after World War II with the development of the hydrogen (H-) bomb, which combined the destructive power of nuclear fission with the one of fusion. In the former process, heavy nuclei are split, with a fraction of the reactants’ initial mass being transformed into kinetic energy, as commonly explained via the $E = mc^2$ formula. In fusion, the opposite occurs: light nuclei merge into heavier ones, with the mass defect before and after the reaction once again released as energy - typically a significantly larger value per unit mass than in fission [2]. Just as the Manhattan Project rapidly led to fission power stations, in the United States and elsewhere, in the

1950's scientists and engineers started working towards controlling the vast energy release from fusion in new reactors. Approximately 70 years down the road, we appreciate that the fusion challenge is more complex than initially expected. Nonetheless, the promise of fusion energy is so enticing and far-reaching that it motivates continuous development, particularly at a time when new experiments, both in public and private ventures, are promising to substantially advance the field. While fusion could play a role in a variety of fields in the future, the applications most directly addressed in this thesis are those related to electric power generation. In this context, the need for diversification of production means calls for investment in basic and applied research on new energy sources. A recent report developed under the auspices of the U.S. National Academies of Sciences, Engineering and Medicine recommended that, for fusion to make an impact on the transition to a low-carbon emission electrical system by 2050, the U.S. should aggressively pursue the creation of a fusion pilot power plant in the 2035-2040 time frame [3].

The fusion reactions that are most promising for energy production are different from the “proton chain” that powers stars like the Sun. In particular [4],



are the simplest processes to produce at sufficiently large scale on Earth. Here, D stands for deuterium (${}^2\text{H}$) and T for tritium (${}^3\text{H}$), both heavy isotopes of hydrogen; n stands for a neutron and p a proton; ${}^3\text{He}$ and ${}^4\text{He}$ are helium isotopes. Fig. 1-1 shows the cross sections for these reactions, multiplied by velocity and averaged over a Maxwellian energy distribution. D-T operation clearly produces much higher fusion rates than other species mixes and is therefore expected to give highest performance in the first fusion power plants. Unfortunately, tritium's radioactivity makes operation significantly more complex and expensive. For this reason, D-D operation is normally favored for scientific experiments. In the reactions above, helium is often referred to as “fusion ash” and must be eliminated from the plasma to avoid fuel dilution, although not before slowing down on (exchanging




Figure 1-1: Maxwellian velocity averages of some fusion cross sections. Data from Bosch & Hale [5].

energy with) D and T ions. Neutrons are not electromagnetically confined since they have no net electric charge. Hence, their energy can be harvested outside the spatial region where fusion occurs, allowing for conversion into electricity not unlike in a fission reactor.

In order for fusion to be a viable source of energy, matter must be first heated to the point of overcoming the binding energy of atoms. The state of matter in which a gas is ionized, creating a cloud of ions and electrons that behave in a collective manner, is referred to as *plasma*. Each particle in a fully-ionized plasma is electrically charged and therefore responds to electromagnetic fields at a distance. It is in the very nature of plasmas to produce emergent behavior, such as turbulence. Plasma physics can therefore be seen as *out-of-equilibrium electromagnetic statistical mechanics*. The ubiquity of turbulence in plasmas is central to many of the challenges in fusion research. This thesis explores properties of thermonuclear plasmas, particularly the way turbulence and collisions affect the motion of particles.

Based on a power balance assessment, fusion reactors operating with a 50-50 D-T mixture are expected to have an optimal operating temperature of approximately 15 keV (more than 150 million degrees Kelvin). This condition, far hotter than the center of the Sun (≈ 1 keV), has already been reached in several experiments. The larger the number of reactants at these high temperatures, the more neutrons can be produced and converted into thermal and electric energy; hence, high density is also necessary. Finally, the plasma must be efficiently confined, so that the energy used to heat the fuel is not dispersed too rapidly.

These principles are concisely captured by the “triple product” inequality [6, 7]:

$$\hat{n} \hat{T} \tau_E \gtrsim 5 \times 10^{21} \text{ keV s / m}^3. \quad (1.1)$$

This describes an approximate condition for the energy from fusion reactions to exceed losses (*ignition*), with \hat{n} being peak density, \hat{T} peak temperature, and τ a measure of energy confinement time (assuming parabolic plasma profiles). The following section describes the principle of magnetic confinement, which has come the closest to overcoming this limit.

1.2 Magnetic Confinement Fusion

The first fusion devices in the 1940s and 1950s attempted to compress (*pinch*) the plasma to reach thermonuclear conditions, but rapid losses from such machines came to be seen as inescapable. Alternative approaches based on confining plasma in linear devices (e.g. in *magnetic mirrors*) or compressing them via high-power lasers (as in *inertial confinement fusion*) have also been explored, but at the time of writing the concept of *magnetic confinement* is at the forefront of the race for fusion energy.

Magnetic confinement was first explored in *tokamaks* (Sakharov and Tamm, 1950) and *stellarators* (Spitzer, 1951), toroidal devices where twisted field lines form a magnetic cage for the plasma [8]. In tokamaks, this is made possible by combining external electromagnets producing a toroidal magnetic field and a transformer that creates a toroidal current (and hence, by Ampère’s Law, a poloidal field). In a stellarator, both toroidal and poloidal fields are externally imposed via twisted magnets. This means that the large internal currents of tokamaks are not needed, resulting in significant advantages for stability and steady-state operation. However, the construction of stellarators poses a number of engineering challenges and their confinement properties have so far proven to be inferior to those of tokamaks. While research into stellarators remains very active, the first fusion pilot plants are expected to use the tokamak design. This thesis focuses on some of the persistent challenges associated with tokamak confinement, examined on two modern-day devices.

Since tokamaks are toroidally symmetric, it is often possible to consider plasma dynamics in 2 rather than 3 dimensions. Depending on the context, one may want the 2 dimensions to describe the Cartesian laboratory frame or the coordinate system described by directions parallel and perpendicular to magnetic field lines. The latter option is useful to describe

plasma dynamics, since particle motion parallel to field lines is mostly unimpeded, while Lorentz forces limit perpendicular motion. Ideally, particles should therefore only move along the twisted field lines if these close upon themselves within the device. The occurrence of particle collisions gives an important deviation from this rule. The “classical theory of transport” shows that this effect produces relatively small cross-field diffusion (see Section 2.2). However, experimental observations in tokamaks always show much higher transport levels. Part of this is due to the complex effects of toroidal geometry on collisional transport, described by “neoclassical theory” (Section 2.3), but the majority of transport is attributable to plasma microturbulence. The latter contribution has historically been dubbed “anomalous”, to highlight its deviation from early understanding of transport processes. As we will discuss in Chapter 2, significant improvements have been made over time in understanding plasma turbulence, although predictive capabilities are clearly incomplete. In this thesis, collisional and turbulence modeling of tokamak regimes is challenged by detailed analysis of experimental particle transport.

In an attempt to compare observations from experimental devices worldwide, a number of “scaling laws” have been obtained, for example the IPB98(y,2) energy confinement time [9, 10, 11]

$$\tau_{\text{IPB98}(y,2)} = 0.0562 I_p^{0.93} B_T^{0.15} P_{\text{loss}}^{-0.69} n^{0.41} M^{0.19} R^{1.97} \epsilon^{0.58} \kappa_a^{0.78}, \quad (1.2)$$

also produced in dimensionless form [10, 11]

$$\tau_{\text{IPB98}(y,2)} \propto \tau_B \rho_*^{-0.70} \beta^{-0.90} \nu_*^{-0.01} M^{0.96} q^{-3.0} \epsilon^{0.73} \kappa_a^{2.3}. \quad (1.3)$$

These formulae relate a macroscopic, global quantity such as the energy confinement time with either engineering parameters, as in Eq. 1.2, or with physically-relevant parameters, as in Eq. 1.3. In Eq. 1.2, I_p is the plasma current [MA], B_T the toroidal magnetic field [T] at the geometric axis R [m], P_{loss} the loss power [MW], n line average density [10^{19} m^{-3}], M the average ion mass (AMU), ϵ the inverse aspect ratio (a/R), a the minor radius [m], and κ_a the plasma vertical elongation. Moreover, in Eq. 1.3 the following definitions were used: $\tau_B = a^2 B_T / T$ is the Bohm timescale, $\rho_* = \rho_i / a$ with ρ_i being the ion gyro-radius, β is the plasma beta (ratio of thermal pressure to magnetic pressure), ν_* is the normalized collisionality, and q is the cylindrical “safety factor” (number of toroidal turns per poloidal turns of magnetic field lines).

It may be argued that scaling laws are not “physical” in nature, meaning that it is always possible to fit a function of the form of Eq. 1.2 using engineering parameters, but it is not necessarily meaningful or illuminating. It is however remarkable that these scalings are *robust* across devices, indicating the existence of macroscopic patterns that are not highly sensitive to experimental uncertainties, details of the plasma geometry, or ways of collecting and analyzing data. One may therefore claim that the purpose of transport studies in fusion devices is to make sense of these scalings on a physical basis, leading to a broad understanding of how one could even *use* transport to improve performance, rather than be subjected to its nefarious effects.

In attempting to understand these phenomena, it has long been clear to the fusion community that comparing and combining data from multiple tokamaks is paramount, for example to reduce covariances between physical variables when constructing scaling laws of the form of Eqs. 1.2 and 1.3. Generally, there is great value in ensuring that conclusions reached with one device can be cross-validated on another. It is also important to create predictive tools that may be used to study future devices. This is the philosophy adopted in this thesis, which examines experimental data and simulations from the Alcator C-Mod and DIII-D tokamaks, and offers some projections to the SPARC and ITER devices, which are undergoing construction in Massachusetts and in France, respectively, at the time of writing. Some example layouts of the magnetic geometry and vacuum vessels for these four experiments are shown in Fig. 1-2. Note that figures are not to scale: ITER is much larger than C-Mod, while DIII-D and SPARC are of similar size. SPARC and ITER are both expected to reach a condition of *breakeven*, where output power will be greater than input power (quality factor $Q > 1$) for the first time. The different dimensions of the devices are explained by their different magnet technologies: high-temperature superconductors (HTS) in SPARC as opposed to low-temperature superconductors (LTS) in ITER. It is beyond the scope of this thesis to describe the approaches of these devices; readers are referred to Refs. [12, 13] for SPARC, and to Refs. [10, 9] for ITER. In the next sections, we will give an overview of Alcator C-Mod and DIII-D, where data in this thesis were collected.

1.2.1 The Alcator C-Mod Tokamak

Alcator C-Mod, a high-field and compact tokamak, operated on the MIT campus between 1993 and 2016. The third of a series of tokamaks at MIT, C-Mod had a major radius of

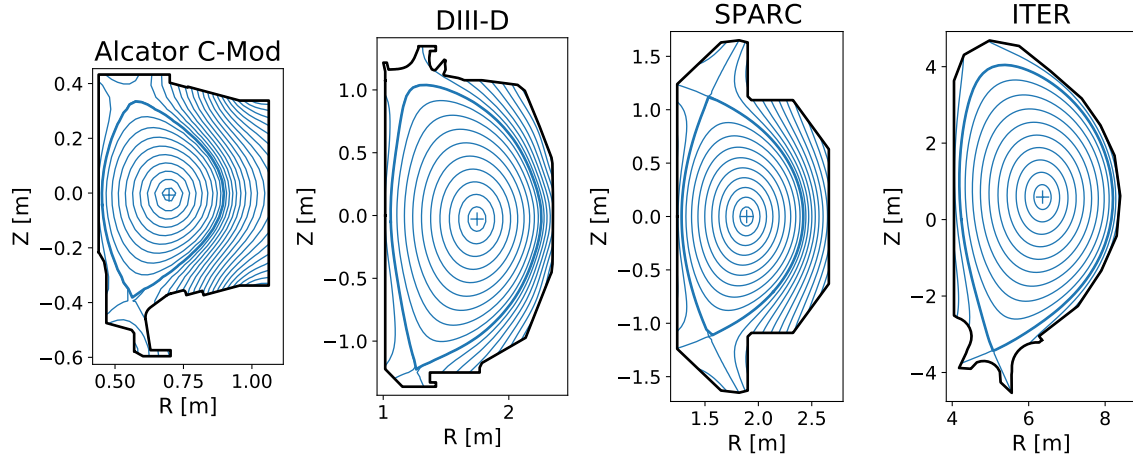


Figure 1-2: Layouts of illustrative magnetic flux surfaces (blue) for Alcator C-Mod, DIII-D, SPARC, and ITER with respective vacuum vessels (black lines). Machine sizes are not shown to scale. The C-Mod and DIII-D equilibria are from discharges #1101014019 ($t = 1$ s) and #174738 ($t = 3$ s), respectively. The SPARC equilibrium is for the V2 device design. The ITER one refers to the ITER baseline scenario.

$R = 0.67$ m, a minor radius of $a \approx 0.22$ m, and toroidal magnetic fields that could reach 8 T on axis, although it typically operated at $B_T \approx 5.4$ T, the expected field for ITER. Plasma currents could go up to $I_p = 2.0$ MA, achieving the highest core densities of any tokamak, up to $n_e \approx 5 \times 10^{20} \text{ m}^{-3}$. Auxiliary power systems used to reach highest performance included ion cyclotron heating (ICH, 50-80 MHz, often used in a H-minority heating scheme) with $P_{ICH} = 6$ MW, and lower hybrid heating (LHH, 4.6 GHz) with $P_{LHH} = 1$ MW, also used for current drive. Notwithstanding its small size, C-Mod was able to push the boundaries of what has been achieved in tokamaks to date, attaining the world record of plasma pressure (approximately 2 atmospheres) on its last day of operation - beating its own preceding record. C-Mod's walls, mostly made of molybdenum, made it an extremely relevant device to test reactor concepts, all while avoiding the complications that would have derived from more intense production of neutrons in a larger plasma. In other words, C-Mod offered an excellent scientific platform with remarkable flexibility. Fig. 1-3 shows an engineering rendering of the device, also illustrating the external cage required to sustain strong forces on the copper magnets and the lower pylons that anchored the device to the ground.

This thesis leverages data collected from Alcator C-Mod experiments over the years; new experiments were conducted and supported at the DIII-D tokamak, which is described next.




Figure 1-3: Engineering rendering of the Alcator C-Mod tokamak, showing an idealized plasma (pink/orange) at its center.

1.2.2 The DIII-D Tokamak

DIII-D is a diverted, medium-size, and moderate-field device with major radius $R = 1.67$ m and typical minor radius of $a \approx 0.67$ m, in operation since 1986. DIII-D is located at General Atomics in San Diego, CA, where a series of doublet-shaped toroidally-confined plasmas were developed for many years, until DIII-D joined the main line of worldwide tokamak progress with diverted “D”-shaped plasmas. The device, shown in Fig. 1-4, can reach toroidal magnetic fields of $B_T = 2.2$ T, but typically operates with $B_T = 2.0$ T. Plasma current can go up to $I_p = 3.0$ MA. Unlike on C-Mod, auxiliary power on DIII-D is provided by electron cyclotron heating (ECH, up to $P_{ECH} = 6$ MW) and by neutral beam



Figure 1-4: Interior (top) and exterior (bottom) of the DIII-D tokamak [14].

injection (NBI). The latter system is powerful and versatile, reaching up to $P_{NBI} = 20$ MW (although lower power is normally available) and maintaining the ability to steer beams both toroidally and poloidally. This allows advanced control of rotation and rotation shear, which has been leveraged to produce turbulence-suppressed advanced tokamak scenarios. At typical DIII-D densities ($n_e < 10^{20} \text{ m}^{-3}$ on axis), the NBI system also constitutes an effective means for core fueling, on top of edge gas puffing. Neutral beams also offer numerous opportunities for diagnostic techniques, as will be described in later chapters.

C-Mod and DIII-D experiments typically have comparable temperature profiles, but C-Mod operated at an order of magnitude higher density. Nonetheless, DIII-D typically produces much higher neutron rates due to its larger volume ($V_p \approx 20 \text{ m}^3$, as opposed to $V_p \approx 1 \text{ m}^3$ for C-Mod). This makes DIII-D an interesting test bed for experiments aiming at higher absolute performance, albeit also making operation more complex. In the next section, we begin describing the subject matter of this thesis and how it will be addressed.

1.2.3 Plasma Transport in Tokamaks

Transport studies have historically focused on heat transport, partly because of its prominent role in establishing high peak temperatures and long energy confinement time, fundamental performance parameters as shown by Eq. 1.1. Achieving high peak density, however, is equally important: it may be easily shown that fusion power scales with central fuel density squared [2]. The study of particle transport is therefore paramount, but complicated by difficult diagnosis of particle sources, neutral particle behavior, and fueling effectiveness – to name a few broad issues. Atomic physics processes are also central to ionization balance and the corresponding radiation in a plasma. As a result, the estimation of particle transport properties and the quantification of uncertainties remain a grand challenge for fusion research.

Particle transport sets not only fuel density, but also impurity concentrations. High levels of contamination pose a significant threat to the prospects of making tokamaks a viable option for energy production, because they both dilute the fuel and also efficiently radiate power away from the core (with a Z^2 approximate scaling, Z being the atomic number) [15]. The need to eliminate core impurities is somewhat in conflict with the apparent requirement of injecting impurities at the edge of the plasma in order to protect material components by radiating energy away before any high-energy plasma strikes the device surfaces [16, 17]. In other words, the challenge is to prevent impurities from leaving the edge regions, where they have a protective function, while also ensuring rapid expulsion of those reaching the plasma center. Alas, helium ash is directly produced by fusion reactions in the core, complicating the matter further. The material composition of device internal surfaces also plays a fundamental role: low- Z materials such as graphite (the main component of DIII-D walls) support high-performance operation in current devices [18], but they are unfortunately not viable in a reactor due to high tritium retention and low threshold energies for physical sputtering by hydrogen isotopes. This results in high erosion yields, which cannot be sustained during the long-term stationary operation envisioned for energy production [19]. On the other hand, high- Z walls such as the molybdenum ones on Alcator C-Mod have favorable hydrogen retention and sputtering yields, but they produce stronger radiation per contaminating ion (due to the aforementioned Z^2 scaling), making the importance of impurity transport all the more evident.




Figure 1-5: Concept of laser blow-off injections into a tokamak plasma, also showing the radial profile of toroidal magnetic field (top panel) and describing some of the nomenclature for tokamak geometries.

In this thesis, we attempt to quantify particle transport through rapid and non-perturbative injections of chosen ions via *Laser Blow-Off* (LBO) [20]. The general concept of LBO injections is shown in Fig. 1-5: a high-power, class-IV laser is used to produce a rapid pulse of radiation directed at a glass slide, which supports (under vacuum conditions) a thin film coating a few μm thick. When the laser is fired, it ablates the film, producing a cloud of neutral particles that proceeds towards the confined plasma (note that the magnetic field will divert any atoms that are unintentionally ionized by the laser). As described in Chapter 4, a number of spectroscopic instruments can then be used to diagnose how neutrals ionize and then transport through the plasma over space and time. This is the basis of how we experimentally attempt to infer transport coefficients in this thesis. The top panel in Fig. 1-5 shows the toroidal magnetic field profile as a function of major radius, scaling approximately as $\sim 1/R$. Red lines in the main figure display magnetic flux contours. Only contours inside the thick red line close upon themselves within the device, thus motivating the nomenclature of *last closed flux surface* (LCFS) for the thick line. The *magnetic axis* is a single point at

the location of maximum flux, indicated with a “+” sign. Outside of the LCFS, open field lines are said to be in the *scrape-off layer* (SOL), appropriately named since plasma along these lines is not confined for long before it hits material surfaces. SOL plasma is expected to have relatively low density and much colder temperatures than plasma in the core region, and is also affected by important interactions with neutral (non-ionized) atoms, as we will discuss in Chapter 6. The plasma in Fig. 1-5 is said to be *diverted* because the LCFS is determined by a magnetic field null (*x-point*, indicated by a black arrow) as opposed to a point of direct contact of closed field lines with a *limiter* surface, which could either be on the inner or outer wall. In this magnetic reconstruction there exists a single x-point, and the region below it (the *divertor*) has the function of eliminating any impurities that build up in the plasma. In practice, in current devices the divertor also plays an important role of energy dissipation before any hot plasma leaving through the LCFS can reach material surfaces.

A central theme that permeates the research in this thesis is *core-edge integration*, i.e. the creation of tokamak scenarios where the core is able to reach high performance (broadly defined by the triple product in Eq. 1.1) while enabling survival of the edge surfaces at full operation. This involves the creation of self-organized plasma states where heat, particles and momentum transport can be effectively tailored. It may seem a contradiction to say that a “self-organized state” can be “tailored” – after all, if there is external control one cannot say that there is self-organization in the common scientific sense. In fact, fusion research has long hinged on the discovery of advanced regimes with varying degrees of control over plasma turbulence, starting with the discovery of the “high-confinement” (H-) mode on the ASDEX tokamak in 1982 [21]. A wide range of advanced regimes has followed, including the Enhanced D-Alpha (EDA) H-mode [22] and the “improved” (I-) mode [23] that will be examined in Chapter 7. Among the wealth of high-performance regimes undergoing exploration at DIII-D [24, 25], this thesis examines one of the most recent and exotic ones, produced in *negative triangularity* plasma shapes, i.e. when the “D” shape of the plasma poloidal cross section is inverted horizontally. The focus of the research presented here is on the experimental inference and modeling of particle transport, rather than the overreaching challenge of core-edge integration, although the broader context of this research is what guided the author’s interest in examining promising regimes for future reactors.

Throughout this thesis, we present advances on multiple aspects of experimental particle

transport. We focus on the quantification of metrics that may be compared between theory and experiments for the purpose of validating state-of-the-art models that support the development of future magnetic confinement fusion devices. Modeling capabilities have been advanced with **Aurora**, a modern package for particle transport and radiation modeling, briefly presented in Chapter 5.3.1 and described in greater depth in Appendix A. **Aurora** represents the means, rather than the objective, of this thesis. Its development was a necessity for this research, and at the time of writing it is already being applied in a number of other projects by researchers worldwide. While the research presented in this work does not include the development of new instrumentation, tools to advance experimental data analysis and interpretation have been developed. Significant effort has been dedicated to demonstrating the importance of charge exchange of neutral particles with ions in the outer part of the plasma. While this phenomenon has long been studied in the region of open field lines of tokamaks, this thesis shows that it cannot be neglected in the study of particle transport and radiation in the confined plasma either. The Alcator C-Mod and DIII-D research presented here offers the most detailed and rigorous methods for the estimation of particle transport coefficients in tokamaks to date. Combined with increasingly sophisticated data analysis techniques and the inclusion of a wealth of measurements, this thesis contributes to validating theoretical transport models for a variety of operational regimes on both devices, examining agreements and disagreements with experiment.

1.3 Summary of Primary Results

We summarize here the most significant contributions made in the research presented in this thesis:

1. **Development of the Aurora package for particle transport and radiation modeling.** Aurora offers important advances for impurity transport simulations with greater fidelity and ease of parallelization with respect to previously available tools. Its open-source routines enable detailed spectroscopic analysis and integration with other specialized modeling suites.
2. **Quantification of the impact of charge exchange between background deuterium neutrals and impurity ions.** This effect has been analyzed via both measurements in the Alcator C-Mod edge and state-of-the-art modeling tools, for both

C-Mod and DIII-D. The importance of charge exchange in pedestal particle transport studies has motivated its consideration in inferences of particle transport coefficients for the first time.

3. **Creation of fully-Bayesian high-performance computing frameworks for the inference of experimental impurity transport.** These have enabled comparisons of inferred transport coefficients to theoretical neoclassical and turbulent models for both Alcator C-Mod and DIII-D discharges. A variety of operational scenarios have been examined, including L-, H- and I-mode discharges on C-Mod, and negative triangularity shots on DIII-D. We have observed relatively good agreement in diffusion in most cases, although turbulence models appear to often underestimate experimentally-inferred diffusion. Apparent discrepancies have been found in turbulent convection, particularly in cases where experimental impurity profiles are flat or hollow.
4. **Forward modeling of the entire K_α spectrum of calcium and extreme ultraviolet line ratios to constrain experimental impurity transport on C-Mod.** The calcium K_α spectrum has been measured at Alcator C-Mod for years, using X-ray Imaging Crystal Spectroscopy, but in this work we constrained impurity transport using all the available spectral features for the first time.
5. **Demonstration of inferences of experimental impurity transport using multi-species constraints on DIII-D.** The capability to infer particle transport coefficients using both quasi-steady and time-evolving impurity measurements provides much greater confidence in radial profiles of experimental diffusion and convection.

1.4 Organization of This Thesis

The first chapters of this thesis aim to provide background on plasma transport (Chapter 2) and spectroscopy (Chapter 3), to the extent needed to appreciate the research presented in later chapters. These background sections are not intended to be complete reviews of their subjects, but references are provided for readers interested in exploring these topics further. In Chapter 4 we describe the experimental diagnostics that were most important to this work. In Chapter 5, we begin to journey into the field of Bayesian statistics as applied in this thesis, providing a short introduction to the subject before describing algorithmic

choices used for inferences of particle transport. In this chapter, we also describe a number of techniques developed to address modeling issues identified in this and previous works on transport analysis. This naturally leads to the description of the new Aurora package. In Chapter 6, we present research on neutral particles through a number of experimental models and experimental measurements. These results are important to the following two chapters on experimental inferences of particle transport, the first (Chapter 7) on C-Mod and the second (Chapter 8) on DIII-D. Finally, in Chapter 9 we provide a summary of results and discuss future avenues of research.

This work builds on numerous others, particularly the Ph.D. theses of N.T. Howard [26] and M.A. Chilenski [27], as well as close collaboration with T. Odstrčil. Reference will be made throughout this document to works in the literature that preceded or indirectly contributed to the advances presented in this thesis.

2

Particle Transport in Tokamaks

In this chapter we introduce important concepts of plasma transport. The neoclassical, gyrofluid and gyrokinetic models described here are the basis of models that will be compared to experimental inferences of particle transport in later chapters. We highlight assumptions and limitations of each model, also characterizing magneto-hydrodynamic phenomena and edge transport as relevant to this thesis. Finally, we discuss parametrizations of transport that permit effective experimental validation of theoretical models.

Introduction

A plasma is in thermodynamic equilibrium only when it is spatially homogeneous and its particle distributions are isotropic and Maxwellian. These conditions are not realizable in a magnetically confined plasma, in which spatial gradients inevitably exist. Moreover, since the presence of magnetic fields leads to faster motion along field lines than across, a strong velocity-space anisotropy exists and constitutes a source of free energy for instabilities leading to non-Maxwellian distributions. While collisions are the reason why Maxwellians are attained as maximum-entropy states in ideal gases (Boltzmann H theorem) [28], they are increasingly rare in hot plasmas, where collisional frequencies decrease with temperature. In this chapter, we will briefly describe the collective micro- and macroscopic phenomena that are particularly relevant to particle transport in tokamaks.

Plasma transport underpins many aspects of tokamak phenomenology; we shall provide two examples. A *radiative collapse* of a tokamak discharge results from impurities being transported towards the plasma center to the point that they radiate more power than

is being supplied by external means and fusion reactions. Such events are of particular concern because neoclassical theory, discussed in Section 2.3, predicts a scaling for the peaking of density profiles that goes with Z (atomic number), meaning that any ion with charge greater than hydrogen may be expected to accumulate near the magnetic axis. MARFEs¹ are another example of problematic interplays between transport and atomic physics: during these events, an edge thermal instability is caused by atomic radiation efficiencies for low- Z ions increasing as one goes to lower temperatures. Neoclassical transport in the core and MARFEs in the edge exemplify two radically different, crucial and complex phenomena related to both plasma transport and atomic physics. The focus of this thesis is to separate these effects and improve understanding of the former. In Section 2.2, we begin by briefly illustrating classical collisional transport theory, before moving to the more tokamak-relevant neoclassical theory of collisional transport in toroidal geometries. Turbulent transport is addressed in Section 2.4, where we describe models of varying complexity to simulate tokamak transport, particularly gyrokinetic and gyro-fluid theories. Quasi-linear descriptions of turbulence are outlined in Section 2.4.3. Section 2.4.4 offers a high-level characterization of turbulent modes in tokamaks. Finally, in Section 2.7 we discuss how the validation of particle transport models has been attempted in the past, particularly via the inference of transport coefficients. This sets the stage for the following chapters of this thesis, where neutrals and impurity transport models are compared to experimental data from the Alcator C-Mod and DIII-D tokamaks.

2.1 Kinetic Theory

The most general statistical description of a plasma is given by the Boltzmann equation,

$$\frac{\partial f_a}{\partial t} + \mathbf{v} \cdot \frac{\partial f_a}{\partial \mathbf{x}} + \mathbf{a} \cdot \frac{\partial f_a}{\partial \mathbf{v}} = \left(\frac{\partial f_a}{\partial t} \right)_c, \quad (2.1)$$

describing the evolution of the distribution function, f_a , in space, velocity and time for a given particle species a . The term on the right hand side of Eq. 2.1 accounts for collisions, sometimes requiring elaborate models; when these are completely omitted, setting $(\partial f_a / \partial t)_c = 0$, the corresponding equation is referred to as the Vlasov (collisionless)

¹The term ‘‘MARFE’’ is said to stand for ‘‘Multifaceted Asymmetric Radiation From the Edge’’. The phenomenon was first detected by Marmor and Wolfe.

equation. In plasmas, one can approximate the collision term on the right hand side of Eq. 2.1 as the sum of many small angle scatterings. Making this approximation, the Boltzmann equation reduces to the Fokker-Planck equation [7].

The acceleration in Eq. 2.1 is dominated by the Lorentz force

$$\mathbf{F}_s = q_s(\mathbf{E} + \mathbf{v}_s \times \mathbf{B}). \quad (2.2)$$

Maxwell's equations specify how electric and magnetic fields are related to each other and to sources (electric charges and currents):

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \quad \nabla \cdot \mathbf{B} = 0 \quad \nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \quad \nabla \times \mathbf{B} = \mu_0 \mathbf{J} \quad (2.3)$$

where ϵ_0 and μ_0 are the vacuum permittivity and permeability, respectively. Note that since the phenomena of interest here are significantly slower than the speed of light, the Maxwell displacement term in Faraday's Law can be taken to be negligible. In Eq. 2.3, ρ and \mathbf{J} are the charge and current density respectively, calculated in terms of moments of the particle distributions via

$$\rho = \sum_a \int d^3\mathbf{v} q_a f_a \quad \mathbf{J} = \sum_a \int d^3\mathbf{v} q_a \mathbf{v}_a f_a. \quad (2.4)$$

Since electromagnetic fields depend on moments of the distribution function and vice-versa, this kinetic description corresponds to a nonlinear advection equation for f , which is the quantity that fully describes our plasma in background fields.

Eq. 2.1 depends on 7 parameters (3 spatial and 3 velocity dimensions, plus time) and is mostly out of our current computational capabilities for problems where a wide range of spatial and velocity scales must be resolved. This motivates the development of reduced (approximate) models that aim to resolve the physics of interest, while ignoring details of lesser relevance.

2.2 Classical Collisional Transport

The classical and neoclassical theories of plasma transport consider the important limit of the kinetic equation (Eq. 2.1) where $\delta = \rho_a/L \ll 1$, ρ_a being the gyroradius (used as a

synonym for Larmor radius) for a species a and L a macroscopic length scales. The classical theory applies in the short mean free path limit, $\lambda/L \ll 1$, in cases where the magnetic field geometry plays no role. In these conditions, by expanding the distribution function as $f_a = f_{a,0} + f_{a,1} + \dots$, one finds that to lowest order the electrons follow a Maxwellian distribution

$$f_{e,0} = n_e \left(\frac{m_e}{2\pi T_e} \right)^{3/2} e^{-\frac{m_e v^2}{2T_e}} \quad (2.5)$$

at rest in the moving frame, with n_e , T_e and m_e being the electron density, temperature and mass, respectively. Higher order terms are normally solved via the Chapman-Enskog procedure [29]. The result, commonly referred as the Braginskii equations [30], provides a set of dependencies and coefficients that have wide applicability in fusion and astrophysical plasmas. However, the classical predictions are nowhere close to observed transport levels in tokamak plasmas, since the aforementioned assumption of short mean free path and the asymptotic assumption in $\delta = \rho_a/L \ll 1$ miss important physics, as we will see in the following sections. Nonetheless, there are important mechanisms in tokamaks where a classical description is appropriate, generally involving transport parallel to the magnetic field. For example, parallel currents are observed to be well described by the Spitzer estimate for electric conductivity

$$J_{\parallel} = -n_e e u_{\parallel} = \sigma E_{\parallel} = 1.96 \frac{n_e e^2 \tau_e E_{\parallel}}{m_e}, \quad (2.6)$$

with the electron collision time being

$$\tau_e = \frac{12\pi^{3/2} m_e^{1/2} T_e^{3/2} \epsilon_0^2}{2^{1/2} n_i Z^2 e^4 \ln \Lambda}, \quad (2.7)$$

$\ln \Lambda$ being the Coulomb logarithm [29]. The parallel conductivity, σ , defined in Eq. 2.6, is important in the tokamak edge, as discussed in Section 2.6. Note that σ scales with $T_e^{3/2}$, but has no n_e dependence, unlike in an ideal gas. Braginskii's $\rho/\lambda \rightarrow 0$ limit also provides an estimate for the force on electrons

$$\mathbf{F}_e = \underbrace{-\frac{m_e n_e}{\tau_e} (0.51 \mathbf{u}_{\parallel} + \mathbf{u}_{\perp})}_{\mathbf{F}_u, \text{ drag/friction}} + \underbrace{-0.71 n_e \nabla_{\parallel} T_e + \frac{3n_e}{2\Omega_e \tau_e} \mathbf{b} \times \nabla T_e}_{\mathbf{F}_T, \text{ thermal force}} \quad (2.8)$$

which in a pure hydrogen plasma requires $\mathbf{F}_i = \mathbf{F}_e$ by momentum conservation. The second term (the *thermal force*) is generally larger than the first term (drag/friction) by $\Omega_e \tau_e \gg 1$, where $\Omega_e = eB/m_e$ is the electron gyrofrequency. Such thermal force is particularly important for impurities in the tokamak edge, providing a mechanism for them to go from the divertor upstream into the SOL [29]. Of course, non-ionized (neutral) particles, for example near the device wall, are also subject to collisions, and these may cause ionization if the projectiles have sufficiently high energy. Modern tokamaks rarely operate at neutral pressures below 1 μ Torr and in these conditions neutral particles always play a non-negligible role on plasma confinement. A classical fluid approximation may be applied to describe neutrals whenever their mean free path is much shorter than the macroscopic scale of their motion. As we discuss in Chapter 6, this condition is often not fully attained in the edge of tokamaks and a fully-kinetic description, e.g. via single-particle (Monte Carlo) simulation, is necessary.

2.3 Neoclassical Transport

The classical theory outlined in the previous section relies on an assumption of *strong magnetization*, $\delta = \rho_a/L \ll 1$, and also on the mean free path being short enough for particles not to be affected by the magnetic field geometry and gradients. In contrast, when the mean free path is long, as typical in tokamaks, a more complex description of transport is needed. Neoclassical theory addresses this situation of collisional transport in toroidal geometry [29].

The assumptions of slow variation in time ($\omega \ll \Omega_{c,i}$) and strong magnetization ($\rho_a \ll L$) are central to the derivation of the *drift-kinetic equation* (DKE) as an approximation of the full kinetic description of Eq. 2.1. The Hazeltine form of the DKE is [31]:

$$\frac{\partial f_a}{\partial t} + (v_{\parallel} \mathbf{b} + \mathbf{v}_D + \mathbf{v}_E) \cdot \nabla f_a - \frac{Z_a e}{m_a} (v_{\parallel} \mathbf{b} + \mathbf{v}_D) \cdot \nabla \phi \frac{\partial f_a}{\partial \epsilon} + \dot{\mu} \frac{\partial f_a}{\partial \mu} = C_a + S_a \quad (2.9)$$

where $\mathbf{b} = \mathbf{B}/B$ is the magnetic field unit vector, μ is the particle magnetic moment, ϕ is the electrostatic potential, C_a is the nonlinear collision operator (summing over binary collisions between all species) and S_a represents sources (particle, energy, etc.). The drift velocities entering this equation capture the presence of currents and pressure gradients in

toroidal geometry [32]:

$$\mathbf{v}_D = m_a \frac{v_{\parallel}^2 \mathbf{b} \times (\mathbf{b} \cdot \nabla) \mathbf{b} + \mu \mathbf{b} \times \nabla B}{Z_a e B} \quad \mathbf{v}_E = \frac{\mathbf{b} \times \nabla \phi}{B}. \quad (2.10)$$

The solution of the DKE must be supplemented by the Poisson Equation, which can be written as [32]

$$-\sum_a n_{a,0} \frac{e^2 Z_a^2}{T_{0,s}} \rho_a^2 |\nabla r|^2 \frac{\partial^2 \phi}{\partial r^2} = \sum_a Z_a e r \int d^3 v f_a, \quad (2.11)$$

where we sum over species with index a , and we defined the gyroradius $\rho_a = v_{t,a}/\Omega_a$ and thermal speed $v_{t,a} = \sqrt{T_{a,0}/m_a}$. As above, the index “0” indicates terms at equilibrium scales. The $|\nabla r|$ is a metric element to correctly account for physical radial distances. Starting from these expressions, neoclassical theory can be derived by considering specific limits, typically by expanding the DKE (Eq. 2.9) in terms of the small δ parameter (or, equivalently, $\rho_* = \rho_i/a$, defined with the tokamak minor radius). Cross-field fluxes can then be computed by evaluating terms at each order of the expansion. For example, the second-order radial particle flux, averaged over a flux surface, is given by [32]

$$\Gamma_{a,2} = \left\langle \int d^3 v \left(f_{a,0} \mathbf{v}_E^{(1)} \cdot \nabla r + f_{a,1} \mathbf{v}_D \cdot \nabla r \right) \right\rangle, \quad (2.12)$$

where $\mathbf{v}_E^{(1)}$ indicates the 2nd order $E \times B$ drift (computed from the 1st order electrostatic potential) and \mathbf{v}_D is the 1st order drift of Eq. 2.10.

Neoclassical transport is typically divided into three regimes: the banana, plateau and Pfirsch-Schlüter ones. These are differentiated by the ion collisionality, defined as

$$\nu_i^* = q R_0 \nu_i / v_{ti} = q R_0 / l_f \quad (2.13)$$

where q is the plasma safety factor, R_0 is the major radius, ν_i is the ion collision frequency, and $l_f = v_t/\nu_i$ is the mean free path for 90 degree scattering of the momentum. The quantity ν_i^* can be interpreted as a measure of the number of collisions that an ion experiences when traveling a significant fraction of a poloidal turn along a field line. Defining the inverse aspect ratio as $\epsilon = a/R_0$, with a being the minor radius, neoclassical transport is distinguished into the banana (collisionless) regime for $\nu_i^* < \epsilon^{3/2}$, the plateau regime for $\epsilon^{3/2} < \nu_i^* < 1$, and the Pfirsch-Schlüter (collisional) regime, for $\nu_i^* > 1$. In the collisionless regime, trapped

particles perform full banana orbits before being untrapped by collisions, hence the name “banana regime”. In the collisional regime, particles only perform a fraction of a full poloidal orbit between collisions, preventing completion of banana orbits. The plateau regime is an intermediate regime.

Among the most important predictions of neoclassical theory is the one of the *bootstrap* (parallel) current, resulting from collisional effects of both passing and trapped electrons in the presence of radial pressure gradients. This current is “self-organized” in the plasma in the sense that it is not externally induced and does not require the use of the external tokamak transformer, thus offering opportunities for long-term operation with no Ohmic current drive. As mentioned in the introduction to this chapter, cross-field neoclassical transport is also of great relevance near the plasma center, where it generally predicts a $\sim Z^2$ trend for impurity peaking that endangers high performance operation. Neoclassical transport is relevant near axis because turbulent instabilities are often stable in this region. Neoclassical transport trends have also been found experimentally in the pedestal region for impurities [33], although this observation awaits further validation.

A number of analytical approximations to neoclassical transport exist. For example, the Chang-Hinton formulae are generally applicable at experimentally-relevant collisionality in *circular plasmas*. Their applicability is obviously reduced in the outer parts of a diverted tokamak plasma, but they are nonetheless useful near the magnetic axis. Many numerical tools have been developed over the years to obtain more elaborate predictions; among these, the NCLASS model employs a combination of numerical and analytical methods, again under the limitation of circular plasmas. Recently, the NEO code [32, 34] has been offering more accurate numerical DKE solutions to first order in ρ_{*i} for arbitrary plasma shapes, and it is therefore the neoclassical model used in this thesis.

We note that although neoclassical theory is valid independently of the collisional mean free path, it relies on the assumption of Larmor orbits being small with respect to the macroscopic length scales ($\delta = \rho_a/L \ll 1$); in some cases, this can make it strictly inapplicable in the steep gradient pedestals region. Moreover, neoclassical transport does not include effects of *finite orbit loss*, i.e. Larmor orbits crossing the LCFS and suffering a collision in the SOL before crossing again. Neoclassical calculations used in this thesis effectively solve a steady-state problem at the scale of the plasma equilibrium, i.e. they do not treat fluctuations on top of this macroscopic equilibrium. While the development

of computational models that may treat both equilibrium and turbulent time scales are obviously desirable, computational demands for such “full- f ” models are still considered formidable. This makes interactions between large- and small-scale phenomena difficult to examine, e.g. in the case of neoclassical radial electric fields providing the $E \times B$ shearing believed to lead to turbulent stabilization at the L-H transition. Feedback mechanisms between neoclassical and turbulent scales are mostly taken to be negligible in this work, considering fluxes arising from neoclassical and turbulent transport to be additive. In the next section, we shall examine turbulent phenomena and describe how this thesis aims at validating some of their state-of-the-art descriptions.

2.4 Turbulent Transport

Microturbulence has long been observed to dominate over neoclassical transport in tokamaks [7]. While collisional toroidal transport is often taken to represent a minimal and irreducible level of cross-field dynamics, turbulence is highly susceptible to a wide range of parameters and experimental actuators. Cross-field transport arises when fluctuating quantities are out of phase; for example, writing the turbulent particle flux as $\Gamma \equiv \langle \tilde{n}\tilde{v} \rangle$, with angle brackets denoting the real component of an ensemble average, we see that $\Gamma \neq 0$ only if the product of density fluctuations, \tilde{n} , and velocity fluctuations, \tilde{v} , gives a real component. This is only the case if there exists a difference in phases, or else the fluctuations may be considered part of a stable wave. Tokamak turbulence theory aims at predicting fluctuation amplitudes and phases to estimate cross-field transport.

In this section, we first discuss the gyrokinetic formalism, which has emerged as the “standard model” of turbulent transport in fusion devices. We will then describe a characterization of important microturbulence modes, before concluding with a discussion of quasi-linear reduced models of turbulent transport, particularly focusing on gyro-Landau fluid theory.

2.4.1 The Gyrokinetic Equation

All formulations of gyrokinetics are based on the assumption that the fast time scales associated with Larmor motion can be integrated over without a loss of fundamental dynamics. However, unlike drift-kinetics, gyrokinetic theory allows electromagnetic fields

to vary significantly across a Larmor radius [7]. The gyrokinetic equation is derived by averaging the Fokker-Planck equation over gyro-orbits [35], eliminating the gyro-phase from the set of coordinates. This reduces the problem dimensionality from 7 dimensions (3 in space, 3 in velocity, 1 in time) to 6 dimensions (3 in space, 2 in velocity, 1 in time). Further simplifications are sometimes made when modeling plasmas with negligible magnetic field curvature; for example, one may consider a “straight” (slab) rather than a “curved” magnetic geometry to ignore some terms in the equations. The condition $k_{\parallel} \rightarrow 0$, often referred to as the “toroidal limit”, is also sometimes invoked to study effects that are intrinsic to magnetic field gradients. One may summarize the key assumptions of gyrokinetics as [36]

- *Strong magnetization*: the Larmor radii of modeled species are taken to be much smaller than the system size, i.e. $\rho_a \ll L$.
- *Low frequencies*: fluctuations are taken to have frequencies much lower than the ion gyrofrequency, i.e. $\omega \ll \Omega_{c,i}$.
- *Anisotropy*: fluctuations perpendicular to the magnetic field must be much smaller than fluctuations parallel to it, i.e. $k_{\perp} \gg k_{\parallel}$. The gyrokinetic equation allows for $k_{\perp}\rho_a \sim 1$, whereas the DKE assumes $k_{\perp}\rho_a \ll 1$.

Moreover, gyrokinetic models that rely on an asymptotic expansion of the particle distribution function, typically referred to as “ δf ”, also assume that fluctuations are small compared to the equilibrium plasma properties. All gyrokinetic approaches also rely on an implicit *weak coupling* assumption, i.e. the plasma density must be sufficiently dense that many particles exist within a Debye sphere, as required for the standard definition of a plasma. An overview of the gyrokinetic formalism is given, for example, in Ref. [36]. In δf treatments, such as those considered in this thesis work, the full particle distribution function is expanded as

$$f_s = f_{0,s}(v, t) \exp\left[-\frac{q_s\phi(\mathbf{r}, t)}{T_{0,s}}\right] + h_s(\mathbf{R}_s, v, v_{\perp}, t) + \delta f_{2s} + \dots \quad (2.14)$$

with $v = (v_{\perp}^2 + v_{\parallel}^2)^{1/2}$ and $f_0 = n_{0s} \exp(-v^2/v_{t,s}^2)/(\pi^{3/2}v_{t,s}^3)$. Here, the 0th order term of f_s corresponds to the Boltzmann factor. The 1st order correction, h_s , is taken to represent all relevant fluctuation-scale dynamics as a function of the guiding centers (\mathbf{R}_s), rather than particle positions along orbits. The operations of gyroaveraging of the full Fokker-Planck

equation result in the nonlinear gyrokinetic equation

$$\frac{\partial h_a}{\partial t} + \underbrace{v_{\parallel} \frac{\partial h_a}{\partial z}}_{\text{parallel streaming}} + \overbrace{\frac{1}{B_0} \{ \langle \chi \rangle_{\mathbf{R}_a}, h_a \}}^{\text{drifts}} = \underbrace{\frac{q_a f_{0s}}{T_{0s}} \frac{\partial \langle \chi \rangle_{\mathbf{R}_a}}{\partial t}}_{\text{adiabatic response}} + \overbrace{\left(\frac{\partial h_a}{\partial t} \right)_c}^{\text{collisions}}, \quad (2.15)$$

where $\chi = \phi - \mathbf{v} \cdot \mathbf{A}$ is the gyrokinetic potential and $\{ \dots \}$ represents a Poisson bracket.² As for the DKE described in Section 2.3, the electromagnetic potentials are related to the fluctuating component of the distribution function via the Poisson equation, representing quasi-neutrality

$$\sum_s q_s \delta n_s = \sum_s \left(-\frac{q_s^2 n_{0,s} \epsilon_0}{T_{0,s}} \phi + q_s \int d^3 \mathbf{v} \langle h_s \rangle_{\mathbf{r}} \right) = 0, \quad (2.16)$$

where the density n_0 has been defined as the first moment of the 0th order particle distribution function. The parallel and perpendicular components of Ampère's law are also needed to evolve magnetic field components, and can be written, respectively, as

$$\begin{aligned} -\nabla_{\perp}^2 A_{\parallel} &= \sum_s \mu_0 q_s \int d^3 \mathbf{v} v_{\parallel} \langle h_s \rangle_{\mathbf{r}} \\ \nabla_{\perp} \delta B_{\parallel} &= \sum_{\mathbf{c}} \mu_0 q_s \int d^3 \mathbf{v} \langle (\hat{\mathbf{z}} \times \mathbf{v}_{\perp}) h_s \rangle_{\mathbf{r}}. \end{aligned} \quad (2.17)$$

with μ_0 being the permeability of free space. The gyrokinetic equation is typically solved to study low-frequency (with respect to Ω_c) turbulent modes. In some cases, slowly evolving phenomena may also be studied with a further reduction in dimensionality by further averaging over the parallel motion of particles, giving 4-dimensional bounce-averaged equations [7].

Numerous gyrokinetic codes have been developed to model plasma microturbulence, both in a local (flux tube, $\nabla_{\parallel} \rightarrow ik_{\parallel}$, $1/\rho_* \rightarrow \infty$) and global (spatially inhomogenous) setting. In this thesis, we employ the CGYRO code [37, 38, 39] as a state-of-the-art numerical tool to predict particle transport both linearly and nonlinearly. CGYRO is a high-performance computing (HPC) code that is particularly optimized for accurate treatment of multi-species plasmas that may have important collisional effects, as for example near the pedestal region of tokamaks. A spectral/pseudospectral representation is used in 4 of 5 phase space dimensions to enable a convenient treatment of advanced collision operators

² $\{U, V\} := \hat{\mathbf{z}} \cdot [(\partial U / \partial \mathbf{R}_s) \times (\partial V / \partial \mathbf{R}_s)]$

and an extended sonic rotation limit [37, 38].

2.4.2 Gyro-Fluid Theory

The gyrokinetic equation offers a valuable reduction in dimensionality with respect to the full kinetic equation, but is still remarkably complex and computationally expensive to solve. This has led to the development of reduced models that can be executed faster, while still retaining at least some of the key physics of interest, to complement more detailed (but limited in number) simulations within the gyrokinetic formulation.

Gyro-Landau Fluid (GLF) equations are velocity moments of the gyrokinetic equation that are closed in such a way as to retain kinetic effects such as linear Landau damping [40, 41, 42]. The latter, in particular, cannot be retained in a set of purely-fluid equations, regardless of how many moments of the distribution function one is willing to compute. This may be understood by considering the kinetic dispersion function $\mathcal{Z}(\zeta)$, where $\zeta = \omega/(kv_t)$ is the ratio of particle thermal velocity and wave phase velocity:

$$\mathcal{Z}(\zeta) = \frac{1}{\pi} \int_{-\infty}^{+\infty} \frac{dv_{\parallel} e^{-v_{\parallel}^2/v_t^2}}{v_{\parallel} - \zeta}. \quad (2.18)$$

Let us Taylor expand this in the fluid ($|\zeta| \gg 1$) and kinetic limits ($|\zeta| \ll 1$) [43]:

$$\mathcal{Z} \approx \begin{cases} i\pi^{1/2} e^{-\zeta^2} - 2\zeta \left[1 - \frac{2}{3}\zeta^2 + \frac{4}{15}\zeta^4 + \dots \right] & \text{for } |\zeta| \ll 1 \text{ (kinetic limit)} \\ i\pi^{1/2} e^{-\zeta^2} - \frac{1}{\zeta} \left[1 + \frac{1}{2\zeta^2} + \frac{3}{4\zeta^4} + \dots \right] & \text{for } |\zeta| \gg 1 \text{ (fluid limit)}. \end{cases} \quad (2.19)$$

The phenomenon of collisionless phase-mixing at the basis of Landau damping, in-existent in standard fluids, arises from the imaginary part of $\mathcal{Z}(\zeta)$. In both the above limits, the imaginary component is seen to scale with $\exp(-\zeta^2)$. Now, recall that an exponential function grows/decays faster than any polynomial of the same arguments. In the fluid limit, the argument of the exponential function ($-\zeta^2$) is large and negative, therefore going to 0 faster than the real polynomial series in the same expansion in Eq. 2.19. On the other hand, in the kinetic limit we have $|\zeta| \ll 1$ and the exponential term may still give important physics [40, 41].

The first gyro-Landau fluid models used to expand the plasma dispersion function as a Padé approximant rather than as a Taylor series in $|\zeta_m|$ [41], but more recent models make

use of the full FLR integrals and expand the fields in a finite series of Hermite polynomials in the poloidal direction [42]. A fluid closure is applied, typically at the 4th or 5th moment, to retain a k_{\parallel} -dependent damping term which mimics the collisionless phase-mixing process underlying Landau damping, fitting a number of coefficients with the results of an exact kinetic response [40]. Gyro-fluid models may in principle be valid all the way from lowest- k trapped ion modes to the ETG modes at high- k [42].

The Trapped Gyro-Landau Fluid [42] (TGLF) model separates the influence of passing and trapped particles; as we will see in Section 2.4.4, this is paramount to resolve some of the critical physics of tokamak microturbulence. The TGLF code, described in the next section, is named after this system of gyro-fluid equations that it solves. TGLF has become a leading tool for comparison of theory and experimental observations, relying on the quasilinear approximation to capture critical physics at limited computational cost.

2.4.3 The Quasilinear Approximation

Fully-nonlinear simulations offer the highest degree of fidelity, but at significant computational cost. Highly-parallelized codes like CGYRO [37] can only be used to compare theory and experiment in a limited number of cases; validation over large databases demands faster computational models. A widely adopted approach for this purpose is to make use of quasilinear approximations.

In quasilinear models, only the linear turbulence response is explicitly computed; nonlinear interactions are approximated via relatively simple numerical rules, which can in principle be based on experimental observations or fundamental principles. Quasilinear codes typically solve an eigenvalue problem for the linear component, for which numerical solutions can be obtained rapidly. Unlike initial value solvers like CGYRO, this allows one to find all the unstable branches of the equations, including subdominant ones that ensure a smooth response near transitions between different turbulent regimes. Each eigenvalue corresponds to a turbulent *mode* whose frequency and wavenumber are related by a *dispersion relation*. The complete mode frequency may be written as $\omega(\vec{k}) + i\gamma(\vec{k})$, where the real part represents an oscillation (wave) and the imaginary part gives the growth (if $\gamma > 0$) or decay (if $\gamma < 0$) rate. To illustrate the quasilinear approximation, let us consider the calculation of turbulent heat fluxes of a species a , $Q_a = \langle \frac{3}{2} \tilde{p}_a \tilde{v}_r \rangle$, and assume that radial velocity fluctuations only arise through $E \times B$ drifts for simplicity. Working

in Fourier space one may write $\tilde{\phi} \sim e^{-i(\omega t - \vec{k} \cdot \vec{r})}$, which corresponds to a decomposition of fluctuations into individual interacting modes; given that $B_\phi \gg B_\theta$ in tokamaks, one may simplify this as $\tilde{v}_r = \nabla \tilde{\phi} \times \vec{B} / B^2 \approx ik_\theta \tilde{\phi} / B$. In tokamak geometry, k_θ should be interpreted as being orthogonal to helically-twisted field lines, i.e. $k_\theta = nq/r$, where q is the safety factor and n is the toroidal wavenumber. We consider an electrostatic approximation and define $\tilde{\Phi} \equiv e\tilde{\phi}/T_e$, $k_y = k_\theta \rho_s$, and $\rho_s c_s = T_e/eB$. This allows us to write the heat flux at each wavenumber as

$$Q_{a,k} = \frac{3}{2} \mathcal{R} \left\langle \tilde{p}_a \frac{ik_\theta \tilde{\phi}}{B} \right\rangle = c_s \frac{3}{2} \mathcal{R} \langle ik_y \tilde{p}_a \tilde{\Phi} \rangle \quad (2.20)$$

where \mathcal{R} indicates the real component. The quasilinear approximation to the heat flux can then be written as

$$Q_a = \sum_k Q_{a,k} = \sum_k Q_{a,k}^{ql} \tilde{\Phi}^2, \quad (2.21)$$

where we defined the quasilinear weights

$$Q_{a,k}^{ql} = \frac{1}{|\tilde{\Phi}|^2 c_s n_e T_e} Q_{a,k} = \left(\frac{3}{2} \frac{1}{n_e T_e} \right) \frac{\text{Re} \langle ik_y \tilde{\Phi}^* \tilde{p}_a \rangle}{\langle \tilde{\Phi}^* \tilde{\Phi} \rangle}. \quad (2.22)$$

Analogously, one can write a quasilinear expression for the particle flux of a species a as

$$\Gamma_a = \sum_k n_a c_s \frac{\text{Re} \langle ik_y \tilde{\Phi}^* \tilde{n}_a \rangle}{\langle \tilde{\Phi}^* \tilde{\Phi} \rangle} \tilde{\Phi}^2, \quad (2.23)$$

with $c_s \equiv \sqrt{T_e/m_i}$ being an approximate ion sound speed, used here simply as a dimensional factor. Eqs. 2.22 and 2.23 show that quasilinear fluxes at each wavenumber depend on the phase relation between $\tilde{\Phi}$ and \tilde{p} (for the heat flux) or \tilde{n} (for the particle flux), which is taken to be entirely determined by linear physics. This key assumption may be justified as a result of tokamak core turbulence being in a *weak turbulence* regime, as opposed to traditional fluid turbulence modelled by the Navier-Stokes equations. In TGLF, the phase relations between fluctuating fields are determined by the solution of an eigenvalue problem including 15 moments per species; of which 12 derive from a 6×2 ‘‘split moment scheme’’ for circulating particles and 3 are for trapped particles. The 6×2 moments (density n , parallel velocity u_\parallel , parallel pressure p_\parallel , total pressure p , parallel energy flux Q_\parallel and total energy flux Q) are used to estimate the circulating particle response by taking the difference between

velocity integrals over the entire velocity space and over only the trapped velocity wedge, $|v_{\parallel}| < f_t v$, where f_t is the trapped particle fraction. This is used instead of a more intuitive integral over the circulating particle region to more accurately model the kinetic density response [42]. The remaining 3 moments for trapped particles (n^t , p_{\parallel}^t , p_T^t) are derived from the bounce-averaged gyrokinetic equation assuming that the Landau resonance is averaged out ($v_{\parallel} k_{\parallel} = 0$), thus making the odd velocity moments trivially zero. The resulting eigenvalue problem is solved using a Hermite polynomial decomposition, analogous to the one described in Section 7.2.3 and Appendix B, for the poloidal variation of electromagnetic fields and neglecting the poloidal variation of the trapped fraction. The result gives γ and ω as a function of k_y , as well as the eigenmodes that are necessary for the calculation of cross-phases between fluctuating fields.

It remains to specify how the fluctuating fields saturate nonlinearly, i.e. what is the amplitude of $\tilde{\Phi}$. In TGLF, the saturation rule for $\tilde{\Phi}$ is given by an educated formulation of dependencies with a number of free parameters that are fit to fully-nonlinear models. The original TGLF ‘‘SAT-0’’ saturation rule is given by [42, 44]

$$\tilde{\Phi} = C_{\text{norm}} \left(\frac{\rho_s \hat{\omega}_{d0}}{a} \right)^2 \left(1 + \frac{T_e}{T_i} \right)^2 \left(\frac{1}{\hat{k}_y^{c_k}} \right) \left[\frac{\hat{\gamma}_{\text{net}}^{c_1} + c_2 \hat{\gamma}_{\text{net}}}{\hat{k}_y^4} \right]. \quad (2.24)$$

with the highlighted parameters (C_{norm} , c_1 , c_2 , and c_k) being free parameters, $\alpha_E = 0.3\sqrt{\kappa}$ (κ being the plasma shape squareness parameter), $\hat{\gamma}_{\text{net}} = \max[(\hat{\gamma} - \alpha_E \hat{\gamma}_E)/\hat{\omega}_{d0}, 0]$ and $\hat{\omega}_{d0} = \hat{k}_y a/R$. Since the free parameters of this model were obtained by fitting a more complete theoretical model rather than experimental data, this has been described as a *theory-based model* [42]. In this thesis, we make use of the more recent ‘‘SAT-1’’ rule [45], developed to capture interactions across ion and electron scales based on multiscale nonlinear gyrokinetic simulations [46]. Its formulation is centered around the balance of linear growth rates with nonlinear damping effects of drift-wave mixing (DWM) and zonal flow mixing (ZFM). Zonal flows are $k_y = 0$ electrostatic modes with finite radial wavenumber (k_x) which are stable linearly but are driven nonlinearly by $k_y \neq 0$ fluctuations. Near linear turbulence thresholds, $\mathbf{E} \times \mathbf{B}$ zonal flows can exceed fluctuation levels, leading to a shift of the nonlinear threshold (Dimits shift) [47]. Zonal flow shearing is now widely appreciated to regulate and partially suppress drift wave transport, also leading to the transition from low(L)- to high(H)-confinement in tokamaks [48]. However, TGLF SAT-1 was developed on the

observation that multiscale turbulence appears to be regulated by the saturation of zonal flows at the maximum of linear growth at ion scales, such that [49, 44]

$$\gamma_{ZFM} = k_y V_{ZF} = k_y \sqrt{\sum_{k_x} dk_x \left| k_x \tilde{\Phi}^2(k_x, 0) \right|} \sim k_y \max_{\text{ion scales}} (\gamma/k_y). \quad (2.25)$$

The SAT-1 rule thus defines an ‘‘effective growth rate’’ at each wavenumber based on this zonal flow mixing velocity, acting differently below and above the wavenumber of maximum growth rate, k_{max} :

$$\gamma_{\text{model}} = \begin{cases} \max [\gamma_{\text{net}} - c_2^{ZF} V_{ZF}(k_{max} - k_y), 0] & \text{for } k_y < k_{max} \\ \sum_{k'_y > k_{max}} dk'_y W(k'_y, k_y) \gamma_{\text{stage1}}(k'_y) & \text{for } k_y > k_{max}, \end{cases} \quad (2.26)$$

where we again highlighted yet-undefined parameters that are set by fitting this model to nonlinear gyrokinetic simulations. In Eq. 2.26, a normalized Lorentzian function $W(k'_y, k_y)$ is used to account for the spectral broadening at high- k due to drift-wave mixing, defining

$$W(k'_y, k_y) = \frac{g(k'_y, k_y)}{\sum_{k'_y > k_{max}} dk'_y g(k'_y, k_y)}, \quad \text{with } g(k'_y, k_y) = \frac{1}{k_y^2 + c_{k_y}(k_y - k'_y)^2}$$

The model also considers the observation that eddies become more radially elongated at electron scales (*streamers*), scaling approximately as $\sim \sqrt{k_y}$. These elements are combined in the saturation potential

$$\tilde{\Phi}_{\text{model}}(k_{x0}, k_y) = \sqrt{C_{\text{norm}}} \frac{\gamma_{\text{model}}}{k_y^2 (1 + a_y k_{x0e}^2)} \left(\frac{\gamma_{12}}{\gamma_1} \right)^{c_\gamma} \max \left[\left(\frac{k_y}{k_{\text{ETG}}} \right)^{1/2}, 1.0 \right] \quad (2.27)$$

with

$$\gamma_{\text{stage1}} = c_1^{ZF} k_{max} V_{ZF} + \max [\gamma_{\text{net}} - c_1^{ZF} k_y V_{ZF}, 0] \quad \text{and} \quad \gamma_{\text{net}} = \frac{\gamma_1}{1 + (a_x k_{x0e})^4},$$

k_{x0} being the radial wavenumber at which the finite- k_y potential is largest. We remark that SAT-1 has more than twice as many free parameters as SAT-0, but captures a wider range of physics. The addition of zonal flow mixing results in interactions between different wavenumbers, making SAT-1 nonlocal in wavenumber, while still maintaining locality in configuration space. Transport estimates may thus be obtained at any flux surface, as we will see in Chapters 7 and 8.

Research over the past decade has shown that the TGLF equations together with the SAT-0 and SAT-1 rules can capture the nature of gyrokinetic results remarkably well, particularly for heat transport [50, 51]. TGLF is now widely adopted as a transport model that can easily be compared with experimental measurements, for example as part of the TGYRO framework [52]. Nonetheless, the fact that TGLF is still a reduced model of the more complete equations and the quasilinear nature of its saturation rules raises questions about its applicability in plasmas where nonlinear interactions dominate over linear responses, as quantified by the Kubo number [53, 54]. In this thesis, we compare experimental particle transport inferences to both TGLF and the more complete gyrokinetic model, attempting to provide evidence for the accuracy (or lack thereof) of CGYRO, as well as for the reduced model of the TGLF code.

In the next section, we will describe a characterization of the main turbulent modes that are observed to dominate tokamak plasmas. Our aim here is not to offer a thorough review of the subject, but rather to introduce some of the basic phenomenology to later discuss the results of this thesis in context.

2.4.4 Characterization of Turbulent Modes in Tokamaks

To illustrate the basic principles of drift waves, consider a plasma slab with a density gradient ∇n in the direction perpendicular to \mathbf{B} . This scenario is represented in Fig. 2-1. Unlike interchange (Rayleigh-Taylor) modes, which have $k_{\parallel} = 0$, drift waves require a finite k_{\parallel} in order to develop. This allows electrons to stream along field lines and respond to cross-field perturbations. If one assumes an adiabatic response for electrons, such that $\tilde{n}_e = n_0 \exp(e\tilde{\phi}/T_e)$, any perturbation in the electron density will correspond to a perturbation of electric potential that will be perfectly in phase with the density itself. Electric fields will therefore develop such that the resulting $\mathbf{E} \times \mathbf{B}$ drift will increase the initial perturbation at locations where $v_{\mathbf{E} \times \mathbf{B}}$ and ∇n are anti-parallel and decrease it otherwise. This causes the formation of convective cells and the propagation of a (stable) wave in the electron diamagnetic direction at a speed given by $v_{d,e} = \nabla p_e \times \mathbf{B} / (n_{0,e} e B^2)$ [55].

When electrons respond adiabatically, drift waves do not drive cross-field transport, but rather result in propagation of the initial perturbation with neither growth nor decay. Growth may occur in the presence of resonant wave-particle interactions that cause a delay of the electron response with respect to the ions. This is because in this case the electron

density and electric potential will not be perfectly in phase, thus giving an $\mathbf{E} \times \mathbf{B}$ drift that amplifies the initial perturbation rather than just propagate it. This is however only the case if ∇B and ∇n are parallel, as at the low-field side (LFS) of a tokamak; on the high-field side (HFS), the opposite occurs and drift waves are normally stable. The tendency of instabilities to concentrate on the LFS is typical of modes that are described as “ballooning”.

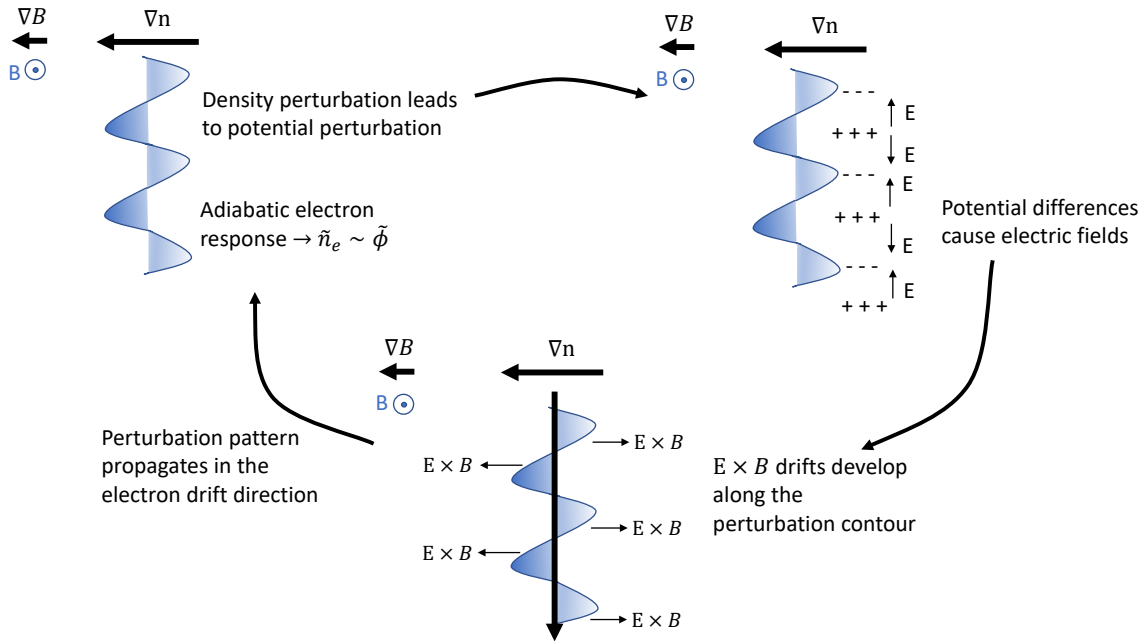


Figure 2-1: Idealized picture of density-driven drift waves in a plasma slab.

The picture above is obviously a simplification, but it nonetheless offers some degree of physical intuition. As one includes realistic magnetic field geometry, temperature gradients, passing and trapped particles, and other departures from the ideal density slab picture, it becomes harder to understand the nature of microturbulent transport in tokamaks. However, a number of modes whose principles are broadly related to the above drift wave picture have emerged as being responsible for most of the observed transport in fusion devices: these are the Ion Temperature Gradient (ITG), Electron Temperature Gradient (ETG) and Trapped Electron Mode (TEM).

ITG is widely believed to be the primary contributor to the anomalous transport levels observed in many tokamak regimes. Its main drive is the ion temperature gradient, thus




Figure 2-2: Sketch of major turbulent modes leading to transport in tokamak core plasmas, showing linear growth rates (γ) vs. perpendicular wavenumber k_{\perp} .

motivating its nomenclature [56]. Defining

$$\omega_n = \frac{T}{qB} \frac{d \ln n}{dx} k_y \quad \text{and} \quad \omega_T = \eta \omega_n = \frac{1}{qB} \frac{dT}{dx} k_y \quad (2.28)$$

as the diamagnetic drift frequencies related to densities and temperature gradients, one can write the dispersion relation for the slab version of ITG in the fluid limit $|\omega/(k_{\parallel} v_{t,i})| \ll 1$ as [57]

$$\omega \simeq \frac{1}{2} (1 + i\sqrt{3}) \left[(k_{\parallel} v_{t,i})^2 \omega_{Ti} \right]^{1/3} \implies \gamma \sim (k_{\parallel}^2 c_s^2 \eta_i \omega_{n,i})^{1/3}, \quad (2.29)$$

where we also assumed a flat density profile and no finite-Larmor-radius (FLR) effects for simplicity. Because of its dependence on the parameter $\eta_i = d \ln T / d \ln n$, this is sometimes called the “ η_i mode”.

Introducing the effects of toroidal geometry, the ITG mode transitions from the ion acoustic (sound) wave nature of the slab mode to interchange-like properties, analogous to the Rayleigh-Taylor instability that normally develops when a denser fluid is above a lighter one. If we define the average drift frequency as

$$\langle \omega_g \rangle = -\frac{2T}{qB} \frac{k_y}{R} = 2\epsilon_n \omega_n \quad \text{with} \quad \epsilon_n = L_n/R \quad (2.30)$$

one can then show that when $\langle \omega_g \rangle \ll \omega_T$ the dispersion relation for the toroidal ITG in the

long wavelength limit ($k_{\parallel} \rightarrow 0$) is [58, 56]

$$\omega = \pm (-\langle \omega_g \rangle \omega_{Ti})^{1/2} \sim k_y \rho_i v_{ti} \sqrt{\frac{\eta_i}{RL_n}} \quad (2.31)$$

This shows that an instability arises only if $\nabla T_i \cdot \nabla B > 0$, which is the case on the LFS but not the HFS.

At high k_y (short wavelength), ions are unable to respond to perturbations, hence their response is effectively adiabatic. On the other hand, electron dynamics become more important and the Electron Temperature Gradient (ETG) mode gives rise to non-negligible electron heat flux. The ETG linear response is completely isomorphic to the ion one, except for the swapped roles of particles; the nonlinear response and saturation of this mode is however radically different from the ITG one because it occurs at very different scales. ETG does not typically play a significant role in ion heat and heavy particle transport in the core of tokamak plasmas; its impact on the investigations presented in this thesis is therefore mostly negligible.

The sketch in Fig. 2-2 shows that at intermediate scales, between ITG and ETG modes, TEM can be dominant. These instabilities result from a resonant interaction between a perturbation and the toroidal precessional drift of trapped electrons. An analogous effect also occurs for ions, leading to Trapped Ion Modes (TIM), which however appear to be of little experimental relevance and will not be discussed here. TEM is often the main candidate invoked to explain electron heat and particle transport in tokamaks. Its spatial scales are contiguous with those of ITG, making the two difficult to distinguish except for different dependencies of their linear growth rates and their saturation mechanisms. A basic picture of trapped particle modes can be gained by considering the bounce-averaged toroidal precession of frequency $\langle \dot{\varphi} \rangle_b$. In the basic gradient-driven drift wave picture in Fig. 2-1 density perturbations were taken to resonate with potential perturbations; in the case of TEM, the toroidal precession frequency $\langle \omega_{\phi} \rangle_{be}$ takes the place of density fluctuations. Taking $k_{\parallel} \rightarrow 0$ as we did for the ITG dispersion relation in Eq. 2.31, we can isolate the pure TEM by neglecting parallel ion drives (taking $k_{\parallel} v_{t,i} = v_{d,i} = 0$) and setting $\langle \dot{\varphi} \rangle_b \approx \varphi$ for simplicity; doing so, one finds [58]

$$\omega \simeq \omega_{n,e} + \sqrt{\frac{3}{2} f_t \omega_{\varphi,e} \omega_{n,e} \left(1 + \frac{d \log T_e}{d \log n_e} \right)} \quad (2.32)$$

where $f_t = \sqrt{1 - B_0/B_{max}}$ is the trapped particle fraction. Eq. 2.32 shows that these waves propagate in the electron diamagnetic drift direction and their growth rate is strongly dependent on trapped particles. As we will discuss in Chapter 8, plasma shaping can strongly influence the resonance between trapped particles and electrostatic potential fluctuations, leading to non-trivial transport phenomena. More generally, TEMs can be destabilized either via collisionless processes (CTEMs) or through dissipative ones (DTEMs).

Fig. 2-1 shows that Micro-Tearing Modes (MTM) are also sometimes important at the low end of the k_y spectrum. These modes, also part of the family of drift waves, are believed to be relevant for pedestal turbulence, particularly in the electron channel [59], but are believed to be negligible for the analysis presented in this thesis and will therefore not be discussed further here.

2.5 Magneto-Hydrodynamic Modes

So far, we have focused on models where kinetic effects are important. A fluid description does however have useful applications in tokamaks, for example to describe equilibrium-scale plasma dynamics. Magneto-hydrodynamic (MHD) theory is normally derived in the standard fluid limit valid at high collisional frequency, but it may be shown to also be valid in the relatively collision-less conditions of core tokamak plasmas to describe the large-scale balance of magnetic and pressure forces [60]. For example, MHD is the appropriate model to describe magnetic flux surfaces in tokamaks. MHD also broadly explains the ultimate limit to plasma pressure in tokamaks, set by ideal non-axisymmetric kink modes, and the occurrence of tearing modes – instabilities where magnetic flux surfaces break and reconnect near rational surfaces with safety factor $q = m/n$, where m is the poloidal mode number and n is the toroidal mode number. In high- β plasmas, such as some of the highest-performance DIII-D discharges, tearing modes can be seeded by the neoclassical bootstrap current $j_{bs} \approx -(\epsilon^{1/2}/B_\theta)dp/dr$ and result in significant cross-field transport and even complete disruption of discharges [61]. Another mechanism for the seeding of neoclassical tearing modes (NTMs) is the triggering of $n = m = 1$ modes at the $q = 1$ surface, typically near the magnetic axis. Such modes are often labelled as “sawteeth” because of their periodic nature resulting in sharp crashes (on μs time scales) of electron temperature on axis, giving a jigsaw-like pattern in some diagnostic measurements. While NTMs are most concerning at high β and constitute

a primary cause of disruptions in DIII-D experiments, sawteeth appear ubiquitously across devices, including in the low- β plasmas of Alcator C-Mod. Sawtooth crashes have been known to not only cause significant cross-field transport of heat, but also flattening of density profiles. This reduces on-axis density peaking, which is typically beneficial for impurities but undesirable for main ions. The plasma radius inside of which sawteeth re-distribute heat and particles is referred to as the *sawtooth mixing radius*, r_{mix} [10, 62], and has been observed to be related to the radius of the $q = 1$ surface (the *sawtooth inversion radius*) via $r_{mix} \sim \sqrt{2}r_{q=1}$ [63]. Multiple attempts have been made at explaining the sawtooth phenomenology, usually in reference to a $n = m = 1$ magnetic reconnection effect, e.g. by Kadomtsev [64], Wesson [65] and Porcelli [66], but this remains an area of active research.

Finally, MHD effects are also believed to determine the stability of the pedestal region of tokamaks. The EPED model [67] provides a prescription to find the height and width of the pedestal pressure from the combined effects of Peeling-Ballooning (PBM) and Kinetic Ballooning Modes (KBM). Observations across a large number of tokamaks have provided clear evidence for the fact that once a pedestal grows to the peeling-ballooning stability limit, an Edge-Localized Mode (ELM) is triggered (see Ref. [68] and references therein). ELMs constitute one of the major problems associated with standard H-mode operation, since they release a large fraction of the plasma heat and momentum over a short time scale, possibly endangering the material surfaces of the device. Tokamak operation without ELMs, either obtained via direct suppression methods or in regimes that are intrinsically free from ELMs, is a major research objective. This thesis mostly addresses particle transport in the latter category of operational regimes without ELMs.

2.6 Scrape-off-Layer Transport

While SOL particle transport is outside the scope of this thesis, the penetration of neutral particles from the edge into the pedestal is a central subject in Chapter 6. SOL transport also sets radial “boundary conditions” for the core. The so-called “2-point model” [69], widely used throughout the fusion community, is applied in this work to constrain the localization of electron temperature (T_e) and density (n_e) measurements with respect to the last closed flux surface (LCFS). This enables more accurate kinetic profile fitting, ensuring that misalignment of experimental data with the magnetic equilibrium reconstruction does

not significantly affect the quality of our analysis. This is particularly important when translating Ly_α emissivities into atomic neutral D densities, as described in Chapter 6, since these strongly depend on accurate estimation of atomic rates as a function of n_e and T_e .

The two-point model is based on power balance at the LCFS. The power going into the SOL, P_{sol} , can be estimated via

$$P_{sol} = P_{\text{Ohm}} + P_{\text{aux}} - P_{\text{rad}} - \frac{dW}{dt} \quad (2.33)$$

where P_{Ohm} is the Ohmic power, P_{aux} is the auxiliary power, P_{rad} is the total radiated power within the LCFS, and W is the plasma stored energy. Assuming that 50% of the SOL power is transported to the divertors by electrons via parallel heat conduction only, one can write the parallel heat flux as $q_{\parallel} = -\frac{2}{7}\kappa_{0,e}\frac{\partial T_e^{7/2}}{\partial s}$, where $\kappa_{0,e}$ is the Spitzer-Härm electron heat conduction coefficient and s is the parallel field coordinate [70]. The simplest 2-point model assumes that all power exits the LCFS at the outer midplane (an exaggeration of ballooning transport) and decays exponentially over a width λ_q . This results in a poloidal heat flux of $q_{\theta} \approx P_{sol}/2\pi\lambda_q R$. The heat flux parallel to the magnetic field lines can then be estimated as $q_{\parallel} \approx q_{\theta}B/B_{\theta} = (P_{sol}B)/(2\pi\lambda_q RB_{\theta})$. Since λ_q is experimentally observed to decay as $1/B_{\theta}$ [16], this is usually taken to imply that q_{\parallel} scales as $P_{sol}B/R$ [71]. An alternative experimental scaling for λ_q obtained across confinement regimes on Alcator C-mod indicates $\lambda_q = (C_f/\bar{p})^{0.5}$, with $C_f = 0.06$ and \bar{p} being the volume-averaged plasma pressure [Pa] [72]. With either λ_q scaling, the 2-point model prediction for the electron temperature at the LCFS is given by

$$T_{e,LCFS} = \left(\frac{7 q_{\parallel} L}{2 \kappa_{0,e}}\right)^{2/7} \quad \text{with} \quad q_{\parallel} = \frac{\frac{1}{2}P_{sol}}{2\pi R\lambda_q \frac{B_{\theta}}{B}} \quad (2.34)$$

with L being an estimate for the SOL parallel connection length. This calculation makes a number of assumptions, not least the fact that SOL radiation and heat transport to limiter surfaces are negligible. More complete models have been developed on both C-Mod and DIII-D, but for the purposes of this work Eq. 2.34 was deemed to be sufficiently accurate.

2.7 Experimental Determination of Particle Transport

Historically, transport model validation for fusion plasmas has mostly focused on heat transport [73, 74]. 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 [75, 76, 77, 78, 79, 80]. Detailed validation of impurity transport theory is more challenging, both from an experimental and a theoretical perspective, since a wider range of physics is involved. The purpose of inferring particle transport coefficients may be seen as twofold: from a practical viewpoint, we wish to experimentally understand the viability of specific tokamak scenarios to scale them to realistic reactor conditions; from a theoretical perspective, the study of heavy ions offers a compelling test case for validation of the theory described in earlier sections of this chapter.

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 as a separation of fluxes into components arising from different thermodynamic forces. In the simplest of descriptions, one may use an “effective diffusivity”, defined as $D_{eff} = -\Gamma/\nabla n$. While such a definition makes it difficult to compare to theory, where non-diffusive terms may be important, its simplicity sometimes justifies the disadvantages. For example, if one were to assume complete dominance of turbulent transport by local low frequency drift waves with time scale of the order of the diamagnetic frequency ($\omega_{*i,e}$ for ions and electrons, respectively) and spatial scale of the order of the gyroradius ($\rho_{i,e}$), the effective diffusivity would be expected to follow the so-called *gyroBohm* scaling $D_{\perp,GB} \sim (\rho_{i,e}/L_T)T/eB$, where L_T is the temperature gradient scale length. On the other hand, whenever transport is due to events of macroscopic size, the conventional *Bohm* scaling normally applies: $D_{\perp,B} \sim T/eB$ [9]. It is interesting to note that in Alcator C-Mod B is higher and L_T is generally smaller than on DIII-D, since kinetic profiles are comparable in flux space; hence, gyroBohm transport suggests that higher field permits going to smaller device size while maintaining similar transport levels, consistently with the ρ_* scaling of transport that is generally invoked to motivate high-field devices [81, 12, 13].

Unfortunately, based on this scaling it is not obvious whether experimental particle transport inferences on C-Mod or DIII-D should generally yield lower or higher diffusivity.

To go beyond the simplest diffusive description, one may define the *transport matrix*

$$\begin{bmatrix} \frac{\Gamma_e}{n_e} \\ \frac{Q_e}{n_e T_e} \\ \frac{Q_i}{n_i T_i} \\ \frac{\Pi}{(m_i n_i v_\phi)} \end{bmatrix} = - \begin{bmatrix} D_e & \cdot & \cdot & \cdot & v_{p,e} \\ \cdot & \chi_e & \cdot & \cdot & v_{h,e} \\ \cdot & \cdot & \chi_i & \cdot & v_{h,i} \\ \cdot & \cdot & \cdot & \chi_\phi & v_\phi \end{bmatrix} \cdot \begin{bmatrix} \frac{1}{n_e} \nabla n_e \\ \frac{1}{T_e} \nabla T_e \\ \frac{1}{T_i} \nabla T_i \\ \frac{\nabla(m_i n_i v_\phi)}{(m_i n_i v_\phi)} \\ 1 \end{bmatrix}. \quad (2.35)$$

where we assumed a 2-species hydrogen plasma, where $\Gamma_i = \Gamma_e$ by quasi-neutrality. Eq. 2.35 separates particle, heat and momentum fluxes from electrons and ions into diffusive and convective coefficients. The 4×5 matrix of transport coefficients is expected to be positive definite and reflect *Onsager symmetries* so as to increase entropy [82]. One can easily extend Eq. 2.35 to not only describe fluxes for electrons and main ions, but also other species.

Diffusive-Convective Ansatz

While the kinetic equation 2.1 and its reduced models are nonlinear in electromagnetic potentials, they are linear in the distribution function itself. In the *trace limit* of a particle species being at low enough concentration as to avoid any nonlinear feedback with electromagnetic fields, its presence is negligible in the Poisson and Ampère equations. This makes fluxes exactly linear in the thermodynamic gradients and thus justifies a diffusive-convective ansatz. This can easily be verified with both neoclassical (e.g. NEO) and gyrokinetic (e.g. CGYRO) simulations. In more realistic finite-density conditions, this linearity breaks down and this ansatz should be interpreted as simply a convenient parametrization of off-diagonal terms in the transport matrix.

Let us write the particle flux of a given species as $\Gamma = -D \nabla n + v n$, with D being a diffusion coefficient and v a convection one. In quasi-steady conditions with no particle core sources, one may expect $\Gamma = 0$, allowing us 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 [83] or via modulation of sputtering sources from external antennae [84], which however do not generalize to higher-Z impurities for which atomic physics plays a key role.

Efforts to estimate radial profiles of D and v have been reported from all major devices, including ASDEX-Upgrade [85, 86, 33, 84], DIII-D [87, 88], JET [89, 90], Tore Supra [91, 92], TCV [93], HL-2A [94], W7-X [95], MAST [96], NSTX [97] and Alcator C-Mod [98, 99, 100, 101, 102, 103]. 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 [20] allows the injection of non-recycling impurities via a time-resolved and non-perturbative laser ablation, whose source function can be more easily characterized. In this thesis, experimental particle transport coefficients are always estimated via LBO injections, for both C-Mod and DIII-D.

The diffusive-convective model described above is often applied to a 1-dimensional description of particle transport. In this thesis, such description refers to *flux-surface averaged* (FSA) quantities, which may be defined on a 1D radial grid based on normalized magnetic flux. However, theoretical transport models typically predict transport at the low-field side (LFS). To compare these results, one may transform LFS quantities (labeled with a “0” subscript) on a minor radius coordinate r_m to those from flux surface averaging on an arbitrary r coordinate using the transformations [78]

$$D = D_0 \frac{n_0}{\langle n \rangle} \left(\frac{\partial r}{\partial r_m} \right)^2 \quad (2.36)$$

$$v = v_0 \frac{n_0}{\langle n \rangle} \frac{\partial r}{\partial r_m} + D \frac{\partial}{\partial r_m} \left(\ln \frac{n_0}{\langle n \rangle} \right)$$

where factors of $n_0/\langle n \rangle$ account for poloidal density asymmetries, expected to be significant for heavy ions, particularly near the LCFS [104, 105]. The $n_0/\langle n \rangle$ ratio may be estimated using a neoclassical code or, more simply, an analytical prediction of centrifugal effects [105]. 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. [78].

Poloidal asymmetries of impurity densities may arise from processes other than centrifugal forces. The concentration of neutral particles near the x-points of divertor

configurations may suggest a clear way to produce poloidal asymmetries, but as we will see in Chapter 6 this effect is averaged out over flux surfaces. Experimental observations have shown in-out impurity asymmetries to arise due to fast ions being produced by Ion Cyclotron Resonance Heating (ICRH) [104], as well as large electron density pedestals [106]. Strong up/down asymmetries of x-ray brightness are also routinely observed on Alcator C-Mod on the side opposite to the ion ∇B direction, exhibiting qualitative features of neoclassical transport but yet unexplained physics [98, 107]. This (likely incomplete) set of asymmetry mechanisms may lead to skepticism as to the value of inferring FSA impurity densities. It is however important to realize that detailed 2D models of tokamak plasmas should always agree with FSA quantities which cannot resolve poloidal asymmetries; in this sense, our 1D estimates offer a necessarily limited but useful description of complex transport processes. From an experimental perspective, the separation of particle fluxes into D and v is the most detailed description that has been successfully attempted to date. From a theoretical modeling standpoint, however, it is useful to separate particle fluxes further into more terms, filling in the gaps in the transport matrix in Eq. 2.35, via [78, 108, 109]

$$\frac{R \Gamma n}{n} = D_z \frac{R}{L_n} + D_T \frac{R}{L_T} + R v_p \quad (2.37)$$

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 the D described previously), a thermodiffusion term D_T and a pure convection term v_p . For clarity, we also define $v_T = D_T/L_T$, 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 (“rotodiffusion”), e.g. Refs. [110, 111]. We shall not consider this explicitly since having species with different toroidal rotation is not possible in CGYRO and this prevents us from rigorously separating this term from others. Rotodiffusion is thus “grouped” in our analysis together with other convection terms.

The quantification of diffusion and convection coefficients requires detailed understanding of particle sources and sinks. This brings us to the subject of the next chapter, where we will introduce key concepts in spectroscopy that allow us to quantify the transport-induced deviation of particle distributions from expected atomic equilibria.

3

Principles of Plasma Spectroscopy

In this chapter we review the basic ideas of plasma spectroscopy, particularly focusing on aspects that are central to transport inferences in tokamaks. We introduce concepts and equations related to ionization equilibria, collisional-radiative modeling, and ion radiation. Finally, we describe K_{α} x-ray spectra that will be central to the Alcator C-Mod analysis in this thesis.

Introduction

Following our discussion on plasma transport in Chapter 2, we now shift our attention to the principles of atomic physics and spectroscopy that are most relevant for the research in this thesis. Spectroscopy is the field that examines spectra produced by the interaction of matter with electromagnetic radiation or emission of radiation by matter itself. Experimental measurement of spectra from plasmas, either partially or fully ionized, can provide a wealth of information about the composition and conditions of interest.

At low densities, typical of the solar corona, the detailed balance of atomic processes is often referred to as “coronal equilibrium” (CE). In this situation, excited states are effectively only populated via collisional excitation and de-populated via radiative decay. This enables important simplifications with respect to more complete atomic models. The opposite limit of high-density conditions is referred to as “local thermodynamic equilibrium” (LTE). In LTE, collisional processes are completely dominant, charge state densities are given by the

Saha equation and the population of excited states by the Boltzmann distribution [112]. Magnetically-confined fusion plasmas live in-between these limits, where computation of atomic rates is unfortunately more challenging. At these densities ($10^{19} \text{ m}^{-3} \lesssim n_e \lesssim 10^{24} \text{ m}^{-3}$), one must account for the complex competition of radiative and collisional processes. This is the domain of “collisional-radiative” modeling [113], or “generalized collisional-radiative” (GCR) modeling when considering the role of metastable states [114], as described in Section 3.3. The Atomic Data and Analysis System (ADAS) [114] framework used in this work offers a collection of atomic data and routines for GCR modeling in fusion plasmas.

Separation of Atomic Timescales

One important concept in GCR theory is the separation of *ordinary* and *metastable* atomic states. Electrons in metastable states have, by definition, much longer lifetimes than the those in ordinary states, which are often assumed to decay instantaneously. In ADAS nomenclature, the ground state of an atom or ion is included among the metastable states. Metastables hold most of the electrons and so it is assumed that the only recombination events that matter are those having a metastable as a collision partner.

More generally, lifetimes in GCR modeling are divided into two groups. The first is the *intrinsic* group, which depends only on atomic parameters. This includes radiative decay from metastable states, τ_m , ordinary excited state radiative decay, τ_o , and auto-ionizing state decay (radiative and Auger), τ_a . The intrinsic group for a particular ion is generally ordered as $\tau_a \ll \tau_o \ll \tau_m$, with typical values in fusion plasmas being $\tau_m \sim 10^1/Z^8 \text{ s}$, $\tau_o \sim 10^{-8}/Z^4 \text{ s}$, $\tau_a \sim 10^{-13} \text{ s}$, with Z being the ion charge. The second group is the so-called *extrinsic* one, which depends on plasma conditions, especially particle density. It includes free particle thermalization (including electron–electron τ_{ee} , ion–ion τ_{ii} and ion–electron τ_{ie}), charge-state change (ionization τ_{ion} and recombination τ_{rec}) and redistribution of population among excited ion states (τ_{red}). The extrinsic group is ordered in general as $\tau_{ion,rec} \gg \tau_{ie} \gg \tau_{ii} \gg \tau_{ee}$. Of course, the degree to which atomic modeling accounts for all these effects defines the validity of results within the field of plasma spectroscopy [114].

Spectral Line Emission

The intensity of an emitted atomic line due to a transition between two states i and j , with energies $E_i > E_j$, is given by

$$I_{ij} = N_i A_{ij} \hbar \omega_{ij}, \quad (3.1)$$

where N_i is the population density of the upper state i , A_{ij} is the transition probability (Einstein coefficient) for the specific transition, $\omega_{ij} = \Delta E_{ij}/\hbar = (E_i - E_j)/\hbar$ is the frequency associated with the energy difference and $\hbar = h/2\pi$ is the reduced Planck constant. Einstein coefficients generally depend on the ion charge and both the initial and final quantum numbers. For relatively simple ion configurations, the interaction between electrons is normally stronger than their individual spin-orbit interaction, making the orbital, L , and spin, S , atomic quantum numbers be independent. This gives rise to the concept of L-S coupling, typically appropriate for light atoms, while for heavy ions the more complex j-j coupling is needed. For configurations that obey L-S coupling, the transitions with the largest Einstein coefficients are those that obey electric dipole selection (Russell-Saunders) rules and are said to be (optically) “allowed”. Other kinds of transitions, e.g. magnetic dipole or electric quadrupole transitions, also contribute to atomic spectra, but are typically much weaker and referred to as “forbidden” [112]. The lifetime of a specific excited state, τ_i , is determined by the inverse of the total probability of that state population decaying to any of the states with lower energy, i.e. $\tau_i = A_i^{-1} = (\sum_{j<i} A_{ij})^{-1}$. The value of A_i generally scales as $A_i \approx 1.6 \times 10^{10} Z^4 n_i^{-9/2} [\text{s}^{-1}]$, with n_i being the principal quantum number of the upper level i [115].

3.1 Fundamental Atomic Processes in Plasmas

In what follows, we briefly summarize the main atomic processes of interest for plasma modeling; a detailed review is available in Ref. [114]. Unless otherwise stated, throughout this thesis we make no distinction between hydrogen and deuterium, since their atomic signatures are effectively identical for our purposes.

Electron impact excitation (inverse: de-excitation) is often the primary population mechanism of excited ion states, described by



In the (de-)excitation process, the incident electron loses (gains) some energy too. Among the most common methods to model this is the Born approximation, where free electrons are modeled as plane waves. In the distorted wave (DW) approximation, incident and scattered electrons as individual wave functions distorted by the nuclear potential. Finally, in the close-coupling approximation, individual bound electrons and their scattering potentials are all considered.

Spontaneous emission occurs when an electron decays from a higher to a lower energy state, emitting a photon whose frequency reflects the energy difference. The process may be written as

$$X^z(i) \rightarrow X^z(j) + h\nu. \quad (3.3)$$

The coronal approximation described above, valid in the low-density limit of a plasma, describes the ion balance of charge states solely on the basis of electron impact excitation and spontaneous emission. At higher densities, typical of tokamaks, the following processes can also be important.

Radiative recombination (inverse: photoionization) is a process which takes place when a positively charged ion captures an electron to one of its orbitals with a simultaneous emission of a photon:

$$e^- + A(q+) \rightarrow A(q-1) + h\nu. \quad (3.4)$$

Dielectronic recombination (inverse: autoionization) is a process in which the capture of an electron into an excited state also promotes a bound electron to another bound orbit:

$$e^- + A(q+) \rightarrow A(q-1)** \rightarrow A(q-1) + h\nu. \quad (3.5)$$

In the second step of the dielectronic recombination process a photon is emitted with a wavelength that is characteristic to the doubly excited state (**) of the q-1 times ionized ion. Dielectronic recombination is a resonant process, because of the discrete energy nature of the bound electron orbits.

Electron impact ionization (inverse: 3-body recombination) is caused by the impact of an electron on an ion with energy greater than the ionization potential of the target:

$$X^z + e^- \rightarrow X^{z+1} + e^- + e^-. \quad (3.6)$$

While the basic process of electron impact ionization is the same as for electron impact excitation, the presence of two unbound (continuum) electrons in the ionized system renders the calculation more complex [116].

Ion-ion excitation and ionization (inverse: 3-body de-excitation and recombination with an electron and a proton) are usually only considered for main plasma ions (from hydrogen isotopes), written in the following as p^+ :



The energy transferred during ion-ion collisions depends on the ratio of masses, and since m_e is at least 1,800 times smaller than m_{ion} , these processes are often negligible. At few-keV temperatures, hydrogen is mostly ionized by electron collisions, but ion collisions have to be considered for excitation [117]. At 10-100 keV energies, both the ionization and excitation of hydrogen atoms occur mostly via ion collisions and charge exchange with neutrals.

Charge exchange (CX) recombination is the process by which an electron is transferred from a neutral atom (donor) to a negatively-charged ion (receiver), with no net momentum change.



A key feature of this process is that the electron capture occurs into a narrow range of principal quantum numbers, n_c , which depend on the receiver's charge Z [118]. This can be understood as a result of energy conservation, resulting in $n_c \sim Z^{3/4}$ [115]. In a semiclassical picture, this condition results in the captured electron being far from the nucleus and the ion charge being partly shielded by other core electrons. The captured electron thus behaves as an isolated charge with hydrogenic behavior, referred to as a *Rydberg state*. These high- n states emit radiation on finite (ordinary) lifetimes, usually via spontaneous de-excitation. Measurements of this emission will be later examined more closely in the context of DIII-D (chapter 8) and Alcator C-Mod (chapter 6).

The standard Rydberg formula for the energy level differences of hydrogen is

$$\frac{1}{\lambda} = Z^2 R_\infty \left(\frac{1}{n_1^2} - \frac{1}{n_2^2} \right) \quad (3.10)$$

Transitions for which Eq. 3.10 applies have wavelengths scaling with Z^{-2} and n^2 . This is referred to as the *equivalence principle* or *reversibility formula* and is often used in codes dealing with neutral beam physics, e.g. FIDASIM [119] (see chapters 6 and 8), to scale cross sections via

$$\sigma(n_f \rightarrow n_i) = \frac{E_i}{E_f} \frac{n_i^2}{n_f^2} \sigma(n_i \rightarrow n_f) \quad (3.11)$$

The CX process is central to some of the contributions of this thesis. We shall consider both resonant CX between neutral and ionized hydrogen, as well as non-resonant CX between neutral hydrogen and heavier ions.

3.2 Ionization Equilibrium

In a quasi-steady plasma, the fraction of an ion's population in each charge state may be approximated as being time independent. This condition is labelled *ionization equilibrium*, and corresponds to the balance of all ionization and recombination processes *in the absence of transport*. At typical tokamak densities, ionization is mostly due to electron impact, but recombination can occur through radiative and dielectronic processes, as well as by charge exchange with neutral particles. For an ion I in charge state Z , we can write the balance of atomic sources and sinks as

$$\begin{aligned} Q_{I,Z} = & -(n_e S_{I,Z} + n_e \alpha_{I,Z} + n_n \alpha_{I,Z}^{cx}) n_{I,Z} \\ & + n_e S_{I,Z-1} n_{I,Z-1} \\ & + (n_e \alpha_{I,Z+1} + n_n \alpha_{I,Z+1}^{cx}) n_{I,Z+1} \end{aligned} \quad (3.12)$$

where S stands for effective ionization rates, α for radiative and dielectronic rates, and α^{cx} are charge exchange rates. In this expression, S and α multiply electron density, n_e , whereas α^{cx} multiplies the background neutral density of hydrogen isotopes, n_n . At equilibrium we have $Q_{I,Z} = 0$. In the absence of transport, Eq. 3.12 uniquely determines the ion population in each charge state. This thesis is focused on investigating how fusion plasmas deviate from this condition, which we therefore carefully define in the following paragraphs.

Let us combine all recombination processes (given above by α and α^{cx} into a single recombination coefficient α_T . The balance of all terms in Eq. 3.12, corresponding $Q_{I,Z} = 0$, defines the condition of ionization equilibrium between neighboring stages. By a simple

process of induction, one may prove that this gives $n_{Z+1}/n_Z = S_Z/\alpha_{Z+1}$. To see this, we first set $Z = 0$, obtaining $-S_0n_0 + \alpha_1n_1 = 0$, which does obey this given relation. Then, assuming the relation to be valid for $Z = i$, i.e. $n_{i+1}/n_i = S_i/\alpha_{i+1}$, one finds for $Z = i + 1$

$$\begin{aligned} -(S_{i+1} + \alpha_{i+1})n_{i+1} + S_in_i + \alpha_{i+2}n_{i+2} &= 0 \\ -(S_{i+1} + \alpha_{i+1})n_{i+1} + \alpha_{i+1}n_{i+1} + \alpha_{i+2}n_{i+2} &= 0 \\ -S_{i+1}n_{i+1} + \alpha_{i+2}n_{i+2} &= 0 \end{aligned}$$

Hence, we find $n_{i+2}/n_{i+1} = S_{i+1}/\alpha_{i+2}$, proving the initial proposition. The accuracy of this balance depends on the number of atomic processes that are accounted for within S and α , as well as the accuracy of their individual calculations.

To estimate the time over which ionization equilibrium is attained, let us consider the balance of reaction rates for all charge states:

$$\begin{bmatrix} -S_0 & \alpha_{T,1} & \dots & & & & \\ +S_0 & -S_1 - \alpha_{T,1} & +\alpha_{T,2} & & & & \\ \vdots & +S_1 & \ddots & & & & \\ & & & \alpha_{T,Z-1} & & \vdots & \\ & & & +S_{Z-2} & -S_{Z-1} - \alpha_{T,Z-1} & \alpha_{T,Z} & \\ & & & \dots & +S_{Z-1} & -S_Z - \alpha_{T,Z} & \end{bmatrix} \begin{bmatrix} n_0 \\ n_1 \\ \vdots \\ n_{Z-1} \\ n_Z \end{bmatrix} = 0 \quad (3.13)$$

The singular values of this matrix correspond to the time scales associated with the equilibration process. The inverse of the largest of these rates may be interpreted as an *atomic relaxation time*, τ_{relax} . In Fig. 3-1 we show this quantity for a number of atomic species over a wide range of electron temperatures. A typical tokamak core density of $n_e = 10^{20} \text{ m}^{-3}$ was used. In real experimental scenarios, the balance of τ_{relax} and the transport time scale determines to what extent ionization equilibrium is actually achieved. Helium (He) is observed to have a short relaxation time at temperatures relevant to core plasmas, while Ar generally takes longer, except at low temperatures. Note that these relaxation times are computed across all charge states of each ion; they are therefore dependent on ionization and recombination rates of all charge states.

Unfortunately, since global confinement times in both Alcator C-Mod and DIII-D are typically of similar order of magnitude to τ_{relax} , ranging from ~ 10 ms in low-confinement regimes to much larger values in high-confinement ones, it is not possible to separate




Figure 3-1: Relaxation times for a number of ions as a function of electron temperature at fixed electron density of 10^{20} m^{-3} . For simplicity, these calculations included effective ionization and recombination rates excluding charge exchange.

transport and atomic time scales in a perturbative approach. Consequently, the ionization equilibrium defined by Eq. 3.13 is generally insufficient to describe the charge state balance in a tokamak plasma. One must therefore determine transport levels in order to quantify particle concentrations and radiation patterns. In the next section, we shall consider how the atomic processes described in Section 3.1 can be modelled within an effective theoretical framework to estimate the ionization and recombination rates discussed in this section, as well as radiated power components.

3.3 Generalized Collisional Radiative Coefficients

Eq. 3.1 shows that the emissivity of a spectral line is linearly proportional to the density of electrons in the upper level of a given transition. Hence, by estimating the probability of electrons being in this level via quantum mechanics, one can compute the density of the emitting ion species. This principle enables the investigation of ion transport by examining spatio-temporal patterns of radiation. Within the framework of generalized GCR theory, this is done via *photon emissivity coefficients* (PECs). As we shall see below, the calculation of PECs is intimately connected to the quantification of the ionization and recombination coefficients discussed in the previous section.

In giving a brief overview of GCR theory, we follow the ADAS nomenclature [114], using Greek characters to identify metastable levels and Latin characters for ordinary levels. Only recombination events starting from metastable levels are considered, taking the population of excited states to immediately decay (*quasi-static assumption*). The coupling between

different atomic levels i and j is denoted by C_{ij} . The *collisional-radiative matrix* for a transition $i \rightarrow j$ is formed by each of these elements, with the addition of ionization and recombination processes with adjacent charge states.

We consider electron transitions between a recombining ion of charge $Z+1$ (parent), a recombined ion of charge Z (child) and an ion of charge $Z-1$ (grandchild). Indices ρ and σ will be used for the metastables of the child ion, ν and ν' for the parent, and μ and μ' for the grandchild. We are interested in examining transitions from ordinary (excited) states, denoted with indices i and j . We then define recombination coefficients $R_{\mu,\sigma}$ between metastable levels μ and σ , and $r_{i,\nu}$ between parent metastable and child ordinary levels. Direct ionization coefficients from child ordinary levels to parent metastable levels are labelled $S_{\nu,i}$. The time evolution of the population number densities for each of these populations can now be written as [114]

$$\frac{d}{dt} \begin{bmatrix} N_{\mu}^{-} \\ N_{\rho} \\ N_i \\ N_{\nu}^{+} \end{bmatrix} = \begin{bmatrix} C_{\mu\mu'} & N_e \mathcal{R}_{\mu,\sigma} & 0 & 0 \\ N_e \mathcal{S}_{\rho\mu'} & C_{\rho\sigma} & C_{\rho j} & N_e r_{\rho\nu'} \\ 0 & C_{i\sigma} & C_{ij} & N_e r_{i\nu'} \\ 0 & N_e S_{\nu\sigma} & N_e S_{\nu j} & C_{\nu\nu'} \end{bmatrix} \begin{bmatrix} N_{\mu'}^{-} \\ N_{\sigma} \\ N_j \\ N_{\nu'}^{+} \end{bmatrix}. \quad (3.14)$$

Here, the ordinary populations of the ions of charge $Z-1$ and $Z+1$ were absorbed into the population of their metastable states. We also assumed that the ionization of stage $Z-1$ only contributes to the metastable population of the state with charge Z .

The aforementioned quasi-static assumption translates into the condition $dN_i/dt \rightarrow 0$, which can be used in Eq. 3.14 to find the rate of change of the metastable population of the Z -times ionized ion:

$$N_j = -C_{ij}^{-1} C_{i\sigma} N_{\sigma} - N_e C_{ji}^{-1} r_{i\nu'} N_{\nu'}^{+} \quad (3.15)$$

Here, the first term refers to ionization and the second to recombination. Substituting this into Eq. 3.14, one finds for the metastable population of the state of charge Z

$$\frac{dN_{\rho}}{dt} = N_e [S_{\rho\mu'}] N_{\mu'}^{-} + \underbrace{[C_{\rho\sigma} - C_{\rho j} C_{ji}^{-1} C_{i\sigma}]}_{\text{QCD}_{\sigma \rightarrow \rho}} N_{\sigma} + N_e \underbrace{[r_{\rho\nu'} - C_{\rho j} C_{ji}^{-1} r_{i\nu'}]}_{\text{ACD}_{\rho \rightarrow \nu}} N_{\nu'}^{+}. \quad (3.16)$$

Similarly, substitution of Eq. 3.15 into Eq. 3.14 for the metastable population of the ion of

charge $Z+1$ gives

$$\frac{dN_v^+}{dt} = N_e \underbrace{\left[S_{v\sigma} - S_{vj} C_{ji}^{-1} C_{i\sigma} \right]}_{\text{SCD}_{\sigma \rightarrow \nu}} N_\sigma + \underbrace{\left[C_{vv'} - N_e^2 S_{vj} C_{ji}^{-1} r_{iv'} \right]}_{\text{XCD}_{\nu' \rightarrow \nu}} N_{\nu'}^+ \quad (3.17)$$

The above expressions show the definition of some of the GCR coefficients, following the ADAS nomenclature. In particular, the ACD and SCD coefficients, summed over all ordinary and metastable states, provide the effective recombination and ionization rates needed in this work to examine charge state balance, as described in Section 3.2. We shall not consider metastable states other than the ground state itself, hence making the QCD and XCD cross-coupling coefficients unnecessary. Once the population in ordinary states is computed via Eq. 3.15, one may then write the emissivity of the spectrum line $j \rightarrow k$ as

$$\epsilon_{j \rightarrow k} = A_{j \rightarrow k} N_e N_j = A_{j \rightarrow k} \left(\sum_\sigma \mathcal{F}_{j\sigma}^{(exc)} N_e N_\sigma + \sum_{\nu'=1}^M \mathcal{F}_{j\nu'}^{(rec)} N_e N_{\nu'}^+ \right). \quad (3.18)$$

which allows us to define PEC components as

$$PEC_{\sigma, j \rightarrow k}^{(exc)} = A_{j \rightarrow k} \mathcal{F}_{j\sigma}^{(exc)} \quad PEC_{\nu', j \rightarrow k}^{(rec)} = A_{j \rightarrow k} \mathcal{F}_{j\nu'}^{(rec)} \quad (3.19)$$

Here the \mathcal{F} factors represent the processes by which the line upper state is populated, either by ionization of a lower level or by recombination from an upper level. By setting some terms to 0 in the GCR matrix in Eq. 3.14, we dropped PEC terms related to the ionization of the ion of $Z-1$ charge, which however may be non-negligible in some cases, as described in Section 3.5. However, under most circumstances the excitation-driven PEC components are by far the dominant ones. In this work, PEC values were computed using ADAS datasets (“ADF04” files) which provide all the required terms in the GCR matrix, Eq. 3.14. The solution to the GCR problem was either obtained via the `adas208` code or the `ColRadPy` solver [120].

3.4 Radiated Power

Having defined some of the GCR coefficients in the previous section, we now consider how these enter the calculation of the total power radiated, P_{rad} , by a plasma. The main contributors to P_{rad} are the following.

Bremsstrahlung (from the German for “breaking radiation”) arises from electron-ion encounters. The continuum power from collisions with an ion species “i” of charge Z can be written as

$$P_{brems}(\nu) = 5.03 \times 10^{-54} n_e n_i Z^2 \left(\frac{T_e}{eV} \right)^{-1/2} e^{-h\nu/T_e} \bar{g}_{ff} \quad (3.20)$$

where all quantum mechanical corrections to the classical result from electromagnetism have been absorbed into the *Gaunt factor* \bar{g}_{ff} [121, 115]. Here, densities are expressed in m^{-3} and temperature is in keV. This formula only strictly applies in spectral regions where there are no radiative recombination contributions, which is typically the case in the visible range. In a realistic plasma where there exist multiple ion species, the total bremsstrahlung is obtained by summing over contributions from each ion. Eq. 3.20 shows that, given some estimates of n_e and T_e , measurement of bremsstrahlung radiation as a function of frequency then allows the quantification of an effective ion charge, $Z_{\text{eff}} = \sum_i (n_i Z_i^2) / \sum_i (n_i Z_i)$. In carbon-wall devices like DIII-D, Z_{eff} is normally dominated by the carbon contribution, hence indirectly permitting an approximate estimation of the carbon density, assuming that carbon is entirely ionized.

Synchrotron radiation is the result of continuous acceleration of electrons as they gyrate about their guiding centers. While this radiation component is usually not a major loss channel in current devices, its scaling with $B^{2.5}$ and volume [122] makes it important in future high-field devices [123].

Line radiation is the result of bound-bound electron transitions. The total line radiation from an ion at a given density and temperature can be computed by summing the product of all PECs with the corresponding energy gap for all transitions. Within ADAS, this sum is labelled as PLT data.

Recombination radiation is the result of electron capture (free-bound processes). Within ADAS, this is grouped together with bremsstrahlung, giving data that is labelled as PRB data.

The definition of PLT and PRB data enables convenient integration of atomic and plasma physics models, by separating the radiated power components that multiply the density of a child charge state and those that multiply the parent charge state.

One common simplification for plasma modeling is to assume that the charge state densities of a given ion are sufficiently close to those given by the ionization equilibrium

(Section 3.2), i.e. ignore the effect of transport in bringing charge states out of ionization equilibrium. Under this approximation, one may combine PLT and PRB data with charge state fractional abundances to compute $L(T_e) = \sum_z P_{rad,z}/(n_e n_z)$ [W·m³], which gives *cooling curves* that are only a function of T_e . These functions represent the total radiation that is expected from an ion at ionization equilibrium. Fig. 3-2 shows cooling curves for a number of ions, computed using PLT and PRB data provided by T. Pütterich based on the work of Ref. [15]. For each ion, the fractional abundance of each charge state was obtained from ACD and SCD data from ADAS. Analogous cooling coefficients may be computed for filtered radiation, providing effective atomic data to compare to soft x-ray diagnostics (Section 4.5). Curves such as those in Fig. 3-2 offer the simplest way to include the effects of radiated power in integrated modeling. As part of this work, the data in Fig. 3-2 were included into the TGYRO integrated modeling code [52] for a public release. However, it is well known that ion transport can cause significant deviations of charge state fractional abundances and distributions from the expectations of ionization equilibrium, as will also be demonstrated in following chapters. This motivates the development of simple-to-use software to interface atomic and plasma physics models with integrated modeling frameworks, as discussed in Section 5.3.1.

In a reactor plasma at equilibrium, with equal mixtures of D and T, fusion alone should provide enough heating (from the slowing down of alpha particles) to balance radiated power as well as power lost via transport, i.e. $P_{fus} = P_{rad} + P_{transp}$. For a given “lumped” impurity species *imp*, one may write such balance as

$$\underbrace{\frac{(n_D + n_T)}{2} \langle \sigma v \rangle E_\alpha}_{\text{fusion heating}} = \overbrace{c_{D+T} n_e^2 L_H + c_{imp} n_e^2 L_{imp} + c_{He} n_e^2 L_{He}}^{\text{radiation losses}} + \underbrace{\frac{\frac{3}{2} k_B T n_e c_{tot}}{\tau_E}}_{\text{transport losses}} \quad (3.21)$$

where we separated the radiation from He, since this will be a direct product of fusion reactions, and approximated the global effect of heat transport via a single heat confinement time scale, τ_E . Based on this expression and using estimated cooling factors, one may compute the maximum central concentration of any ion in a burning plasma. For example, for tungsten on ITER, setting the operational limit to be at the point where radiative losses reach 10% of the heating power density (~ 0.5 MW/m³), then for central electron density of $n_e = 1.4 \times 10^{20}$ m⁻³ and temperature $T_e = 20$ keV, one finds $L_z \approx 5 \times 10^{-32}$ W/m³




Figure 3-2: Cooling factors over electron temperature for a variety of ions of interest for fusion modeling. Data derived from Pütterich et al. [15] and selected ACD and SCD data from ADAS [114].

and a maximum W concentration of $c_W^{max} \approx 5 \times 10^{-5}$. This is obviously an extremely small number, showing that the plasma must be maintained at high purity levels. Note that such amount of W is negligible from the standpoint of fuel dilution, since $Zn_z/n_e \approx 3 \times 10^{-3}$ for W^{64+} ; other low-Z impurities in the plasma are expected to contribute more to dilution and less to radiation [19]. Based on these simple estimates, it is clearly imperative that we understand how to prevent unfavorable transport of high-Z ions to avoid a radiative collapse and that we limit core concentrations of low-Z ions to minimize fuel dilution.

3.5 X-Ray Spectroscopy

The x-ray region of the electromagnetic spectrum holds invaluable information about high-energy atomic transitions in plasmas. In this section, we introduce concepts that are important to the x-ray spectroscopy work in this thesis, leaving a description of related diagnostics to Section 4.4.

X-ray spectroscopy sees its origins in the work of W.L. and W.H. Bragg, jointly awarded the 1915 Nobel Prize for their use of high-energy electron beams to excite x-ray emission from a variety of crystals [124]. This made it possible to study the atomic structure of these materials because the electron beam energies were comparable to the energy gaps in the examined atomic structures. The same principle applies in plasmas, where high-energy electrons may excite core atomic electrons, enabling examination of radiation emitted by their subsequent decay. Measurements of x-ray spectra often rely on the principle of Bragg diffraction at gratings, described by the common formula $n\lambda = 2d \sin \theta$, with n being the diffraction order, λ a certain wavelength, d the spacing between the diffraction crystal atomic layers, and θ the diffraction angle.

In this thesis, high-resolution x-ray spectra are analyzed in detail, examining line ratios and broadening profiles, described in the next section. The K_α spectrum has a central role in our analysis on Alcator C-Mod and is therefore discussed below. For clarity, we first remind readers of some standard spectroscopic notation.

Spectroscopic Notation #1

It is typical to use roman numbers to indicate the ion charge plus one, e.g. He I is neutral helium; on the other hand, numerals are used to represent the ion charge itself, e.g. He¹⁺ is the first ionized stage of helium.

Spectroscopic Notation #2

We label atomic levels as $n l^x 2S+1 L_J$, where l is the orbital of the electron (s, p, d, etc.) and x is the number of electrons in that orbital (1 or 2 for the s orbital, 1 to 6 for the p orbital, etc.). $n l^x$ is referred to as the *electronic configuration* of the outermost electron shell.

Spectral broadening

Multiple processes may contribute to the shape of measured spectral lines; among these, we shall discuss Doppler broadening and instrumental effects as being most important. We take pressure broadening, also known as collisional or Stark broadening, and natural broadening, due to the natural lifetime of atomic excited states, as being negligible for the research discussed in this work. Doppler broadening results from the thermal spread of energies of

the ions whose emission we are interested in, hence its estimation can give a measure of ion temperature, as will be discussed in Chapter 4. Doppler broadening may be simply modeled by using [125]

$$\theta(\nu) = \frac{1}{\sqrt{\pi}\Delta\nu_D} e^{-\left(\frac{\nu-\nu_0}{\Delta\nu_D}\right)^2} \quad (3.22)$$

with the Doppler profile half-width being

$$\Delta\nu_D = \frac{1}{\nu_0} \sqrt{\frac{2T_i}{m}} \quad (3.23)$$

The numerator in the exponential of Eq. 3.22 is the difference between the rest and measured frequencies of a certain transition, referred to as *Doppler shift*. One can write this as a wavelength difference as

$$\Delta\lambda_D = \lambda \cdot \left(1 - \frac{v \cdot \cos(\alpha)}{c}\right) \quad (3.24)$$

where v is the plasma velocity and α is the angle between the line-of-sight and the direction of plasma rotation. Knowing α for a given diagnostic and measuring $\Delta\lambda_D$, one may therefore estimate the bulk velocity of the emitting species.

Precise measurements of Doppler broadening are often made more difficult by instrumental effects, an undesirable result of diagnostic imperfections giving photon counts at slightly different energies from those truly incident on detectors. It is often reasonable to model instrumental effects as giving a Gaussian profile; the overall broadening shape is then the convolution of the Doppler-induced Gaussian shape with the instrumental one, giving another Gaussian whose width is the result of adding Doppler and instrumental widths in quadrature [115].

3.5.1 K_α Spectra

When an electron transitions from the 2p orbital of the second (L) shell (principal quantum number of 2) to the innermost (K) shell (principal quantum number of 1), the resulting line emission is referred to as K_α radiation, in Siegbahn notation. This usually corresponds to the strongest x-ray spectral features. The K_α line of hydrogen-like ions is often called Ly_α (pronounced as “Lyman alpha”). In Sections 6.2 and 6.3 we shall discuss Ly_α radiation from D edge neutrals, which due to the small charge of hydrogen actually falls into the ultraviolet range.

In this section, we shall focus on K_α radiation arising from He-like ions, where one of the two bound electrons is always in the ground state. He-like ion spectra in the x-ray region are dominated by the resonance line w, the forbidden line z, and the intercombination lines x and y, identified using Gabriel notation [126]. Between these lines, there exists a large number of Li-like ion lines that are labelled as *satellite lines* of the He-like ion lines, since their wavelengths are only slightly shifted from the w, x, y and z lines. These lines correspond to transitions from doubly-excited states, where one electron effectively acts as a “spectator” that slightly perturbs the nuclear potential perceived by the transitioning electron [127, 128]. Table 3.1 shows the electronic configurations of the He-like lines and the satellites of main interest in this work.

Line name	Transition	Line name	Transition
w	$1s2p (^1P_1) - 1s^2 (^1S_0)$	z	$1s2p (^3S_1) - 1s^2 (^1S_0)$
x	$1s2p (^3P_2) - 1s^2 (^1S_0)$	y	$1s2p (^3P_1) - 1s^2 (^1S_0)$
a	$1s^2 2p (^2P_{3/2}) - 1s2p^2 (^2P_{3/2})$	m	$1s^2 2p (^2P_{3/2}) - 1s2p^2 (^2S_{1/2})$
q	$1s^2 2p (^2S_{1/2}) - 1s2p^2 (^2P_{3/2})$	r	$1s^2 2p (^2S_{1/2}) - 1s2p^2 (^2P_{1/2})$
s	$1s^2 2p (^2S_{1/2}) - 1s2p^2 (^2P_{3/2})$	t	$1s^2 2p (^2S_{1/2}) - 1s2p^2 (^2P_{1/2})$
j	$1s^2 2p (^2P_{3/2}) - 1s2p^2 (^2D_{5/2})$	k	$1s^2 2p (^2P_{1/2}) - 1s2p^2 (^2D_{3/2})$

Table 3.1: Electronic configurations for transitions in He-like spectra, including dielectronic satellites $1s^2 2l - 1s2l2l'$, identified using Gabriel notation [129].

Table 3.2 shows the wavelengths and energies of He-like Ar and Ca lines, as measured using the Alcator C-Mod XICS diagnostic (see section 4.4).

	λ_w [Å] (E_w [keV])	λ_x [Å] (E_x [keV])	λ_y [Å] (E_y [keV])	λ_z [Å] (E_z [keV])
Ar XVII	3.94912 (3.1395)	3.96581 (3.1263)	3.96934 (3.1235)	3.99417 (3.1041)
Ca XIX	3.1773 (3.9022)	3.1892 (3.8876)	3.1928 (3.8832)	3.2111 (3.8611)

Table 3.2: Wavelengths (Å) and energies (in parentheses, keV) of the resonance (w), intercombination (x, y), and forbidden (z) lines for Ar and Ca as measured on Alcator C-Mod.

Fig. 3-3 shows the energy level structure (Grotrian diagram) of Li-like and He-like ion stages giving rise to the lines of Table 3.1. We also indicate $n = 3$ and $n \geq 4$ satellite transitions since these are also present within the spectral range of interest. Color and line styles represent different upper level population processes: radiative recombination (RR), dielectronic recombination (DR), electron impact excitation (EIE), inner shell excitation (ISE), and inner shell ionization (ISI).

The detailed modeling of He-like Ca spectra in this thesis was made possible by the

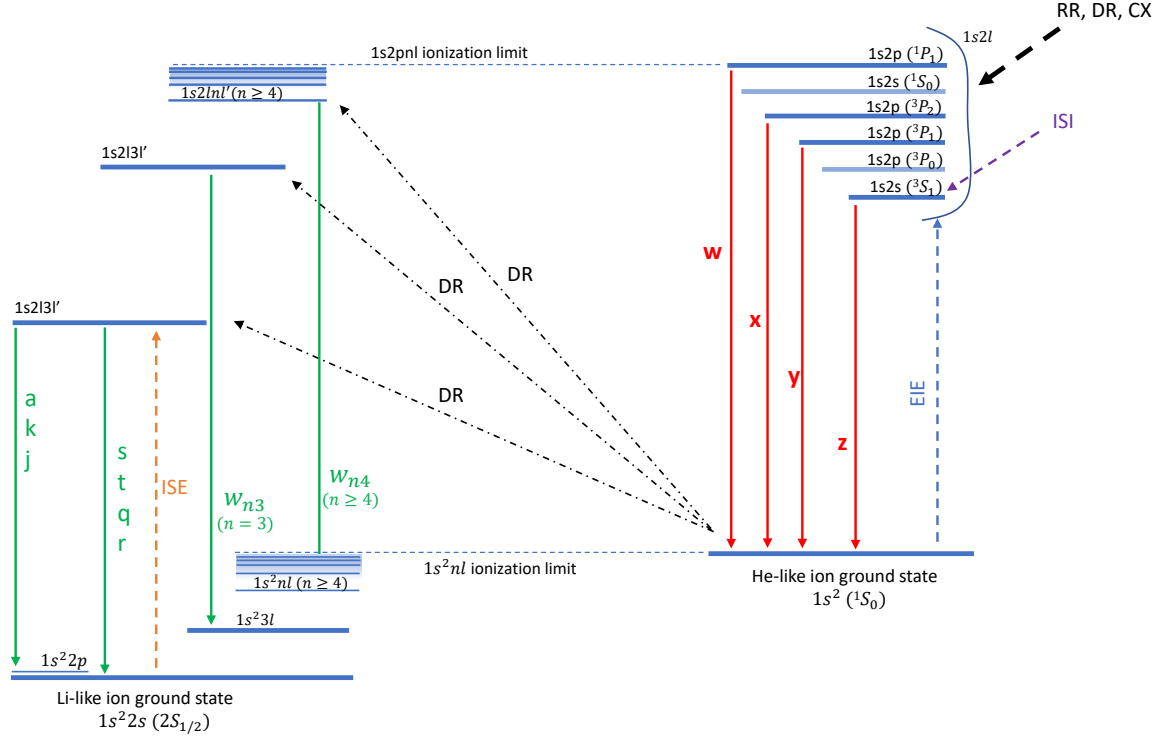



Figure 3-3: Grotrian diagram showing the energy levels (not to scale) and transitions leading to K_α spectra, deriving from He-like and Li-like ions. Transitions between the $1s2l$ levels and two-photon decays are not shown for simplicity (see Fig.1 of Ref. [130]). Colors and line styles differentiate different atomic processes: radiative recombination (RR), dielectronic recombination (DR), charge exchange (CX), electron impact excitation (EIE), inner shell excitation (ISE) and inner shell ionization (ISI). Adapted from Ref. [117].


inclusion of atomic rates for all the relevant processes within the atomDB database [131]. These rates were mostly derived from the existent literature [127, 132, 133, 134, 135, 136, 137] and complemented by A. R. Foster [138] with calculations using the Flexible Atomic Code (FAC) [116].

Use of this compilation of atomic data in atomDB allows the creation of synthetic He-like Ca spectra such as those shown in Fig. 3-4. These spectra were each computed at a combination of electron density and temperature ($n_e = 10^{20} \text{ m}^{-3}$ and $T_e = T_i = 1 \text{ keV}$ for Fig. 3-4a; $n_e = 10^{20} \text{ m}^{-3}$ and $T_e = T_i = 3.5 \text{ keV}$ for Fig. 3-4b) assuming pure Doppler broadening, using Eqs. 3.22, 3.23 and 3.24, and charge state fractional abundances in conditions of ionization equilibrium.

From these figures, collisional excitation (blue) is seen to account for most of the upper level population of the w, x and y lines, while the z line may be equally dependent on excitation and radiative recombination (green) at high temperatures. The j satellite is



(a)



(b)

Figure 3-4: Local K_α Ca spectrum at $n_e = 10^{20} \text{ m}^{-3}$ and (a) $T_e = T_i = 1 \text{ keV}$, and (b) $T_e = T_i = 3.5 \text{ keV}$.

nearly degenerate with the z line, meaning that the region between 3.208 and 3.213 \AA is dependent on Li-like, He-like and H-like Ca stages.

Fig. 3-5 shows the same spectrum as in Fig. 3-4a (at $T_e = T_i = 1 \text{ keV}$), but now on a log scale to allow better visualization of the hundreds of satellite lines that contribute to the spectrum, particularly in the colder part of the confined region. Most of these satellites have amplitude smaller than 1% of the w line emissivity, but some of them significantly




Figure 3-5: Local K_α Ca spectrum at the same conditions as in Fig. 3-4a, now on a log scale, permitting visualization of many satellite lines within the spectral range.

contribute to the main features of these spectra. Comparing the spectra in Fig. 3-4, radiative recombination appears to be only significant at core-like temperatures; this is a result of the low fraction of H-like Ca at the lower temperatures. Note that this observation does not need to hold in real experimental plasmas where particle transport can bring H-like Ca out of ionization equilibrium and into the edge region, thus enhancing the radiative recombination components of the w, x, y and z lines. It is also worth remarking that the spectra shown in Figs. 3-4a and 3-4b have only a weak dependence on electron density, except for the x and y lines, whose ratio to z is known to constrain possible n_e values [130]. Investigations on the TEXTOR tokamak [117] have also shown that CX contributions to K_α spectra are not negligible whenever background D densities are significant, but this is not the case for our data since our lines of sight mostly focus signal integration in the core.

He-like ion spectra have been the subject of extensive work, both in astrophysics [126, 139, 140, 141, 142, 130] and in tokamaks [143, 144, 145, 146, 147, 148, 149]. Detailed knowledge of the population dynamics of the upper levels of these transitions has enabled measurements of electron temperature from line brightness ratios on multiple fusion devices [150, 147, 145], as well as constraints on edge background neutral densities through their impact via charge exchange [151, 152, 146]. In most cases, tokamak experimental investigations have focused on Ar, due to the simplicity of working with inert gases and

the large Ar He-like state fractional abundance in typical core plasma conditions. In Ref. [146], He-like Ar spectra were also used to investigate impurity transport on the TEXTOR tokamak, linearly scaling radial diffusion profiles from a numerical transport code until matching measured spectra.

Previous work on impurity transport on Alcator C-Mod focused on time histories of the brightness (line integrated emissivity) of the w line, since this usually has the most photon counts following Ca LBO injections [153, 102, 154]. This approach was extended during this thesis to also isolate the z line brightness and constrain inferences with the w/z brightness ratio [155]. In Ref. [156], we also considered the impact of including the x and y intercombination lines and the j and k satellite lines. The latter approach was later found to be problematic in most experimental settings due to the presence of m, s and t satellites in the same region where x and y exist. However, this issue does not seem to be widely appreciated in the solar astrophysics community, where the ratio of intercombination lines to the forbidden line is commonly used as an electron density diagnostic. These shortcomings motivated the development of a much more comprehensive forward model of XICS spectra, involving line integration of local spectra such as those in Figs. 3-4a and 3-4b to model the entire propagation of LBO injections into a variety of regimes. Improvements and limitations of these modeling techniques will be described in detail in Chapter 7.

4

Particle Transport Diagnostics

In this chapter we briefly review the principles of plasma diagnostics that are central to this thesis. Key instruments on both Alcator C-Mod and DIII-D are discussed, with a focus on spectroscopy that has been used in this thesis for experimental impurity transport model validation.

Introduction

This thesis is focused on the interpretation of diagnostic measurements for the purpose of investigating particle transport. An overview of the main diagnostics used in this research on Alcator C-Mod and DIII-D is therefore in order. We omit detailed descriptions of systems that were not the focus of this work. Among these, it is worth mentioning the case of bolometry and Visible Bremsstrahlung (VB) measurements. On both C-Mod and DIII-D, bolometers allow estimation of spatio-temporal patterns of total radiation, which are intrinsically connected to particle transport. However, their interpretation was deemed unfeasible for the inferences of particle transport described in Chapters 7 and 8 due to the signal dependence on edge dynamics and the integration of emission from all ion species. Its use was therefore limited, on both devices, to the assessment of power balance, necessary for both integrated plasma modeling and for the 2-point model described in Section 2.6. Similarly, VB estimates of the effective ion charge, $Z_{\text{eff}} = \sum_i (n_i Z_i^2) / \sum_i (n_i Z_i)$ (see Section 3.4), offer useful quantification of plasma purity and a macroscopic metric of impurity retention, but they do not constrain particle transport effectively as a result of the

Z_{eff} sum over all plasma species and all charge states.

Fig. 4-1 shows the spectral ranges of interest for the diagnostics that are most important to this work. Local measurements of electron temperature and density were obtained on both C-Mod and DIII-D via Thomson Scattering (TS) and Electron Cyclotron Emission (ECE) systems, described in Section 4.1, which work in the infrared and microwave regions. The key *actuator* employed in this work, the Laser Blow-Off (LBO) system, uses a Nd-YAG laser at 1064 nm wavelength, like the TS diagnostics. In Section 4.2, we describe principles and key features of ultraviolet spectroscopy on both devices. Section 4.3 describes the Charge Exchange Recombination (CER) diagnostic on DIII-D, which provides localized measurements of ion emissivity, temperature and rotation from the visible spectral range. This system also allows estimation of impurity densities over space and time in neutral beam injection (NBI) heated discharges and has been central to the DIII-D studies presented in Chapter 8. Section 4.4 describes measurements of impurity brightness, temperature and rotation via X-ray Imaging Crystal Spectroscopy (XICS) used on C-Mod, at the high-frequency end of the spectrum in Fig. 4-1. We conclude this chapter in Section 4.5 with a description of Soft X-Ray (SXR) arrays, which filter radiation in such a way as to focus signal integration from the core region of both devices.

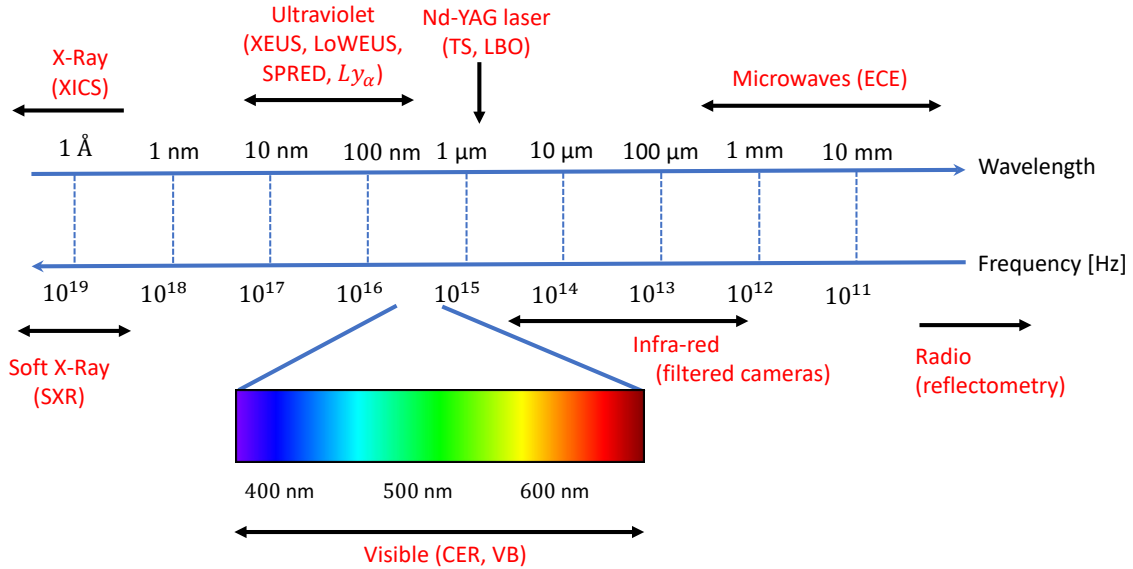


Figure 4-1: Electromagnetic spectrum, illustrating regions of interest for various tokamak diagnostics discussed in this chapter.

4.1 Electron Density and Temperature Diagnostics

The estimation of electron density (n_e) and temperature (T_e) radial profiles is central to much of this thesis work. On both C-Mod and DIII-D, we rely on Thomson Scattering and Electron Cyclotron Emission diagnostics to estimate local values of n_e and T_e , which are assumed to be constant over flux surfaces.

The principle of TS diagnostics is the incoherent scattering of electromagnetic radiation on electrons, usually relying on Nd-YAG ($\lambda = 1064$ nm) pulsed lasers. Scattered radiation is collected at a number of optical fibers whose views intersect the laser path. Gaussian fits to the energy-resolved spectra provide n_e and T_e estimates at each sampled location. The same type of class-IV laser is also used by the LBO systems on C-Mod and DIII-D [153]. Both the C-Mod [157, 158, 159] and DIII-D [160, 161] TS diagnostics have been extended over the years to reach mm-resolution in pedestal measurements, also providing divertor coverage in DIII-D [160]. Although these systems estimate plasma parameters almost instantaneously at each laser pulse, their time resolution is limited by a repetition rate of $\sim 10 - 20$ ms, comparable to typical transport time scales. To examine faster events (e.g. MHD sawtooth instabilities), ECE diagnostics complement TS with resolutions down to ~ 10 μ s. On both C-Mod [158] and DIII-D [162], ECE systems use second-harmonic extraordinary (X) mode emission, which prevents cut-offs under most operating conditions and enables accessibility at typical densities and magnetic fields of each device.

Over the course of this thesis, we made use of different radial kinetic profile fitting methods. For C-Mod transport modeling we adopted the Gaussian Process Regression (GPR) techniques developed by Chilenski [154], which minimize user bias by inferring hyper-parameters corresponding to the fits of highest probability[154]. We quantify GPR uncertainties by propagating diagnostic uncertainties and data scattering in time via the Law of Total Variance and Markov Chain Monte Carlo (MCMC) sampling [163], using a Gibbs kernel with tanh length scale as suggested in Ref. [154]. While GPR offers a principled approach to profile fitting, the time-independent nature of the GPR tools available to us made them of limited use in the context of particle transport simulations that attempt to resolve time dependent kinetic backgrounds. The use of parametric Radial Basis Function (RBF) routines was found to offer the most effective means to accomplish spatio-temporal fits of n_e and T_e , as in Ref. [88]. Additional routines were developed for the purpose of

pedestal neutral studies, where accurately matching sharp gradients has previously been shown to be easier and more reliable using a “modified-tanh” parametric formulation (see Chapter 6).

4.2 Ultra-Violet Spectroscopy

Diagnosing spectral lines in the ultra-violet range offers opportunities both for tokamak operations, *e.g.* to assess impurity content and possible vacuum leaks, and for physics research. In this work, spectral lines in the Vacuum Ultra-Violet (VUV) range have been analyzed for both C-Mod and DIII-D, developing an analysis framework that attempts to identify, fit and model their emission in a unified, device-agnostic framework.

On C-Mod, VUV emission could be measured via both the X-ray and Extreme Ultraviolet Spectrometer (XEUS) and the Long-Wavelength Extreme Ultraviolet Spectrometer (LoWEUS). Both were flat-field grating spectrometers developed at Lawrence Livermore National Laboratory and deployed on multiple devices [164, 165]. The C-Mod XEUS used a 2400 lines/mm grating and covered the spectral range $10 < \lambda < 70\text{\AA}$, while LoWEUS had 1200 lines/mm and viewed the range $100 < \lambda < 300\text{\AA}$. Fig. 4-2 shows the two views of the spectrometers (blue for XEUS, green for LoWEUS), going through the midplane at slight angles. Red lines here show the views of the XICS chords, described in Section 4.4. Contours in the background show the EFIT [166] magnetic equilibrium reconstruction, with contour color peaking at the poloidal field coil locations outside the vacuum vessel.

On DIII-D, a dual-range SPectrometer Recording Extended Domain (SPRED) [167, 168] is routinely used to monitor low- and medium- Z impurities with a single line of sight going tangentially through the midplane. Two optical systems are arranged back to back with individual entrance slits, gratings and detectors, covering the EUV spectral ranges 10–29 nm and 25 – 110 nm. Each detector has 1024 pixels and typical time resolution of 2 ms, like the C-Mod systems. A similar spectrometer is available for divertor measurements on DIII-D; for clarity, we shall adopt the common nomenclatures of “core-SPRED” and “div-SPRED” to distinguish them.

VUV measurements of low-/medium- Z ions typically provide information on the outer shells of the plasma, where fractional abundances of partially-ionized charge states are most significant. This makes VUV data particularly interesting to complement spectroscopy




Figure 4-2: Lines of sight of the Alcator C-Mod XEUS (blue) and LoWEUS (green) spectrometers, as well as for the spatially-discretized 32 chords of the XICS crystal used for Ca measurements (Section 4.4).

focused on core impurity transport.

VUV Line Identification

Over the course of this thesis, a package for atomic line identification, named `atomID`, was developed for use in control room settings as well as detailed analysis. `atomID` provides a fast and automated comparison of detected atomic lines with those available in the National Institute of Standards and Technology (NIST) Atomic Spectra Database (ASD), widely regarded as one of the most complete collections of its type, as well as other databases, including ADAS [114], CHIANTI [169] and curated selections of commonly observed lines on both DIII-D and C-Mod.

Fig. 4-3 shows an example VUV spectrum where `atomID` was used to identify lines from Ca, F, Ar and Mo only. The top labels show the attempted identification, with a percentage giving a measure of Gaussian uncertainty based on the expected accuracy of the




Figure 4-3: VUV spectrum from the C-Mod LoWEUS diagnostic for a discharge where Ar was puffed and CaF_2 was injected with the laser blow-off system approximately 2 ms before the displayed time slice. A selection of atomID identifications for Ca, Ar and Mo from 3 databases is shown.

diagnostic wavelength calibration and uncertainties associated with each database. Three databases were used here. The first one, labelled “DIII-D SPRED database”, is a collection of approximately 100 lines that are often observed with the core-SPRED system on DIII-D. The “NIST” database comprises all the lines in the NIST ASD system, here sub-selected to only include resonance lines (i.e. lines with the ground state as the lower stage of the transition). The third database is a transcription of the Ca and Ar lines recently identified by Träbert *et al.* [170]. The gray and black superposed spectra in Fig. 4-3 represent two time slices recorded before and after an LBO injection of CaF_2 , respectively. The difference between these two spectra is helpful to identify spectral features that directly or indirectly result from the CaF_2 injection. Ar was puffed at the beginning of this discharge and is continuously recycled; indeed, the Ar XV line at $\lambda \approx 22.1$ nm is the same for the black and gray spectra. Mo is an intrinsic impurity in Alcator C-Mod, being a substantial component of the device plasma facing components. The LBO injection appears to significantly modify the brightness of a well-known Mo XXXII line at $\lambda \approx 12.8$ nm, but closer inspection and differences between time slices following the LBO injection clearly identify an F VI line at 12.69 nm present in the NIST ASD. Many identifications were omitted from this figure for clarity of visualization, particularly in the region $\lambda \approx 15.5 - 17.0$ nm where many Ca lines overlap. For medium- and high-Z ions, it is often the case that individual lines cannot be resolved and modeling of their emission requires consideration of all the observed

components, as will be discussed in Chapters 7 and 8. The `atomID` tools have been developed to enable line identification via a set of machine learning techniques inspired by the field of computer vision; future work will attempt to extend this promising approach.

Multi-Line Constraints

Neither the C-Mod XEUS nor the DIII-D core-SPRED diagnostics are absolutely calibrated, but the relative amplitude of lines within the spectrum can be used to constrain the density of all the species that produce identifiable emission on the detector if an external reference density is available. For example, the fully-stripped carbon density measured on DIII-D by Charge Exchange Recombination Spectroscopy (Section 4.3) can be compared, after applying the appropriate PECs, to the brightness of core-SPRED CX lines that also derive from carbon, since the core-SPRED line of sight also crosses one of the neutral beams. This allows the estimation of a single scalar factor that can then be applied to all other lines of the spectrum in order to estimate the density of each emitting species. One critical element in this approach is the availability of accurate PECs. Approximate estimates may be sufficient for “control room assessments” of impurity densities during tokamak operation, but it is challenging to obtain sufficiently accurate PECs to enable reliable use of the relative calibration across a detector. Efforts to make this a practical option for the inference of particle transport coefficients will be discussed in Chapters 7 and 8.

4.3 Charge Exchange Recombination Spectroscopy

In tokamaks where Neutral Beam Injection (NBI) heating and current drive are routinely applied, diagnostic techniques based on charge transfer between beam neutrals and plasma ions have long been used to estimate kinetic properties [118]. This is the case of the Charge Exchange Recombination (CER) system at DIII-D [171], routinely used to estimate carbon density, temperature and rotation by measuring the local emissivity of H-like C $n = 8 \rightarrow 7$ transitions. The upper level of this transition is known to be mostly populated via charge exchange (see Chapter 3) and is therefore related to the density of beam neutrals at the location of interest, at the crossing-point between the diagnostic lines of sight and the neutral beam path. In practice, some loss of localization is expected due to the finite width of the beam and lines of sight [172] as well as beam halos [173] and beam plumes (particularly

important for He [174]). Assuming the plasma to be optically thin, the total emissivity and knowledge of atomic cross sections allow one to compute the density of the emitting species. From the line shift with respect to the rest wavelength of the transition, the velocity of the emitting species can be inferred via $u = c(\lambda_c - \lambda_0)/\lambda_0$. Moreover, assuming Doppler broadening to be dominant, one can obtain the temperature of the species by fitting simple Gaussians and using the relation $T = (s^2/\lambda_0^2) mc^2$ [175, 115]. Here, λ_c is the center wavelength of the Gaussian line shape, λ_0 is the rest wavelength of the transition, c is the speed of light, s is the Gaussian standard deviation, and m is the mass of the emitting ion. CER measurements are normally made in the visible spectral range to allow use of simpler detector technology and extension to a large number of diagnostic views (48 tangential, 32 vertical on DIII-D). This typically limits CER measurements to low- and medium-Z ions ($Z \leq 20$). CER measurements on DIII-D are routinely made for the intrinsic C impurity. Over the course of this thesis research, measurements of Li, F, Al and Ca were also made following LBO injections. Fig. 4-4 shows the CER measurement locations (green squares and red diamonds) in a diverted negative triangularity discharge that will be examined later, in Chapter 8. The spatial localization of CER chords can be varied between discharges and can, in principle, extend across the high- and low-field sides, also crossing the separatrix. For the analysis of LBO injections, signals are typically integrated over approximately 5 ms.

Recently, CER measurements on DIII-D have been extended also to the main ion (D) species [176, 177, 178], particularly in the pedestal and scrape-off layer (SOL), finding striking differences from carbon ion measurements [179]. In this thesis, we did not make extensive use of these data, although this would be beneficial for detailed analysis of pedestal and SOL particle transport in future work.

Charge exchange recombination spectroscopy has also been pursued on C-Mod, both via a diagnostic neutral beam [158] and via D gas puffs [106]. These measurements offered insights on the poloidal distribution of intrinsic boron impurities, but were not designed to support LBO operation. Their use has therefore been relatively limited throughout this thesis.




Figure 4-4: Magnetic equilibrium reconstruction from a kinetic EFIT [166] of DIII-D negative triangularity discharge 180526 showing the location of localized measurements by the CER vertical (green squares) and tangential (red diamonds) chords, and Thomson Scattering (cyan pluses) diagnostics. Note that CER measurement localization can vary between discharges. Blue X's indicate the lower (primary, active) and upper (secondary) x-points. Flux surface contours are equally spaced in the square root of normalized poloidal flux.

4.4 X-Ray Crystal Spectroscopy

High-resolution x-ray spectroscopy can reveal a wealth of information about ion species in the plasma through detailed measurements of spectral features. In this thesis, the X-ray Imaging Crystal Spectroscopy (XICS) system on Alcator C-Mod, also labeled *High Resolution X-ray spectrometer with Spatial Resolution* (HIREXSR), was used to observe spectra of H- and He-like argon ($Z_{Ar} = 18$) and He-like calcium ($Z_{Ca} = 20$). Ar was commonly puffed into C-Mod for the purpose of estimating ion temperatures and rotation via XICS. Ca, on the other hand, was only introduced via LBO injections (with CaF_2 slides) for the purpose of studying impurity transport by modeling time histories of Ca line brightnesses.

XICS diagnostics use detectors capable of capturing x-ray photons over 2-dimensional pixel arrays, permitting the measurement of spatially- and spectrally-resolved spectra. Spectral resolution is achieved via Bragg diffraction at a spherical quartz crystal, described by the standard formula $n\lambda = 2d\sin\theta$, with n being the diffraction order, λ a given wavelength, d the spacing between crystal atomic layers, and θ the diffraction angle.

The spherical bending results in radiation being focused at a certain radial distance from the crystal (the Rowland circle), with one detector axis corresponding to different wavelengths and the other axis mapping to different spatial locations in the plasma (Johann configuration). The spatial localization corresponds to integrals along the lines of sight projected from arbitrarily discretized regions of the detector, mapping these to the crystal via the diagnostic optics [175]. The XICS system on C-Mod used Dectris PILATUS 100K detectors, each containing 487×187 pixels acting as single photon counters; the diagnostic could give up to 5 ms effective resolution. Four detectors were available for the experiments analyzed in this thesis, three of them used to image Ar emission in the range $3.94 < \lambda < 4.00 \text{ \AA} (\text{Ar}^{16+})$ and one used to view Ca emission in the range $3.17 < \lambda < 3.225 \text{ \AA} (\text{Ca}^{18+})$. To facilitate data analysis, the detector projection onto the plasma is typically discretized into a number of views, typically 32 for the Ca measurements. Fig. 4-2 shows these views (red), on top of the magnetic equilibrium. While much of the emission along each line of sight is collected from regions near the chords' tangency radii to the flux surfaces, some signal is always collected from the entire view. Line integration of local emissivity for each chord can be computed via a set of quadrature weights that depends on the specific magnetic geometry of each experiment.

The line shape measured by XICS is dominated by Doppler broadening, but due to line-of-sight integration the spectra are not exactly Gaussians. Instrumental broadening also results from imperfect Bragg reflection at the crystal, possible displacement of the detector from the Rowland circle and the finite-size mirror effects known as Johann error; these effects have been quantified on C-Mod to give an effective broadening corresponding to 100-300 eV, which must therefore be subtracted from estimated ion temperatures. We note also that natural line broadening due to the finite lifetime of atomic excited states is negligible for the spectra of interest here [175, 115].

Prior to this thesis, the THACO framework [180] was used for spectral analysis and Doppler tomography of all XICS signals. In Section 7.2.3 (and, more thoroughly, in Appendix B), a modern alternative to the THACO spectral fitting methods will be presented.

4.5 Soft X-Ray Arrays

Spectrally-resolved measurements such as those of the VUV, CER and XICS diagnostics are often complemented in tokamaks by diagnostics that integrate over photons at a wide range of energies. As discussed at the beginning of this chapter, this is the case for bolometry systems, which aim at estimating the total radiated power from the plasma. Photodiode arrays with a Si substrate are also often employed to obtain high time resolution (μs scale) radiation measurements integrated over the x-ray range and eliminating visible radiation. This can be achieved by placing a Be foil of thickness ranging between 10-150 μm along the axis connecting the pinhole and photodiode pixels. The measured radiation is then mostly from the soft x-ray (SXR) range, thus motivating the nomenclature of these diagnostics.

On Alcator C-Mod, SXR arrays were set up since the beginning of the device operation using 50 μm Be filters, giving 50% transmission at 2 keV energy [158]. This system was later extended with edge-focused chords viewing through a thin 10 μm filter, with 50% cutoff near 500 eV, to study pedestal radiation [181], but data from this array were not available for the experiments discussed in this thesis. All C-Mod arrays were contained in ultrahigh vacuum and the Be filters were bent to make sure that each line of sight through the plasma would go through the same foil thickness [158].




Figure 4-5: Tomographic reconstruction of SXR radiation from the DIII-D diverted negative triangularity discharge #180526 discussed in Chapter 8. White dashed lines show selected magnetic flux surfaces, while the thick continuous line shows the LCFS as inferred from EFIT.

An analogous SXR system is also available on DIII-D, with 2 poloidal and 3 toroidal arrays. The poloidal arrays each have two linear 16-element Si (AXUV) photodiode arrays, for a total of 64 chords, while the toroidal ones each provide 12 chords, for a total of 36 chords. Unlike the C-Mod system, the DIII-D SXR arrays have Be filters of variable thickness (commonly 12 and 125 μm) placed on an adjustable filter wheel, but none of these filters is spherically-bent; as a result, different lines of sight have slightly different filtering effects. This is a cause of significant uncertainty when attempting to predict the signals that should be measured by photodiodes for some given impurity population in the plasma. Fig. 4-5 shows a tomographic reconstruction for a DIII-D negative triangularity discharge, obtained using the PYTOMO software tools [182]. Tomography allows for the estimation of localized emission via regularized inversion techniques, but was not used for the inferences of impurity transport described in Chapter 8. Instead, we made use of line integrated signals without the additional processing step of localization.

In the next chapter, we introduce the key principles of Bayesian inference and describe their application for particle transport studies on both Alcator C-Mod and DIII-D.

5

Bayesian Inference of Particle Transport

In this chapter we present the approach developed in this thesis for the inference of experimental particle transport coefficients on both Alcator C-Mod and DIII-D. After a short review of previous efforts, we review key concepts of Bayesian statistics and nested sampling. We then present our forward model based on the Aurora package and a number of techniques that have been developed to better model both likelihood and priors.

Introduction

The inference of D and v profiles relies on an iterative process in which predictions of charge state densities and emissivity profiles are computed for given sets of parameter 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.

The need for computational statistics to infer transport coefficients derives from the lack

of direct and unequivocal constraints from experimental data. This is generally described as an “ill-posed problem”. It is interesting to contrast the case of particle transport inferences with the problem of plasma tomography, where local quantities are recovered from line-integrated measurements via fast regularized optimizations [182]. Tomography’s main purpose is to localize the quantity that is directly measured, *i.e.* to obtain local emissivities from line-integrated brightness. In order to succeed, tomography requires sufficiently constraining data; in this regard, the problem is similar to particle transport inferences. However, in the context of the latter, one typically only has relatively few measurements and, most importantly, these only offer a weak constraint on the quantities of interest. While diagnostics like CER (Section 4.3) measure charge state densities directly (still with a number of modeling choices and assumptions related to neutral beam penetration and atomic rates), CER time and spatial resolutions are usually insufficient to uniquely quantify the full time history of charge states within the plasma. Other measurements are typically even more weakly constraining, *e.g.* XICS (Section 4.4) measures line-integrated radiation with complex contributions from multiple charge states. As we will describe in this chapter, it is therefore important to combine information from multiple measurements and techniques in order to constrain transport inferences as much as possible. It seems however unrealistic to expect that enough measurements could be obtained to uniquely and rigorously offer comparisons to theoretical transport models on any current device. One must therefore “let the model fluctuate around the data”, in the words of J. Byers [183], attempting to reproduce experimental observations while varying free parameters, numerical models, and assumptions.

On both Alcator C-Mod [26] and DIII-D [87], the inference of experimental impurity transport coefficients has been previously attempted using a χ^2 minimization via the Levenberg-Marquardt algorithm [102]. This approach suffers from the risk of finding *local* χ^2 minima rather than *global* ones, depending on initial guesses of model parameters. A step forward in this respect was offered by Villegas [92], who made use of genetic algorithms that enable wider exploration of the parameter space. More recently, Odstrčil [88] applied Markov Chain Monte Carlo (MCMC) algorithms to improve uncertainty estimates from Levenberg-Marquardt optimization. While these techniques may be appropriate in well-constrained problems, they still generally lack rigorous uncertainty quantification and model selection, making the avoidance of under- or over-fitting to experimental data a difficult and somewhat

arbitrary process. Using synthetic data, Chilenski [184] demonstrated the importance of model selection in impurity transport studies, showing that this requires a much larger number of forward model evaluations than was possible at the time. Unfortunately, synthetic data make it difficult (arguably impossible) to also explore issues of *model inadequacy* that may result from over-simplified experimental analysis, inaccurate atomic rates, discrepancies of detector models and other important factors [185]. For example, the omission of sawtooth modeling or the time-dependence of atomic rates, as in previous C-Mod work [102, 184, 103], may prevent the optimal model complexity to emerge.

This thesis offers advances in both forward and inverse modeling of particle transport. Forward modeling has been particularly improved via the development of the Aurora package [186] (Section 5.3.1). Aurora includes a *generative* (simulation) model of particle transport in fusion plasmas and is the cornerstone of all synthetic diagnostics created in this work, for both Alcator C-Mod and DIII-D. Aurora’s development has enabled a shift to High-Performance Computing (HPC), which is also central to our inverse modeling enhancements via advanced algorithms such as nested sampling (Section 5.2). These have allowed us to explore model selection on real tokamak data for the first time, and to depart from a number of idealizations often adopted in experimental analysis.

In this chapter, we describe methods for both forward and inverse modeling developed in this thesis to infer particle transport coefficients from experimental data. The statistical techniques discussed here generally fall within the realm of Bayesian statistics, where distributions are interpreted as probabilities rather than event frequencies, as in the “frequentist” paradigm. We make no attempt to describe differences between Bayesian and Frequentist statistics; interested readers are referred to the review by Von Toussaint in Ref. [187]. In the following section, we begin by reviewing some of the nomenclature used in the rest of this work.

5.1 Key Concepts of Bayesian Inference

In Bayesian frameworks, we are generally interested in the inference of parameters θ (omitting vector symbols for simplicity) given experimental data \mathcal{D} and a model \mathcal{M}_i . Bayes’ formula reads

$$p(\theta, \gamma | \mathcal{D}, \mathcal{M}_i) = \frac{p(\mathcal{D} | \theta, \gamma, \mathcal{M}_i) p(\theta, \gamma | \mathcal{M}_i)}{\mathcal{Z}(\mathcal{D} | \mathcal{M}_i)}, \quad (5.1)$$

where

$$\mathcal{Z}(\mathcal{D}|\mathcal{M}_i) = \int p(\theta, \gamma|\mathcal{M}_i)p(\mathcal{D}|\theta, \gamma, \mathcal{M}_i)d\theta d\gamma \quad (5.2)$$

is the *evidence* (or *marginal likelihood*) of model \mathcal{M}_i . Here, θ is a vector containing the parameters that we wish to infer and γ is another vector containing nuisance parameters, i.e. parameters that are needed by our model evaluation, but whose values are not of physical interest. Any “choice” of parameters $\{\theta, \gamma\}$ may be referred to as a parameter *sample*. For simplicity, we shall often avoid an explicit separation of θ and γ in our notation. The term $p(\theta, \gamma|\mathcal{M}_i) \equiv \pi(\theta)$ (a probability density for the parameters, given a specific model) is generally referred to as the prior (see Section 5.4), 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), one 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.

In recent years, improved computing capabilities have led to increased popularity of MCMC algorithms, which can undertake a wide exploration of the parameter space. These methods typically suffer from the need of approximate independence from starting conditions, significant parameter tuning and issues of multimodality [188, 189]. MCMC variants such as ensemble sampling MCMC and Parallel Tempering MCMC [189] offer advantages, but still require careful choices to be made for chains to explore the parameter space (“mix”) effectively.

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 [188]. We note that the χ^2 is a metric describing the accuracy of signal matching 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- χ^2* , is a deceptively simple and often unjustifiable metric for model selection [190].

Fig. 5-1 shows a diagrammatic overview of the inference framework developed in this thesis for both Alcator C-Mod and DIII-D analysis. The following sections describe the main elements of this diagram, particularly those that differ from previous efforts to infer transport coefficients in tokamaks.

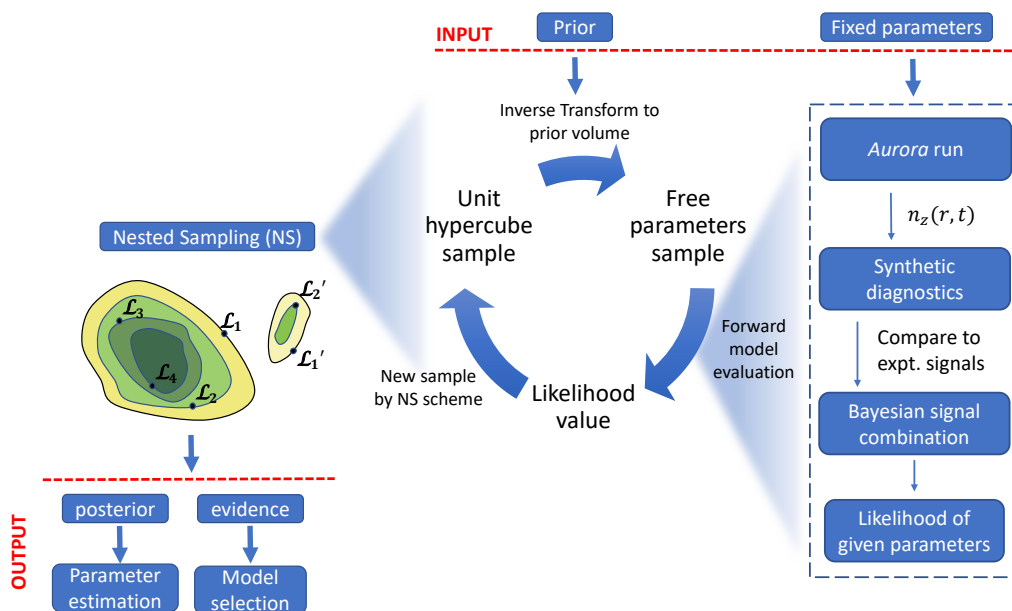


Figure 5-1: Overview of the nested sampling approach taken for the inference of impurity transport for both Alcator C-Mod and DIII-D plasmas. Figure adapted from Ref. [155].

The left hand side of Fig. 5-1 is a representation of how sampling in our inferences of particle transport is “orchestrated” via a *nested sampling* algorithm. Each sample corresponds to a single set of free parameters, which is passed as input to the forward model, shown within the right hand side box. For both C-Mod and DIII-D inferences, sampling is run in parallel on approximately 300 CPUs¹ (depending on the problem complexity) on nodes that are part of the *engaging* cluster at the Massachusetts Green High Performance Computing Center.

¹Intel(R) Xeon(R) CPU E5-2683 v4, 2.1 GHz.

In Section 5.2, we first present nested sampling and its use in this work. In Section 5.3 we describe our likelihood model, including the Aurora package used for forward modeling. Finally, Section 5.4 describes methods and choices related to prior distributions. We note that nested sampling has also been adopted in the Bayesian Spectral Fitting Code, described in section 7.2.3 and Appendix B, hence the descriptions of the next section are relevant more broadly than just in the inference of particle transport.

5.2 Nested Sampling

Impurity transport inferences in this thesis make use of the Nested Sampling (NS) method [191] for model selection and parameter estimation. The primary objective of NS is to evaluate the evidence integral in Eq. 5.2, which can be re-written as

$$Z = \int \mathcal{L} dX \quad (5.3)$$

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

$$X(\lambda) = \int_{\{\theta: \mathcal{L} > \lambda\}} \pi(\theta) d\theta \quad (5.4)$$

where $\pi(\theta)$ is the prior distribution. $X(\lambda)$ 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. 5.3, 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 compute Eq. 5.3 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, typically via the inverse Cumulative Distribution Function (CDF) of the specific prior used for each dimension. As a result of this transformation, unit-hypercube samples are mapped to regions of higher prior probabilities more frequently than to regions with lower prior probabilities.

Parameters are then passed to the forward model, whose comparison to experimental data gives a likelihood value. As will be described in Section 5.3, our forward model is constituted by a single simulation with the Aurora package and application of synthetic

diagnostics. The NS algorithm 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 rejected otherwise. 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. 5.3, 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, the estimate of \mathcal{Z} is further improved by using live points that were discarded at every iteration [192].

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 [193, 192] does so by fitting (potentially overlapping) ellipsoids to the set of live points and sampling from within their union. One attractive feature of `MultiNest` is its speed and effectiveness in exploring parameter spaces with up to 30-50 dimensions (depending on the problem of interest). While other algorithms, such as `PolyChord` [194], perform better at higher dimensionality, `MultiNest` is an excellent tool for our transport inferences extending to a maximum of 30 free parameters. For standard 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, using $n_{live} = 200 + 25 \times D$ [194]. Future work will explore varying n_{live} during an inference using the dynamic nested sampling algorithm [194].

`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 inferences of particle transport, we conservatively set $f = 0.1$ when using vanilla NS and $f = 0.01$ for INS. Increasing n_{live} or decreasing f from our defaults 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 report on separate 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 a data-driven approach allows us to understand ambiguities of our data that may lead to unphysical local minima in posterior distributions.

5.3 Likelihood Model

The development and refinement of a forward model are clearly central to the value of an inference. Comparing the results of a forward model to observed data results in a likelihood evaluation, which may be interpreted as a metric for the “quality” of the match between a certain set of forward model inputs and the data. In the context of Bayesian inference, the likelihood appears in Bayes’ formula, Eq. 5.1, in the same way as the prior distribution, which is described in Section 5.4.

In inferences of particle transport, the forward model involves the simulation of cross-field density evolution following LBO injections and the application of synthetic diagnostics. This requires evolving the density of a chosen ion “I” over time and space (normally a 1.5D radial coordinate defined using normalized magnetic flux-surface volumes) and solving coupled continuity equations for each charge state “Z” in cylindrical geometry:

$$\frac{\partial n_{I,Z}}{\partial t} = -\frac{1}{r} \frac{\partial}{\partial r} (r \Gamma_{I,Z}) + Q_{I,Z}. \quad (5.5)$$

The radial (cross-field) particle flux $\Gamma_{I,Z}$ is separated into diffusive and convective terms with the Fick’s Law ansatz

$$\Gamma_{I,Z} = -D \frac{\partial n_{I,Z}}{\partial r} + v n_{I,Z}, \quad (5.6)$$

where D is the diffusion coefficient and v is the radial convective velocity. Both D and v can be set as functions of the radial coordinate, time, and also charge state of the ion. In this work, we focus on steady-state background plasma conditions and assume time independence of these coefficients. Typically, D and v are also taken to be the same for all

charge states, although linear dependence of v on Z (the charge of each ion stage) is expected from neoclassical theory and turbulent transport is likely to have some weak dependency on Z too. In Chapter 7, the dependence of pedestal transport coefficients on Z is explored, but this is otherwise omitted.

The $Q_{I,Z}$ term in Eq. 5.5 is the same given in Eq. 3.12, and we reproduce it here for convenience:

$$\begin{aligned}
Q_{I,Z} = & -(n_e S_{I,Z} + n_e \alpha_{I,Z} + n_H \alpha_{I,Z}^{cx}) n_{I,Z} \\
& + n_e S_{I,Z-1} n_{I,Z-1} \\
& + (n_e \alpha_{I,Z+1} + n_H \alpha_{I,Z+1}^{cx}) n_{I,Z+1}.
\end{aligned} \tag{5.7}$$

$Q_{I,Z}$ accounts for sources and sinks from atomic processes of ionization and recombination, whose rates are estimated using the Atomic Data and Analysis Structure (ADAS) database [114] at measured values of electron density (n_e) and temperature (T_e) as a function of the radial coordinate. Charge exchange also requires specification of the background atomic density of hydrogen-isotopes, n_H . It is through $Q_{I,Z}$ that continuity equations for different charge states are coupled.

Forward modeling of impurity transport provides time histories of $n_{I,Z}$ radial profiles starting from an initial condition representing the LBO injection. Charge state densities can then be used to compute the emissivity of each transition $\epsilon_{i \rightarrow j}$ by evaluating the appropriate photon emissivity coefficients at local electron densities and temperatures, as discussed in Section 3.3.

Once the emissivity $\epsilon_{i \rightarrow j}$ of each line has been computed, this can be directly compared to local data, e.g. from CER, or it can be line integrated along the lines of sight of measuring instruments. This completes the definition of our synthetic diagnostics. Thus, each forward model evaluation generally consists of a single forward model simulation for some choice of input parameters (relating to plasma geometry, kinetic profiles, D and v profiles, etc.), the calculation of line emissivities of interest over space and time, and finally the appropriate operations of line integration.

As described in Section 5.1, to infer radial profiles of transport coefficients that correspond to observed signals we iterate over this forward model while varying D and v (and any nuisance parameters) until reaching satisfactory agreement. The choice of D

and v parametrization has an important impact. One may easily infer a single value of D to give a flat diffusion profile and a single value of v to give a convection profile that is 0 on axis and grows linearly with the radius; while this choice makes inferences fast and simple, it does not offer a level of detail that may be useful for transport model validation. In order to infer radially-detailed profiles, we make use of splines for both D and v at a variable number of knots; the exact number is determined via model selection in each case by comparing the Bayesian log-evidence to other spline complexities. An additional Gaussian feature is included in the pedestal region of the v profile, with free amplitude, width and location being freely varied within chosen priors. This allows for an effective reduction of model complexity when trying to make inferences with a limited number of free parameters, based on the physical expectation that radially sharp features may be present in the v pedestal.

In the next section, we present Aurora, a software package that was developed as part of this thesis work to offer a modern forward model of particle transport for research on both Alcator C-Mod and DIII-D. The following sections then address a number of other aspects of forward modeling that significantly improve likelihood modeling with respect to previous work on inferences of particle transport.

5.3.1 The Aurora package

Bayesian inferences via nested sampling are significantly more rigorous and complete than more traditional methods involving non-linear optimizers, but they also typically require many more iterations. For this reason, using previous computational packages for particle transport as a forward model was found to be unfeasible. Among these, the STRAHL [195] and MIST [196] codes have a history of being used for Alcator C-Mod and DIII-D analysis. Others, including SANCO [197] and ITC [91], have found applications in numerous works over the past decades, with widespread adoption beyond the research groups of the original developers. None of these packages appeared to be appropriate for the purpose of iterating during Bayesian inferences for the following (non-exhaustive) set of reasons:

- They are generally hard to adapt to high performance computing due to their need to repeatedly access a hard disk for input/output files;
- Limited interaction of users with the source code makes it difficult to assess atomic data and spectroscopic predictions;

- Modification of the code to explore new algorithms or physical effects can be challenging.

Over the course of this thesis, a new package was developed to offer a modern alternative to these tools. The result of this effort has been named Aurora [186] and has evolved into a toolbox for particle transport, neutrals and radiation modeling. Its initial development was based on the STRAHL code [195], against which Aurora has been fully benchmarked.

Aurora is well integrated with Python tools originally developed within the One Modeling Framework for Integrated Tasks [198] (OMFIT) and recently released for independent installation. This allows Aurora to leverage a wide community effort to generalize common analysis routines and access data from multiple devices. For example, reading of files in standardized formats, e.g. from EFIT [166] or GACODE routines, is easily accomplished, permitting non-trivial manipulations of experimental data without the need for detailed expertise. Integration of Aurora within integrated modeling frameworks is made particularly straightforward by its simple installation procedures, multi-language interfaces and device-agnostic implementation. Aurora is an entirely open source project, hosted on *GitHub* and welcoming contributions from users with a range of interests. Documentation is available at <https://aurora-fusion.readthedocs.io>.

Aurora’s forward model couples transport and atomic processes by solving continuity equations for each charge state over time and a 1D radial coordinate. Fig. 5-2 illustrates the main categories of inputs and outputs of the code, whose objective is to provide charge state densities (n_z), emissivities (ϵ_z) and radiated power components (L_z) as a function of radial location, time and charge (Z).

Starting from an initial condition (an impurity neutral source from a LBO injection, a gas puff, a pellet, etc.), Aurora can automatically retrieve from the internet, load and process atomic rates, typically from the Atomic Data and Analysis Structure (ADAS) database [114]. It can work with magnetic geometries from EFIT [166] files and interact with machine databases via the MDS+ data acquisition system [199]. Plasma profiles (n_e , T_e , T_i) can be set arbitrarily or loaded from files of several standardized formats. Background (D) neutral densities, shown in Chapter 6 to be extremely important for the interpretation of particle transport in the edge, can be given as a free input based on spectroscopic analysis (e.g. the D Ly $_{\alpha}$ data discussed in Chapter 6) or taken from Aurora interfaces with external neutral modeling suites such as KN1D [200], FIDASIM [201, 202] or SOLPS-ITER/EIRENE [203,

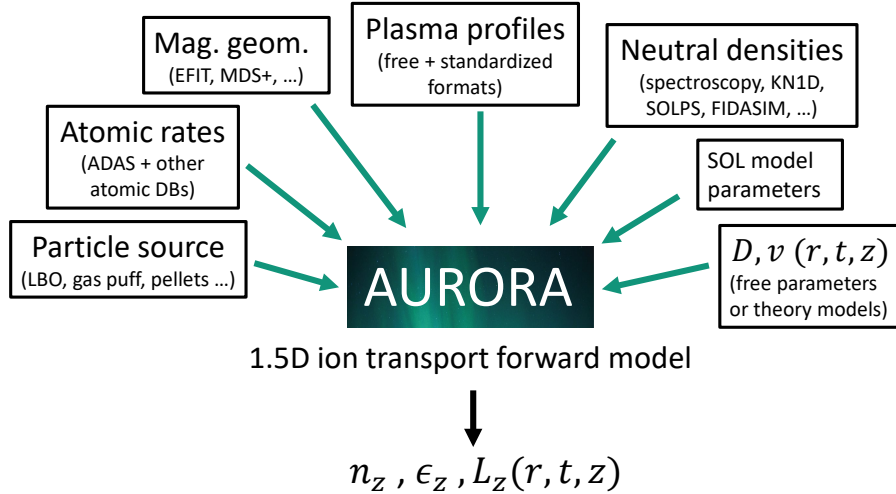


Figure 5-2: Overview of Aurora’s 1.5D ion transport forward model. All inputs can be obtained either from experimental data or modeling, except for atomic rates (normally from ADAS) and transport coefficients (either from a theoretical transport model or considered “free parameters”). Aurora can calculate charge state densities (n_z), emissivities (ϵ_z) and radiated power components (L_z). Figure adapted from Ref. [186].

204, 205]. A small set of parameters is also needed as input to Aurora’s simple Scrape-Off Layer (SOL) 0D model, which allows basic considerations on wall and divertor recycling.

Fig. 5-2 highlights that D and v are *inputs*, rather than outputs, to Aurora’s forward model of ion transport. Since in this work we are interested in inferring D and v profiles from experimental data, one must therefore iterate over these D and v inputs until the output of Aurora matches observations. D and v can be provided as functions of time, space and charge state (i.e. Z value). Typically, it is hard to determine what Z dependence should be assigned to transport coefficients, but in some cases this capability can enable validation efforts, e.g. to test the predictions of neoclassical transport in the pedestal region [155]. It is also possible to provide transport coefficients for each charge state as obtained via a neoclassical and/or turbulent transport code.

Currently Aurora solves coupled continuity equations for all charge states using by default a first-order, vertex-centered, finite-volume scheme recently developed by Linder [206]. This offers particle conservation and better numerical stability with respect to finite-difference schemes in the presence of large convection, as in pedestal regions of tokamaks. The scheme uses adaptive upwinding for the spatial discretization of the advective

terms based on local evaluations of the parameter

$$K_i = \max(0, 1 - 2/|\mu_i|) \cdot \text{sgn}(\mu_i), \quad (5.8)$$

where $\mu_i = |v(r_i)| \Delta r_i / D(r_i)$ is the local Péclet number. In the limit of $\mu_i \rightarrow 0$ ($K_i = 0$, diffusion-dominated case) a purely central scheme is used, while for $\mu_i \rightarrow \infty$ ($K_i = 1$, advection-dominated case) pure upwinding is adopted.

The temporal discretization is performed using the θ -method, equally weighing contributions from densities at the previous time step and the new time step (i.e. $\theta = 1/2$). Ionization and recombination rates are set to act on each charge state using the Lackner method [207], which considers ionization and recombination in two half steps. In the first half step, the ionization term is computed using the density at the new time step (implicitly) and recombination with the density at the previous time step (explicitly); in the second half step, ionization is computed at the previous time step and recombination at the new time step. This method is unconditionally stable and allows one to use the fast Thomas algorithm to solve the resulting tri-diagonal system of equations for the charge states. Readers are referred to Refs. [206] and [207] for details on the numerical scheme. Additionally, as proposed by Linder [206], Aurora can evolve the neutral stage in the same way as ionized stages, even though the specification of transport coefficients for neutrals is not trivial. Alternatively, users can choose not to evolve neutrals, in which case recombination for the first ionized stage is artificially set to zero to prevent particle losses.

The 1st order finite differences scheme described in the STRAHL manual [195] has also been implemented in Aurora to benchmark routines against a more established code. Such benchmarks have proven to be perfectly satisfied when using identical modeling choices to those available in STRAHL.

Aurora also allows the use of *superstages* as a way to reduce computational complexity by effectively “bundling” charge states together. The technique used in Aurora is analogous to one explored in past JET [197, 208, 209] and ITER [210] modeling. The Aurora implementation has been developed with a focus on simplicity and flexibility, avoiding the need for additional ADAS files to be created since all required adaptations are handled internally. To the author’s knowledge, this is significantly simpler and more flexible than methods applied in any other particle transport code. More details on superstaging are

given in Appendix A.

In concluding this section, we summarize and highlight key advantages of the Aurora 1.5D forward model with respect to pre-existing codes:

- In iterative frameworks, Aurora can be more than $10\times$ faster.
- The current forward model algorithm gives better particle conservation at large v/D with respect to finite-difference schemes, even with half as many grid points.
- Interfacing with other modeling frameworks is made easier via dedicated modules and OMFIT tools.
- The *superstaging* method (Appendix A) offers additional speed improvements and a simple test bed for atomic physics in edge transport codes.

Aurora is a fundamental toolbox for most of the work presented in this thesis, including the analysis and modeling of neutral particles in Chapter 6. Its interfaces to ADAS data make it a valuable package for simulations of impurity transport, as well as for experimental plasma spectroscopy. The forward models used for the Bayesian inferences of particle transport presented in Chapters 7 and 8 are fundamentally based on Aurora.

5.3.2 Treatment of the Plasma Background

Impurity transport inferences in previous Alcator C-Mod work [26, 27] “averaged” over sawteeth (Section 2.5), effectively ignoring the observed time-dependent behavior of the plasma background. This caused inferred transport coefficients to “compensate” for the lack of MHD effects in simulations, making their values near the magnetic axis not easily interpretable. However, since validation efforts were focused at midradius, rather than near the axis, this was argued to be of negligible importance. This is indeed the case for the L-mode discharge analyzed in these works and revisited in Chapter 7 of this thesis. However, averaging over sawteeth is not appropriate for high-performance discharges where these MHD processes can have a much larger impact on particle and heat transport. In this thesis, time-dependent kinetic profiles are always used to appropriately capture the modulation of on-axis T_e as a result of MHD activity. Experimental observations on Alcator C-Mod [211] have suggested that sawteeth are unlikely to affect T_i and n_e as much as T_e .²

²Electron density profiles are typically quite flat in C-Mod’s high collisionality profiles, hence making the eventual impact of sawtooth flattening less significant.

Based on the relaxation times of Fig. 3-1, which are 1-10 ms for most ion species at typical core temperatures, one would expect SXR brightness to only reflect T_e crashes on a ms time scale. However, measurements show a very rapid response ($\ll 1$ ms), suggesting that variations of SXR brightness after sawteeth must be mainly caused by the radial redistribution of individual impurity charge states. Therefore, the effect of sawteeth on impurities must be imposed as an additional process to the diffusive-convective transport. For this purpose, a phenomenological model of particle redistribution following sawtooth crashes has been implemented in Aurora, analogously to the one used in STRAHL. Both are based on experimental observations on Alcator C [63, 62], where charge state density profiles were observed to flatten inside of the sawtooth mixing radius. Aurora models this behavior using complementary error functions with

$$n_z(r) \leftarrow \frac{\bar{n}_z}{2} \operatorname{erfc}\left(\frac{r - r_{\text{mix}}}{w}\right) + \frac{n_z(r)}{2} \operatorname{erfc}\left(\frac{r_{\text{mix}} - r}{w}\right), \quad (5.9)$$

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 radius, r_{mix} , and $w \approx 1$ cm is a “crash region width”, which sets the spatial smoothness of the sawtooth crash. Such smoothness permits one to avoid the particularly small time steps that would otherwise be necessary to maintain particle conservation following sharp crashes.

In the presence of fast changes to background kinetic profiles, it is also important to account for the finite signal integration over time of each spectroscopic diagnostic. This is also crucial if the diagnostic time resolution is of the same order as the rise time following a LBO injection, as it is the case for most spectrally-resolved measurements used in this thesis. On Alcator C-Mod, we also allow for a free parameter to account for uncertainties in the triggering time of detector modules.

We conclude this section by remarking the importance of including background D neutral profiles for the evaluation of total recombination rates of impurity ions. This is discussed in detail in Chapter 6, where background neutrals are shown to have a strong impact on charge state balance via the effect of charge exchange. It is extremely important to account for this in the pedestal of any current fusion device, but also in the core of low density discharges, for example in DIII-D L-mode scenarios. To the author’s knowledge, work in this thesis has been the first to ever include the effect of charge exchange in impurity transport inferences.

5.3.3 Automatic Rescaling of Synthetic Signals

Previous work by Chilenski suggested including rescaling factors for synthetic signals as free parameters in order to account for the lack (or imperfection) of the absolute calibration of real experimental data. In this scheme, one free parameter would be needed for each diagnostic signal whose amplitude should be independently rescaled, i.e. not be considered relative to other signals. Here, we show a simpler solution to this normalization problem by applying an analytic scale factor α to each set of synthetic signals that should be rescaled with respect to each other. Our objective is to minimize the χ^2 metric between the real signal and a synthetic signal, i.e.

$$\min_{\alpha \in \mathbb{R}} \sum_i \left(\frac{y_i - \alpha \hat{y}_i}{\sigma_{y_i}} \right)^2. \quad (5.10)$$

This can be solved analytically in a few steps, first setting the partial first derivative of the summation with respect to α to 0:

$$\begin{aligned} \frac{\partial}{\partial \alpha} \sum_i \left(\frac{y_i - \alpha \hat{y}_i}{\sigma_{y_i}} \right)^2 = 0 &\implies \sum_i \left(\frac{y_i - \alpha \hat{y}_i}{\sigma_{y_i}} \right) \left(-\frac{\hat{y}_i}{\sigma_{y_i}} \right) = 0 \implies \sum_i \frac{-y_i \hat{y}_i + \alpha \hat{y}_i^2}{\sigma_{y_i}^2} = 0 \\ \therefore \alpha &= \frac{\sum_i \left(\frac{y_i \hat{y}_i}{\sigma_{y_i}^2} \right)}{\sum_i \left(\frac{\hat{y}_i^2}{\sigma_{y_i}^2} \right)} \end{aligned} \quad (5.11)$$

In these expressions, the summation over index i covers all the data points that should be rescaled together. Applying Eq. 5.11 to these synthetic data points achieves the objective of Eq. 5.10 with minimal computational cost, as opposed to the use of free parameters for each diagnostic. The two approaches are not strictly equivalent: the least-squares approach effectively reduces a high-dimensional inference to a lower-dimensionality one, plus a fast optimization within a subset of the parameter space. Nonetheless, application of Eq. 5.11 appears justified in our case in order to make inferences more computationally manageable and reduce chances of multimodality.

5.4 Prior Distributions

As discussed in Section 5.2, nested sampling is an inherently Bayesian method. The use of priors, on top of likelihoods, allows one to provide valuable information to the algorithm and avoid unphysical regions. On the other hand, it also demands 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.

Typically, applications of Bayesian inference make use of independent priors for each dimension, so that the total prior is the product of the distributions for each free parameter:

$$\pi(\theta) = \prod_i \pi_i(\theta_i) \quad (5.12)$$

In order to work in a “whitened” space of prior samples, one can use a unit hypercube and inverse cumulative distribution functions (CDFs), which for a general 1-dimensional distribution $f(\theta)$ are defined as

$$F(\theta) = \int_{-\infty}^{\theta} f(\theta') d\theta' \quad (5.13)$$

The inverse of the CDF function $F(\theta)$ is well defined for some of the most common prior distributions (e.g. uniform or Gaussian). This allows uniform random variables $x \sim U(0, 1)$ to be easily mapped to samples from arbitrary distributions, if we know their inverse CDFs. This is commonly referred to as *inverse transform sampling*.

Independent priors (Eq. 5.12) offer simplicity, but cannot encapsulate relations between free parameters. In Sections 5.4.1 and 5.4.2 we describe two cases of “correlated” priors applied in this research.

Whenever possible, we avoid the use of uniform priors over a parameter value. These would imply that we have no indication of what is the most likely region of parameter space where true physical solutions lie - something that is typically not true. For example, we apply a Gaussian prior, rather than a uniform one, over the sawtooth inversion radius, estimated with ~ 1 cm spatial accuracy via ECE or SXR diagnostics on both Alcator C-Mod and DIII-D. 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} m²/s) to much larger values, possibly approaching 100 m²/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* [212, 213].

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 have explored the use of *Gaussian copulae*, through which we set correlations between the sampled 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, thus making it impossible to separate priors as in Eq. 5.12. Gaussian copulae are defined by the expression

$$C_R^{Gauss}(\vec{u}) := \Phi_R(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d)) \quad (5.14)$$

where R is a correlation matrix, Φ_R is the joint cumulative distribution function of a multivariate normal distribution with zero mean vector and covariance matrix R , and Φ^{-1} is the inverse cumulative distribution function of a standard normal. The result of the transformations in Eq. 5.14, $C_R^{Gauss}(\vec{u})$, 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 [214]. In practice, we have found the use of copulae for our problem not to be very impactful. Options to also impose maximum gradients within the radial profiles of transport coefficients were also tested, but they were found to lack the generality needed to investigate impurity transport across devices, since sharp variations in D and v are indeed possible in some circumstances (e.g. when localized electron cyclotron heating is applied).

In early work leading up to publication of Ref. [155], we chose to apply prior constraints over $\{D, v/D\}$ rather than $\{D, v\}$ 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. More recently, the technique described in Section 5.4.2 has offered an elegant

solution to constrain D , v and v/D simultaneously.

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 . Typically, sampling of $v/D \sim \mathcal{N}(0, 3)$ [m^{-1}] is appropriate. By this notation, we mean that v/D is taken to be distributed as a normal distribution of mean 0 and standard deviation 3 m^{-1} . This allows enough freedom to explore positive and negative peaking within a range that spans measurements on other devices [215, 88, 216, 92] and encompassing observed values of $1/L_{ne} \sim 1 \text{ m}^{-1}$ at midradius in all the discharges analyzed in this work. Near-axis, we allow significantly more variation with $v/D \sim \mathcal{N}(0, 10)$, since neoclassical transport in this region could lead to much greater profile peaking (large v/D). In the pedestal, where we have only weak experimental constraints but large inward convection may be expected, 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).³ For C-Mod, we also impose that a “non-negligible” fraction of LBO-injected particles should enter the confined plasma in Aurora simulations. Previous work (e.g. Ref. [144]) reported LBO penetration fractions greater than 5-10%, suggesting that one may conservatively set a minimum penetration fraction of 1-2%. This simple condition typically appears to limit the range of physically reasonable transport coefficients at the LCFS to $D \lesssim 5 \text{ m}^2/\text{s}$ and $v \lesssim -50 \text{ m/s}$. Note that this explicit constraint on particle penetration fractions is not necessary on DIII-D, since the absolute calibration of the CER diagnostic sufficiently constraints particle penetration, preventing sampling of unphysical edge parameters that would result in too many particles leaving the Aurora simulation domain too rapidly.

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.

5.4.1 Free spline knots

Previous research on impurity transport inferences on Alcator C-Mod [27] demonstrated issues with the *identifiability* of free spline knots when these are all allowed to be free within the radial range of simulations. This can be understood as knots not having an

³Truncated normal priors are used when appropriate to prevent sampling of unphysical parameters, e.g. negative radial “widths”.

individual “identity” due to the operation of sorting them based on their radial location after the original sampling. In other words, since two knots can effectively swap places with no practical difference, they are not “identifiable”, leading to complications in learning the posterior landscape.

In this thesis, this problem has been addressed in a number of ways. The simplest solution is to set knots to be free within some restricted bounds. For example, one may initially set knots to be equally separated within the grid, and then allow their location to vary near their initial location, maintaining a minimal distance between them to ensure that extreme gradients cannot be sampled. This gives some flexibility to knot locations, but does not achieve the complete freedom that one may desire.

An alternative, more elegant and flexible method is to “force” identifiability by remapping unit hypercube samples for spline knot locations in such a way as to always retain their numerical order. Following Handley [217], we define a uniform prior in the hyper-triangle defined by $\theta_{min} < \theta_1 < \dots < \theta_n < \theta_{max}$ with

$$\pi(\theta) = \begin{cases} \frac{1}{n!(\theta_{max}-\theta_{min})^n} & \text{for } \theta_{min} < \theta_1 < \dots < \theta_n < \theta_{max} \\ 0 & \text{otherwise.} \end{cases} \quad (5.15)$$

This leads to CDFs of the form

$$F_i(\theta_i | \theta_{i-1}, \dots, \theta_0) = \left(\frac{\theta_i - \theta_{i-1}}{\theta_{max} - \theta_{i-1}} \right)^{n-i+1}. \quad (5.16)$$

Each value of the CDF, $x_i \equiv F_i(\theta_i | \theta_{i-1}, \dots, \theta_0)$, can be taken to correspond to a unit hypercube sample. By inverting this relationship, we can map unit hypercube samples $\{x_i\}$ to parameter samples $\{\theta_i\}$. A simple re-arrangement of Eq. 5.16 gives

$$\theta_i = \theta_{i-1} + (\theta_{max} - \theta_{i-1}) x_i^{1/(n-i+1)} \quad (5.17)$$

Eq. 5.17 offers the mapping between unit hypercube samples and sorted uniformly-distributed spline knots between θ_{min} and θ_{max} as given by Handley *et al* [217]. One can clearly notice the similarity of Eq. 5.17 with the equivalent linear remapping for a simple 1-dimensional uniform prior

$$\theta_i = \theta_{i-1} + (\theta_{max} - \theta_{i-1}) x_i. \quad (5.18)$$

For our application to transport inferences, we are interested in finding a distribution for free knots between the magnetic axis at $\rho = 0$ and the end of our radial grid at $\rho = \rho_{\max}$. However, it is necessary to ensure that knots are not sampled too close to each other, since this can lead to large gradients in transport coefficients and resulting numerical instability. We therefore wish to impose a minimum distance, $\Delta\rho$, between sampled knots. This can be achieved in multiple ways. We have found it convenient and effective to simply limit the value of θ_{\max} to be $\rho_{\max} - \Delta\rho \times (n - 1)$, where n is the number of free knots, rather than setting $\theta_{\max} = \rho_{\max}$. Every knot sampled via Eq. 5.17 is then shifted by $\Delta\rho \times i$, where i is the ordered index of each knot within the sampled set. In other words, we first sample knots within a reduced radial range $\{0, \rho_{\max} - \Delta\rho \times (n - 1)\}$ and then shift each knot by a distance that is sufficient to enforce $\Delta\rho$ between each knot.

5.4.2 Physically-correlated D - v - v/D priors

When sampling free spline values for D and v parameters, one must ensure not only that such parameters are physically reasonable, but their ratio, v/D , is physical too. This is important because v/D is a physically meaningful quantity, indicative of impurity profile peaking in the absence of particle sources and sinks, and therefore we must exclude the possibility of our forward model exploring unreasonable parameter space. For example, if one were to sample D and v values of $0.1 \text{ m}^2/\text{s}$ and 10 m/s , respectively, both individually reasonable parameters, the v/D ratio would be 100 m^{-1} , which is not itself physical. Clearly, sampling of reasonable D and v values may be achieved by tuning hyper-parameters of each individual prior, but constraining the ratio of two fully independent priors is not possible. Correlations between priors must be introduced. The same issue occurs regardless of whether one chooses to sample the $\{D, v\}$ pair, the $\{D, v/D\}$ one, or $\{v/D, v\}$, since the same two degrees of freedom are effectively used in sampling.

In this thesis, we developed a simple, yet effective method to overcome this difficulty. Rather than sampling any of the parameter pairs listed above, we sample two abstract coordinates in a polar plane. We interpret one of these to correspond to the tangent of an angle, $\tan(\tilde{\Theta})$, and the second to a radial coordinate, \tilde{r} . The $\tan(\tilde{\Theta})$ variable is taken to be distributed according to an arbitrary prior distribution for the product $\chi \cdot v/D$, which we denote by $\mathcal{P}_{\chi v/D}$. Here, χ is a parameter that allows for scale separation of D and v priors. We sample a value of \tilde{r} according to a chosen distribution \mathcal{P}_r . The \tilde{r} and $\tilde{\Theta}$ samples

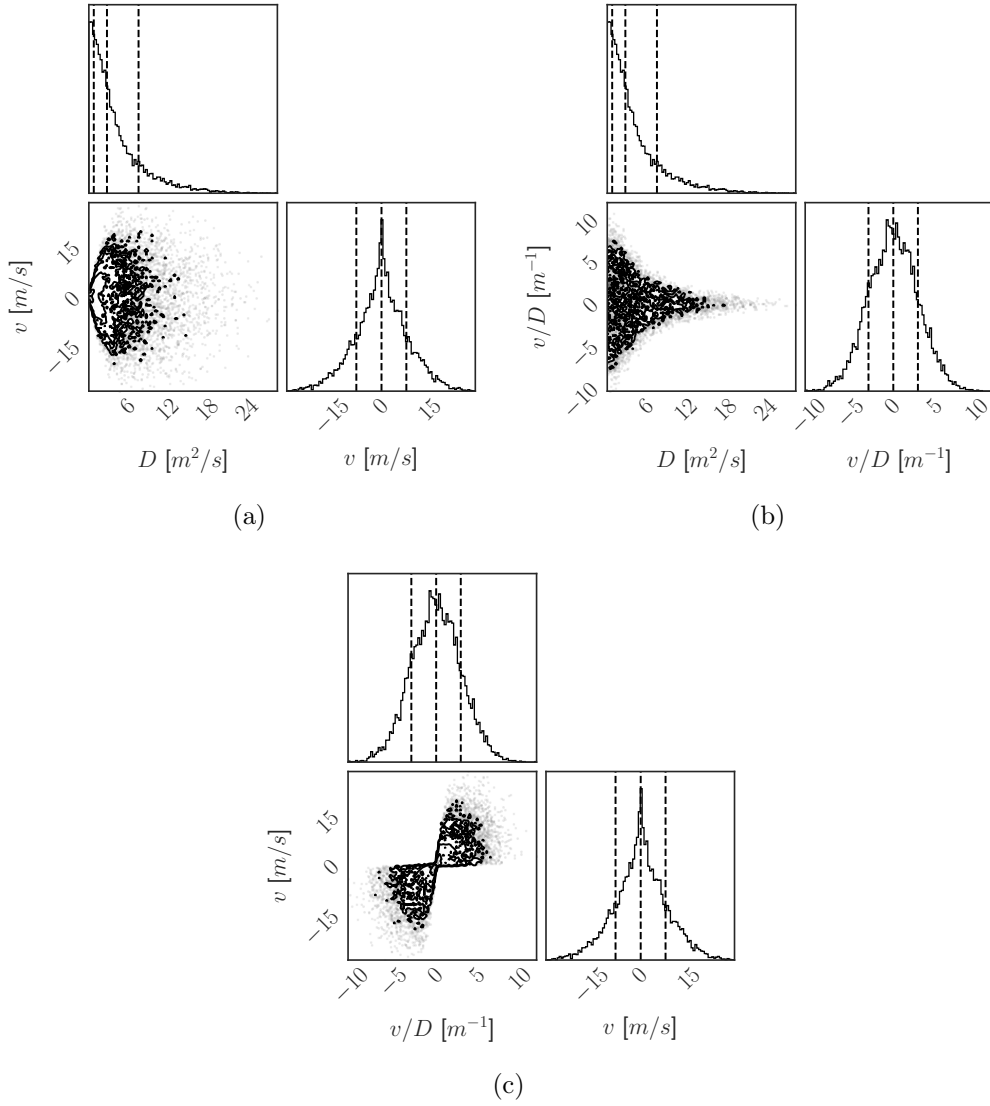


Figure 5-3: Physical sampling of diffusion (D), convection (v) and their ratio (v/D) for the case of $\chi = 1$, $\mathcal{P}_{\chi v/D} = \mathcal{N}(0, 5)$ and $\mathcal{P}_r = \mathcal{N}(0, 10)$.

can then be simply combined to give

$$D = \tilde{r} \cos(\tilde{\Theta}) \quad v = \tilde{r} \sin(\tilde{\Theta})/\chi \quad (5.19)$$

This scheme can be interpreted geometrically as follows: we use one degree of freedom to sample the ratio of χv and D (represented by the $\tan(\tilde{\Theta})$ variable) and another degree of freedom to sample the value of $\tilde{r} = \sqrt{|\chi v|^2 + D^2}$. Eq. 5.19 then projects the \tilde{r} radial distance onto the v and D dimensions. The parameter χ allows one to set different scales

for D and v while still sampling their \tilde{r} amplitude from the same distribution. A value of $\chi < 1$ would indicate that sampling of the magnitude of v must be allowed to explore large values. Fig. 5-3 offers a visual representation of this sampling scheme for the case of $\chi = 1$, $\mathcal{P}_{\chi v/D} = \mathcal{N}(0, 5)$ and $\mathcal{P}_r = \mathcal{N}(0, 10)$. These corner plots clearly show that the scheme effectively prevents sampling of unphysical values, trading prior freedom in D and v so as to maintain a reasonable v/D . One trade-off is that D and v amplitudes must use the same prior distribution. This is partly a limitation for D , because a rigorous Jeffreys' prior for D may be more appropriately set as a log-uniform distribution. However, this is a relatively arbitrary choice in our case, and the practical difference is expected to be always negligible in realistic inferences. All inferences presented in this work, for both Alcator C-Mod and DIII-D, used this technique, setting $\chi = 1/3$.

5.4.3 Combination of Constraints from Multiple Diagnostics

Given multiple instruments that provide partial constraints on the same quantities (in our case, particle transport coefficients), their combination is not trivial. One may choose to simply add the likelihood from each, or to assign a larger weight to one or another diagnostic – none of these choices is inherently more rigorous than the others if these measurements are completely independent. Therefore, one may want to let the data dictate which choice is most appropriate.

Inspired by the astrophysical literature (e.g. Refs. [188, 218, 219, 220]), we adopt a Gaussian likelihood over measurement errors and reconsider the common practice [102, 88] of weighing χ^2 from different diagnostics with fixed factors as

$$\chi^2 = \sum_k \alpha_k \chi_k^2 \tag{5.20}$$

where k is an index identifying an experimental signal to be matched. The α_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 α_k . By a maximum-entropy argument (see Appendix C 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. For the problem of inferring diagnostic weights, one must

also ensure that values are always positive. While no analytic solution to this constrained entropy maximization is available [219], the gamma distribution (of which the exponential is a special case) generally 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 a gamma prior for α_k analytically (see Appendix C). This allows us to effectively substitute the simple χ^2 metric with

$$\ln P(D|\theta) = \sum_{k=1} \left[\ln \Gamma \left(\frac{n_k}{2} + \nu \right) - \left(\frac{n_k}{2} + \nu \right) \ln \left(\frac{\chi_k^2}{2} + \nu \right) - \ln \Gamma(\nu) + \nu \ln \nu \right]. \quad (5.21)$$

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 χ^2 one, but sums over diagnostic weights that follow a gamma probability distribution with $\nu := a = 1/b$. By setting ν to different values, we allow α_k to be more or less free to weigh diagnostics differently. Appendix C shows some possible choices of ν that we considered. Note that there is one value of α_k for each diagnostic k , but since we analytically marginalized the likelihood over the α_k prior there is no direct sampling of α_k to be done: inclusion of these parameters in an inference does not incur 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. This shifts the balance from our *prior* expectations (personal bias) on appropriate diagnostic weights to the *likelihood*, empowering the data to determine the best weighing between diagnostics. Of course, corrupted data from one instrument can drive this method to unreasonable results, since the algorithm cannot possibly determine which signal is wrong unless it is a clear outlier from a large set of measurements. This technique must therefore be used with care and is by no means a panacea for the problem of weighing multiple measurements. Alternative options based on *effective sample sizes* for each instrument, computed from time series correlations in the data, were also explored. Generally, diagnostic weights based on this concept were found to be difficult to apply in a fully-Bayesian inference framework. On the other hand, a simple rescaling of the likelihood contribution from each diagnostic based on its number of data points was found to give a useful, albeit not rigorous, method to avoid biasing inferences

towards matching one instrument more than another. When working with diagnostics with widely different numbers of data points, this fixed rescaling was combined with the analytically-marginalized exponential prior or the gamma prior of Eq. 5.21, which provide some degree of freedom based on observed data.

5.5 On The Use of Manufactured Solutions

Previous work by Chilenski [184] examined the problem of inferring impurity transport coefficients using synthetic data, i.e. signals that were first created via a single run of the forward model with a chosen set of parameters. This is sometimes referred to as the *method of manufactured solutions*. Here, we wish to briefly comment on the value and limitations of this during the development of an experimental inference framework.

Undoubtedly, it is important to test an inference framework using synthetic data, as a way to *verify* the integrity and correctness of the code. This may also be a valuable approach to examine systematic issues with a given problem, e.g. to identify what size of uncertainties may be affordable in experimental data before inferences stop being effective. Synthetic data may also be helpful during the development of a framework to avoid bias towards specific data sets that will later be analyzed.

In a realistic setting, experimental data will necessarily deviate from the expectations of a synthetic data set. Typically, such deviations make the experimental analysis significantly more challenging. While significant emphasis must be put on the realistic size of uncertainties, as well as their appropriate characterization, the concept of *model inadequacy* is often of critical importance. Attempting to estimate parameters for a model that cannot accurately represent the observed physics can bias an inference task and make results harder to interpret. The inclusion of time-dependent background, sawtooth modeling, free knots, and charge exchange of impurities with background neutrals constitute significant improvements with respect to previous work. In the next chapter, we shall examine measurements and modeling of neutral dynamics in both Alcator C-Mod and DIII-D. Predictions of neutral densities will later provide the basis to assess the impact of charge exchange on particle transport inferences, presented in Chapters 7 and 8.

6

Neutral Particle Studies

This chapter discusses atomic deuterium (D) neutral measurements and modeling. Data from a D Ly_α midplane diagnostic on Alcator C-Mod are shown to be in good agreement with the state-of-the-art SOLPS-ITER model. A database of Ly_α signals is used to create empirical scaling laws for neutral density and penetration depth into the pedestal. We demonstrate how estimations of deuterium cross-field flux inferred from Ly_α data on C-Mod can form the basis for robust studies of pedestal particle transport. Finally, we illustrate how typical neutral densities, both from the edge and neutral beam injection, affect impurity charge state balance via charge exchange for both C-Mod and DIII-D.

6.1 Introduction

Charge eXchange (CX) is known to be of great importance in the scrape-off-layer and divertor regions, but is usually ignored for core and pedestal studies. In this chapter, we assess this approximation, considering both thermal and fast neutrals on Alcator C-Mod and DIII-D. The label “thermal” is used to describe particles that have a Maxwellian distribution, although they may not be in equilibrium with respect to the background plasma. This is the case for neutrals in the plasma edge. Another population of neutrals is present whenever Neutral Beam Injection (NBI) is used to heat and/or drive current in a tokamak discharge. While NBI operation was rare and only used for diagnostic purposes on Alcator C-Mod, neutral beams are active in most DIII-D discharges, where the lower plasma density allows better penetration. Since NBI neutrals are injected at energies near 75 keV, i.e. much higher than the background temperatures, these particles are referred to as “fast”. However,

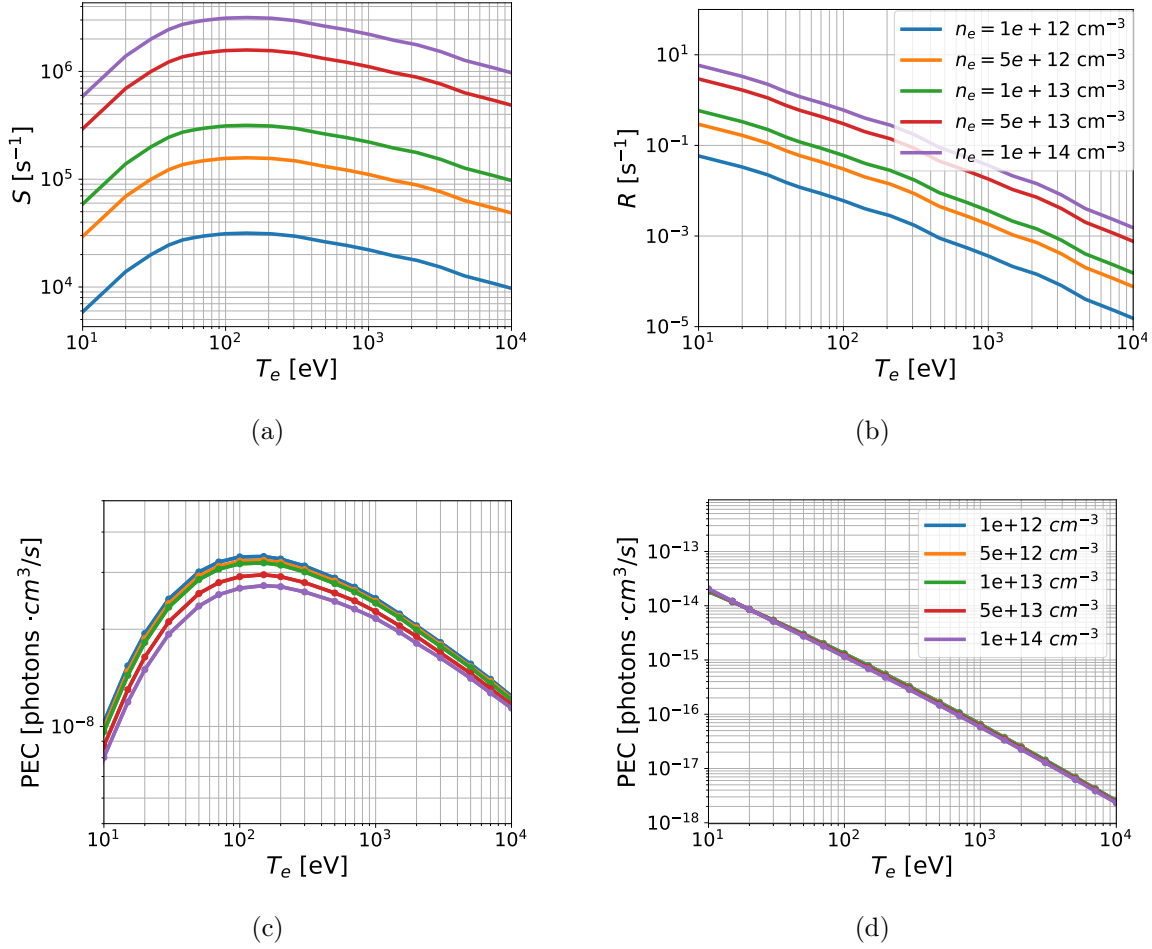


Figure 6-1: H rates of (a) ionization, (b) recombination, and Ly_α photon emissivity coefficients for (c) ionization-driven and (d) recombination-driven processes.

neutrals that arise from multiple generations of CX reactions starting from the fast NBI neutrals, commonly labelled “halos”, are at the same temperature as the plasma background from which they originate and are therefore also taken to be thermal. In this chapter, experimental measurements and modeling of neutral populations, as well as their interactions with impurity ions, will be discussed in depth. All the analysis discussed in this chapter makes use of Aurora, either to process the appropriate atomic rates or to post-process results from other codes (SOLPS-ITER and FIDASIM).

Fig. 6-1 shows effective (a) ionization (S) and (b) recombination (R) ADAS [114] rates of H, also applicable to D and T isotopes, as a function of electron temperature for multiple electron densities¹. The dependence on electron density is approximately linear. On the

¹ R here includes both radiative and dielectronic recombination contributions. Throughout this thesis, we denote the total recombination, including charge exchange, as α .

other hand, the peaking of S near 100 eV has important consequences: there is very little variation in ionization rate as neutrals penetrate further than the last closed flux surface (LCFS), which typically has $T_e \approx 75$ eV on both C-Mod and DIII-D. Balancing ionization and recombination, one finds that the fractional abundance of the neutral H stage over the proton density is typically higher than $\sim 10^{-8}$ even in the core plasma of existing tokamaks. The lower panels of Fig. 6-1 show the photon emissivity coefficients (PECs) from (a) excitation- and (b) recombination-driven processes for the H Ly $_{\alpha}$ ($n = 2 \rightarrow 1$) transition. The former of these is most significant, and, similarly to S , it peaks near 100 eV. These coefficients are central to the analysis of Ly $_{\alpha}$ brightness discussed in Sections 6.2 and 6.3.

In this chapter, we present analysis of experimental D Ly $_{\alpha}$ data from the Alcator C-Mod edge midplane, compare results to the state-of-the-art models for neutral dynamics, and broadly consider the impact of CX between impurity ions and background neutrals. We examine thermal neutrals in the Scrape-Off Layer (SOL) and near the edge of the confined plasma, as well as in the core via modeling of Neutral Beam Injection (NBI). For clarity, we use the symbol n_n to denote atomic neutral densities of deuterium (or, more generally, hydrogen isotopes) and n_0 to denote the neutral stage of any other ion. In the next section, we present a validation effort to compare Ly $_{\alpha}$ experimental measurements from the edge midplane on Alcator C-Mod with the comprehensive EIRENE [203] Monte Carlo neutral model within SOLPS-ITER [205, 204].

6.2 Comparison of D Neutral Measurements and SOLPS-ITER Predictions with EIRENE

The Alcator C-Mod D Ly $_{\alpha}$ data discussed here are from one L-mode, one Enhanced D-Alpha (EDA) H-mode and one I-mode discharge that have very similar engineering and physics parameters to those of discharges that will be later discussed in the context of impurity transport in Chapter 7.

Electron density, n_e , and temperature, T_e , in the pedestal and SOL region have been measured via Thomson Scattering [221, 222], complemented by Electron Cyclotron Emission [158] T_e data further into the core. To ensure reliable profile fitting across the LCFS, we make use of a parametric function that closely resembles the one originally introduced by Groebner and Carlstrom [223] and adopted by a large number of previous

pedestal studies. This can be written as

$$f(R) = b + \frac{h}{2} \left[\tanh \left(\frac{R_0 - R}{d} \right) + 1 \right] + m \times (R_0 - R - d) \times H(R_0 - R - d) \quad (6.1)$$

with baseline b , height h , position R_0 , full width D , and interior slope m . Here, $d = \Delta/2$ is the pedestal half-width. The Heaviside function, $H(R_0 - R - d)$, is used to account for a finite radial slope, m , in the pedestal region. The base of the pedestal is located at $R = R_0 + d$ and has value $f \approx b$, while the pedestal top has value $b + h$ at $R = R_0 - d$. The maximum radial derivative is at $R = R_0$ and has magnitude $|\nabla f|_{R_0} = h/\Delta$ [224].

The 2-point model described in Section 2.6 is used to constrain the electron temperature at the LCFS and thus limit the impact of misalignment between diagnostic measurements and magnetic equilibrium reconstructions. P_{Ohm} is calculated from plasma current, loop voltage and internal inductance of the plasma, all computed via magnetic diagnostics and the EFIT algorithm used for reconstructions of magnetic geometry [166]. P_{aux} is estimated by assuming that the radio frequency (RF) heating has an efficiency of 80% [225]. Radiated power is taken from cross-calibrated bolometry measurements.

We make use of Ly_α signals measured with 20 spatial chords going across the low-field side (LFS) midplane. Radially-localized emissivity estimates, ε , are obtained by assuming toroidal symmetry and using a regularized Abel inversion. We calculate atomic D neutral densities by dividing local emissivities, ε , by the Ly_α transition energy and the appropriate rate coefficients from ADAS (Fig. 6-1):

$$n_n = \frac{\varepsilon}{E_{ij} \cdot (n_e \mathcal{P}_{\text{Ly}_\alpha}^{exc} + n_i \mathcal{P}_{\text{Ly}_\alpha}^{rec})}. \quad (6.2)$$

Here, we have indicated how PECs for both excitation- and recombination-driven population of the $n = 2$ contribute to Ly_α emission, with n_i being the ionized D density, although the recombination component is in practice always negligible as seen from Fig. 6-1. Details of the Ly_α data analysis procedure are described in Section E.2 of Appendix E.

Previous research by Hughes *et al.* [159] compared measurements of n_n from Ly_α with the Kinetic Neutral 1D (KN1D) model [200], particularly focusing on pedestal structure and core fueling. Relatively good agreement was found between the *shape* of radial neutral density profiles from KN1D and Ly_α measurements, but KN1D cannot predict absolute neutral sources. Here, we revisit some of the Ly_α data to compare to the more advanced

EIRENE Monte Carlo neutral model included within the SOLPS-ITER code [203].

The SOLPS-ITER results discussed here have been obtained in collaboration with R. Reksoatmodjo by iteratively modifying input heat and particle diffusivities until matching TS n_e and T_e profiles. The setup of these runs is similar to the one presented in Ref. [226], except for the fact that we have attempted to match kinetic profiles further into the core, since the focus of this validation effort is the penetration of neutrals and their effect on pedestal ionization balance. Up to 50,000 Monte Carlo neutrals have been run in each EIRENE simulation for good statistics. For simplicity, the SOLPS-ITER discussed here include only 2 species, they do not use explicit fluid drifts and they do not aim to match measured divertor heat fluxes. These simplifications are expected to have negligible influence on the inferred penetration of D neutrals through the pedestal. In this sense, the present comparison to Ly_α data offers a validation exercise that we believe is unaffected by the most complex numerical challenges of SOLPS-ITER simulations.

Fig. 6-2 shows the SOLPS-ITER prediction for atomic D neutral density for the I-mode discharge #1080416025, which has very similar engineering and physical parameters to the I-mode discharge (#1101014030) that will be discussed in Chapter 7 in the context of impurity transport inferences. The Aurora post-processing tools for SOLPS-ITER allow one to extract SOLPS-ITER results both on the B2(.5) grid and the EIRENE grid. Interpolation of the 2D distributions of any output permits comparison to the C-Mod Ly_α data at the outer midplane, as well as calculation of flux surface averages, as will be discussed below.

Fig. 6-3 shows the comparison of midplane Ly_α measurements with the SOLPS-ITER results, interpolated to the low-field side, for the 3 shots of interest. Red lines represent SOLPS-ITER results, while the black regions represent the experimental uncertainty range, with color intensity representing the Gaussian probability density function. Figures on the left show the base-10 logarithm of atomic D neutral density, those on the right use a normalization by the local n_e . In evaluating these comparisons, we remark that uncertainties here have been evaluated by propagating uncertainties in n_e , T_e and Ly_α emissivity profiles. By comparing the ADAS rates to those from the Janev & Smith database [227], we have found $\sim 10\%$ discrepancies in neutral densities, which may therefore be interpreted as an additional source of uncertainty. Most importantly, an uncertainty on the absolute calibration of the Ly_α system for the three shots, taken over a span of 2 years, is not available and is likely to affect absolute neutral densities by up to a factor of 2. Since




Figure 6-2: Spatial distribution of atomic D neutral density from SOLPS-ITER modeling of the Alcator C-Mod discharge #1080416025. Orange contours on top of the neutral densities represent poloidal flux contours, with the thickest of the lines showing the LCFS.

the Thomson Scattering diagnostic measures n_e and T_e near the top of the plasma and these quantities are not strictly flux-surface functions in the SOL, we have checked whether the SOLPS-ITER results suggest significant discrepancies between the true location of TS measurements and the midplane locations to which these are mapped via flux surfaces. If so, this would affect both the Ly_α experimental analysis, since n_e and T_e are used to process signal emissivities, and the SOLPS-ITER results, since transport coefficients in the code were tuned to match n_e and T_e . However, negligible differences have observed between the

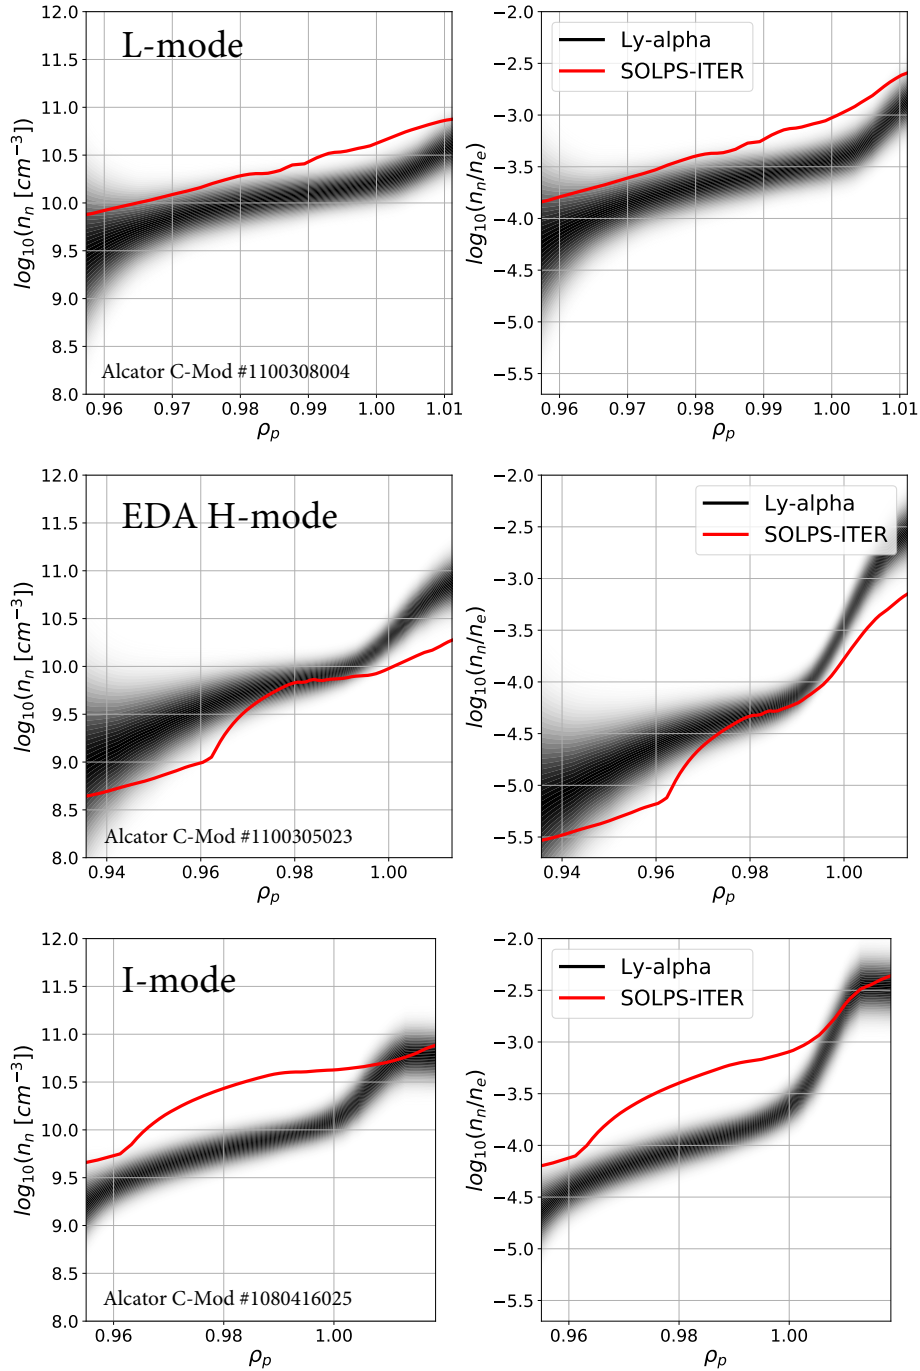


Figure 6-3: Comparison of SOLPS-ITER (EIRENE) predictions for neutral midplane radial profiles with Ly_{α} spectroscopy in the C-Mod (a) L-mode, (b) EDA H-mode, and (c) I-mode shots discussed in the text. Note that radial ranges differ for the three cases to focus on regions that are well covered by Ly_{α} and Thomson scattering data.

two locations, suggesting that this is not a significant source of error.

Thus, Fig. 6-3 offers a valuable comparison between high-quality experimental data and

state-of-the-art models. We note that the SOLPS-ITER results are themselves affected by significant uncertainties, for example those related to the recycling model at the wall and the choice of transport coefficients used to match experimental n_e and T_e data points.

Overall, SOLPS-ITER appears to be rather successful at matching the experimental data in Fig. 6-3. Discrepancies by up to a factor of ≈ 5 are found; nonetheless, for the aforementioned reasons, profile shapes may be argued to be more important than absolute densities themselves. While in the L-mode case the profile gradients are very well recovered, we are unable to explain some of the discrepancies found in the other SOLPS-ITER runs at present.

The comparisons of Fig. 6-3 should be considered “spot checks” at the LFS midplane. In practice, to assess the impact of neutrals on impurity transport studies, flux-surface average (FSA) neutral densities are more appropriate than LFS values. Using the same Alcator C-Mod I-mode discharge as in Fig. 6-3, Fig. 6-4 motivates this statement by comparing the effective ionization rate for each charge state with the characteristic transit time that passing and trapped impurity ions take to travel a parallel distance $L = qR$, where q is the safety factor and R is the major radius. As in recent work by Dux *et al.* on AUG [228], we define a normalized ionization rate as

$$\nu_{ion}^* \equiv S\tau_t = S \frac{qR}{v_{th}} \equiv \frac{\sum_z n_z S_z^{ion}}{\sum_z n_z} qR \sqrt{\frac{m_{imp}}{2k_B T_i}} \quad (6.3)$$

where S_z is the ionization rate of each charge state of density n_z and S is taken to be the sum over ionization rates weighted by the density of each charge state. If $\nu_{ion}^* < 1$, flux surface averaging of background asymmetries (e.g. from edge or beam neutrals) may be taken as a good approximation to reality; in this case, 1.5D simulations of impurity transport are expected to be valid. If, on the other hand, $\nu_{ion}^* > 1$ then local effects may be too important to ignore. The radial profiles of Fig. 6-4 show that in practice $\nu_{ion}^* \ll 1$ everywhere except very close to the LCFS, where the validity of a 1.5D model of particle transport may be questionable regardless. This suggests that, in agreement with the AUG results of Ref. [228], charge exchange of neutrals and impurities may be treated as a flux-surface averaged process. This justifies our use of charge exchange in Aurora’s 1.5D forward model even when edge neutral densities display strong poloidal asymmetries near the LCFS or in the presence of a directed energetic neutral beam (as discussed in Section 6.5).




Figure 6-4: Radial profile of the normalized effective ionization rate ν_{ion}^* , as defined in the text. We also plot the square root of the inverse aspect ratio, $\sqrt{\epsilon}$, directly proportional to the trapped particle fraction, similarly to Dux *et al.* [228].

In the next section, we delve more deeply into the experimental database of Ly_α signals from Alcator C-Mod and describe a number of scaling laws aiming to elucidate key dependences of neutral density in the pedestal and at the LCFS.

6.3 C-Mod Deuterium Ly_α Database

While a comparison of experimental D Ly_α data and SOLPS-ITER results (Section 6.2) has been possible only in a few cases, due to the complexity of running SOLPS-ITER, a D Ly_α database has been constructed for a much larger number of shots. A key objective of this study is to extract regressions for the neutral density at the LCFS and for the neutral density decay length in the pedestal. Based on these two quantities, one may in principle estimate the entire profile of background D neutrals in the confined plasma. This allows one to quantify the impact of charge exchange on impurities in any Alcator C-Mod discharge whose parameters are sufficiently near those within the present database. In this section, we present the main results from this study.

6.3.1 Database Development and Methodology

The database described here has been developed based on discharges from 2008 with no ELMs. While much more data were available in principle from other years, this sub-selection was made to ensure that the calibration of the Ly_α diagnostic would be unchanged within the database. Discharges were filtered based on the availability of edge Thomson Scattering data for n_e and T_e , as well as reliable Ly_α emissivities and D_2 edge pressure measurements.

The latter were obtained via ionization gauges at several points in the plasma periphery, including at several locations in the divertor regions. The final database comprises 262 discharges in total, of which more than half (159) are Ohmic L-mode shots, 38 are RF-heated L-mode shots, 61 are H-modes, and 4 are I-modes. In each case, we average over 100 ms time windows identified as “quasi-steady” based on small variability of engineering parameters, radiated power, electron density, electron temperature, and wall pressure within this time range. This results in approximately 1,000 data points. Future work will extend this analysis to more discharges, particularly in high-performance regimes. The current database is biased towards discharges with $I_p = 0.8$ MA, since approximately two thirds of the selected discharges are at this plasma current, although the remaining one third is well distributed between 0.4 and 1.3 MA. The data span a range of Greenwald fractions (defined via density volume averaging) between 0.1 and 0.5, with an approximate Gaussian distribution centered near 0.25. Past research on Alcator C-Mod has shown that the outer divertor is robustly detached at Greenwald fractions greater than ≈ 0.3 in Ohmic L-mode discharges (possibly at higher values in other regimes) [229].

As discussed in Sections 6.2 and E.2, Ly_α emissivities are processed using ADAS rates and modified-tanh fits to n_e and T_e to obtain atomic D neutral densities. Approximately one third of the discharges in this database also included A-port Scanning Probes (ASP) data, which provide detailed measurements of n_e and T_e in the SOL. The availability of these data give greater confidence in kinetic profile reconstructions outside of the LCFS, but affects our pedestal analysis relatively weakly. Scanning probes would be much more critical if one were to examine gradient scale lengths near the LCFS, as in past work by Hughes [159], but doing so for our database is left for future work. Given that the 2-point model provides a robust estimate for T_e at the LCFS, we choose to examine neutral densities from Ly_α at that location. This effectively reduces the degrees of freedom and uncertainties of our regression analysis, based on a clear physical criterion (see Section 2.6). Fig. 6-5 shows histograms and distribution densities for some of the key parameters of interest for this study. On the left, the plasma current, I_p , is seen to be dominated by 0.8 MA shots, although approximately one third of the discharges are at I_p ranging from 0.5 to 1.2 MA. Other displayed parameters include the safety factor at $\Psi_n = 0.95$ (q_{95}), the midplane D_2 pressure ($p_{D2,mid}$), the non-radiated power crossing the SOL normalized by the LCFS area (P_{sol}/A), the Greenwald density fraction (f_{GW}), the difference between the upper and lower

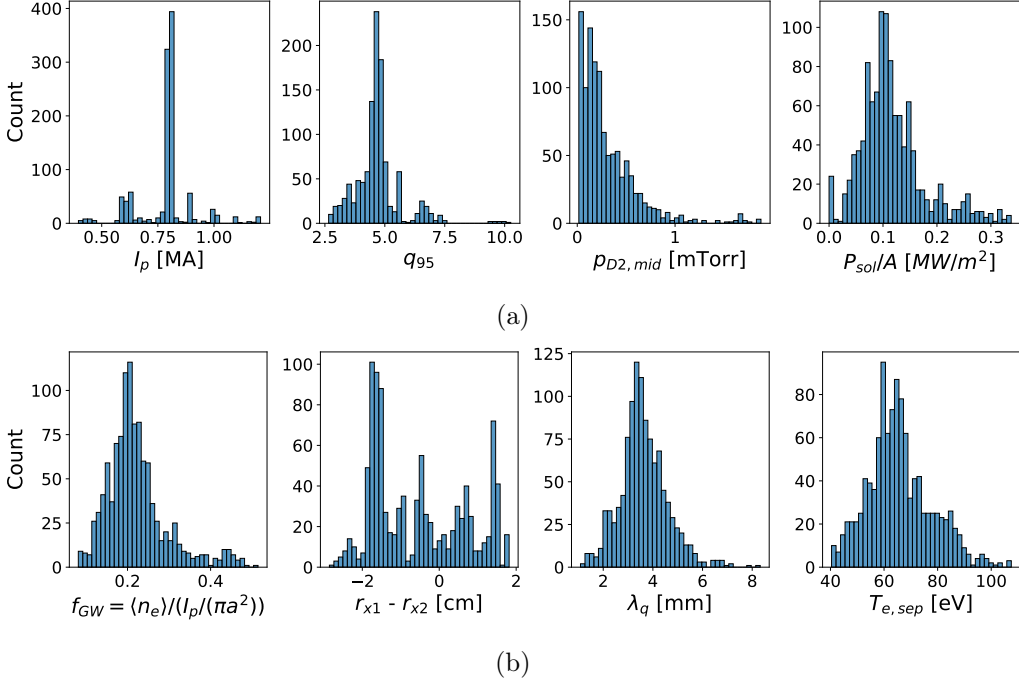


Figure 6-5: Histograms of some key parameters within the C-Mod Ly_α database. The distribution of plasma current is strongly clustered near $I_p = 0.8$ MA, although approximately 1/3 of values are at other currents between 0.5 and 1.2 MA. The other parameters, described in the text, have relatively broad distributions.

x-point radial locations mapped to the midplane (the EFIT *seep* parameter), the divertor heat flux width (λ_q , taken from the Brunner scaling [72], see Section 2.6), and the separatrix electron temperature from the 2-point model ($T_{e,sep}$). Except for I_p , all parameters have relatively broad distributions within the database. The $seep \equiv r_{x1} - r_{x2}$ parameter is positive for Upper Single Null (USN) discharges and negative for Lower Single Null (LSN) ones. We define “near-Double-Null” (DN) cases as having $-0.1 < seep < +0.1$ [cm]. Unlike in AUG work by Kallenbach *et al.*, on C-Mod the D_2 pressures measured in the upper and lower divertors have only a weak impact on the neutral and electron densities at the separatrix. This finding, clearly reflected in our database, is central to the “main-chamber recycling” paradigm that has been investigated for many years on Alcator C-Mod [230, 231, 229].

On C-Mod, the upper divertor was “open” (with flat surfaces), while the lower divertor was “closed” (with a near-vertical target plate). Closed divertors aim at better confining neutrals and impurities introduced for divertor heat flux reduction [229]. In the database presented here, the toroidal magnetic field has a fixed direction that results in the ion ∇B particle drifts always pointing down. As a result, all LSN discharges have ∇B drifts pointing

towards the divertor, which is generally referred to as the “favorable direction” since it makes transitioning to H-mode easier. On the other hand, all USN discharges in this database have ∇B drifts pointing away from the divertor, which is known to be an unfavorable condition for H-mode access. This means that within the presented database there is a complete covariance of “favorable vs. unfavorable” and “LSN vs. USN”. Consequently, one cannot distinguish dependences on the ∇B direction vs. the divertor geometry based on this database alone.

In conformity with standard practices in the field, we quantify relationships in our regressions in the form of power laws. These offer simplicity, sometimes at the cost of being physically difficult to interpret. Over this study, the following normalizations are used: densities are in units of 10^{20} m^{-3} , I_p is in MA, pressures are in mTorr, magnetic fields in T, P_{sol}/A (power going from the core into the SOL divided by the LCFS area) in MW/m^2 , pedestal scale lengths are in mm, particle fluxes are in $10^{20} \text{ particles}/(\text{m}^3 \cdot \text{s})$, and cumulative gas fueling in $\text{Torr} \cdot \text{l}$. The regressions shown below make use of the Trust Region Reflective algorithm, available within the *scipy* Python package [232]. In order to select appropriate regression variables, we supplement simple statistical coefficients of determination (R^2) with metrics that aim to prevent over-fitting to the data. We make use of the common Bayesian Information Criterion (BIC) and the Akaike Information Criterion (AIC), both approximations of the Bayesian evidence defined in Chapter 5 and defined, respectively, as

$$\text{BIC} = k \ln(n) - 2 \ln(\hat{L}) \rightarrow k \ln(n) + 2\hat{\chi}^2, \quad (6.4)$$

$$\text{AIC} = 2k - 2 \ln(\hat{L}) \rightarrow 2k + 2\hat{\chi}^2, \quad (6.5)$$

where k is the number of estimated parameters in the model, \hat{L} is the maximum likelihood value for the model, and n is the total number of data points. The arrows above indicate our reduction of these formulae for the case of Gaussian likelihoods, which introduces square residuals ($\hat{\chi}^2$). The BIC formula (Eq. 6.4) results from the maximization of a posterior model probability, whereas the AIC (Eq. 6.5) aims at minimizing the Kullback-Leibler divergence between the true distribution and the estimate from the model. The latter is typically recommended for cases where the true model is infinite-dimensional and is not strictly part of the set of tested models [233]. For completeness, we indicate R^2 , BIC and AIC values for each of our regressions, although we only make direct use of the AIC for model selection.

We make use of a recursive model selection scheme that involves computing regressions for all permutations of variable combinations within a chosen set, evaluating a power law relating these to the dependent variable (*regressand*) of interest, and eliminating independent variables, one by one, at each iteration. The initial set is chosen to include approximately 10 variables that are physically related to the regressand and could therefore be expected to be part of the final model-selected power law. Analogously to the way model selection is made based on evidence ratios for Bayesian inferences of particle transport, we compute the difference between the AICs of all pairs of regressions, defining the metric $\Delta_i = AIC_i - \min(\{AIC\})$. The evidence ratio is then [234]

$$ER_i = \frac{\exp(-\frac{1}{2}\Delta_{\text{best}})}{\exp(-\frac{1}{2}\Delta_i)} \quad (6.6)$$

and one can define an Akaike weight for a given model as

$$w_i = \frac{\exp(-\frac{1}{2}\Delta_i)}{\sum_{r=1}^R \exp(-\frac{1}{2}\Delta_r)}. \quad (6.7)$$

We compute such weights for all models (combinations of regression parameters) and sum them so as to estimate the relative importance of all variables under consideration [234], eliminating the one with the lowest total weight from our set. We repeat this elimination process until only 3 variables are left. At this point, we compute the AIC for all combinations of remaining parameters and simply select the model that gives the lowest AIC. We have found this procedure to be effective at identifying the most important variables, but it can suffer from ambiguity when strong correlations exist, e.g. between B_t , q_{95} and I_p . In this study, when such issue has arisen, we have individually inspected the regression models of lowest AIC and selected the one with the highest R^2 .

6.3.2 Neutral Density Regressions

Fig. 6-6 shows a power law regression for the measured neutral density at the separatrix, $n_{n,sep}$, in terms of D_2 pressures at the midplane, $p_{D2,mid}$, in the lower divertor (at the B port), $p_{D2,divB}$, and in the upper divertor (near the cryopump), $p_{D2,cryo}$. The ordinate gives measured values; the abscissa shows the regression result. The dashed line represents a 1-to-1 relationship, corresponding to the power law shown at the top of the figure. All

displayed data points are from L-mode discharges (Ohmic and RF-heated) without cryopump in operation. The colors and symbols for the data points in Fig. 6-6 indicate whether each case was in LSN (green downward triangles), USN (blue upward triangles) or DN (red diamonds). Recall that, due to the fixed direction of B_T in this database, the $B \times \nabla B$ vector was always pointing downwards, so LSN shots always have ion ∇B drifts in the favorable direction and a closed divertor, and vice versa. Too few DN cases are present in the database to make clear statements about DN geometries. While most of the data points with highest neutral densities are in LSN (favorable ion ∇B drift, closed-divertor) conditions, no clear separation from USN data points is found. We note that the 2-point model constrains T_e at the LCFS, and since T_e is measured by the same diagnostics as n_e , the 2-point model indirectly reduces mapping inaccuracies in the LCFS reconstruction for n_e too.

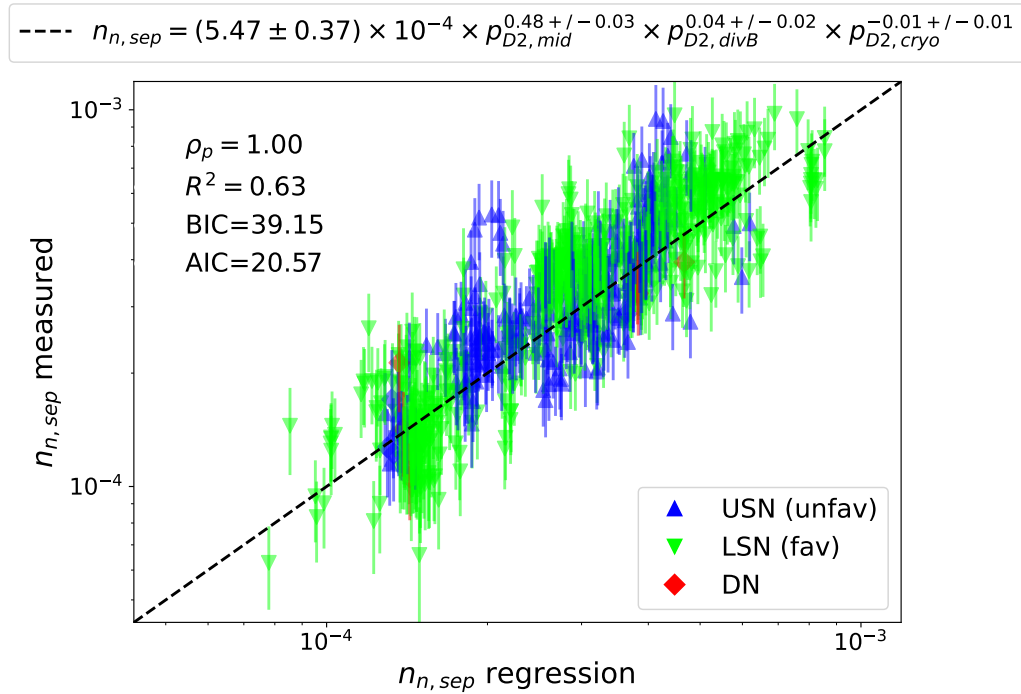


Figure 6-6: Regression for the atomic D neutral pressure measured at the C-Mod midplane in terms of midplane ($p_{D2,mid}$), lower-divertor ($p_{D2,divB}$), and upper-divertor ($p_{D2,cryo}$) wall D_2 pressures. Colors and symbols distinguish USN, LSN and DN cases.

The regression in Fig. 6-6 is *not* a result of the model selection procedure described in Section 6.3.1; rather, it was obtained with an arbitrary choice of variables in order to illustrate the much greater importance of the midplane pressure over divertor pressures

on C-Mod, indicative of main-chamber recycling being the main source of neutrals. In no circumstances in this study we have found results deriving from the outboard midplane Ly_α data to correlate with divertor pressures better than with midplane pressure.

Fig. 6-7 shows a regression over the same dataset of L-mode discharges without cryopump operation, but now only in terms of the midplane wall pressure, which is the only variable that is selected by the procedure of Section 6.3.1. The result of this regression is

$$n_{n,sep} = (6.06 \pm 0.11) \times 10^{-4} \times p_{D2,mid}^{0.50 \pm 0.01}. \quad (6.8)$$

Whereas data symbols distinguish USN, LSN and DN cases, as in Fig. 6-6, colors here indicate the amount of gas injected in the discharge up to the time window of interest, S_{gas} . There appears to be a trend of greater $n_{n,sep}$ with S_{gas} , but this is weaker than the correlation with $p_{D2,mid}$. For this reason, $p_{D2,mid}$ was automatically selected as a better regression variable than S_{gas} itself. We note that while S_{gas} and $p_{D2,mid}$ are clearly positively correlated, the value of $p_{D2,mid}$ can vary by one order of magnitude for a given value of S_{gas} in our dataset.

Adding other variables to the regression of Fig. 6-7 is strongly discouraged by our model selection procedure (Section 6.3.1) and does not change the R^2 value significantly. However, H-mode cases appear to have a larger exponent for $p_{D2,mid}$ and a stronger dependence on $T_{e,sep}$, with the regression

$$n_{n,sep} = (5.20 \pm 6.05) \times p_{D2,mid}^{0.89+/-0.05} \times T_{e,sep}^{-2.13+/-0.26} \quad (6.9)$$

achieving $R^2 = 0.60$. The fact that $T_{e,sep}$ is found to be more important for $n_{n,sep}$ predictions in H-mode shots suggests significant differences in SOL temperature profiles between confinement regimes, but we note that our H-mode database here only includes approximately 150 time slices, mostly at a single magnetic field intensity (5.5 T), hence this conclusion may not hold universally. Interestingly, the H-mode cases appear to have stronger correlations between $n_{n,sep}$ and divertor pressures, although the midplane pressure is always the dominant variable. Comparing cases with and without cryopump operation, one finds that the latter reduces the correlation between $n_{n,sep}$ and $p_{D2,mid}$, and consequently the R^2 of any model-selected regression, for any magnetic geometry.

The separatrix neutral density, $n_{n,sep}$, may be expected to correlate with the SOL




Figure 6-7: Model-selected regression for the atomic D neutral pressure measured at the C-Mod midplane in terms of midplane wall pressure, $p_{D2,mid}$, in L-mode discharges without cryopump operation. Symbols distinguish between USN, LSN and DN cases, whereas colors indicate the amount of injected gas.

electron temperature decay length, λ_T , since this determines how far neutrals must travel in the radial direction before they reach the LCFS at the T_e value given by the 2-point model. As described in Section 2.6, T_e at the LCFS depends only weakly on P_{sol} ($T_{e,LCFS} \propto P_{sol}^{2/7}$), hence the separatrix electron temperature itself is only weakly varying between shots (see the histogram in Fig. 6-5b). On the other hand, for larger λ_T , more neutrals get ionized in the SOL rather than inside the LCFS. In the Spitzer-Härm regime (attached and non-radiative divertor, dominated by electron-heat conduction), $\lambda_q = 2/7 \lambda_{T_e}$. From both the Brunner [72] and Silvagni [235] λ_q scalings, one sees that higher pedestal pressures mean smaller λ_q , and therefore smaller λ_T . In view of these relationships, we have explicitly tested whether n_n is correlated with λ_q . We find that the R^2 is hardly affected by the inclusion of λ_q among the regression parameters, and both the BIC and AIC strongly disfavor it.

As shown by Eq. 3.12, the balance of radiative and dielectronic recombination with charge exchange for heavy ions is set by the ratio of the background D atomic neutral density to the electron density, n_n/n_e . In order to quantify the impact of CX over the pedestal, this

quantity must be determined as a function of radius. It is therefore interesting to combine the $n_{n,sep}$ predictions presented in this section with estimates of neutral penetration depth into the pedestal, which are the subject of the next section.

6.3.3 Neutral Penetration Depth

We obtain an estimate for the D neutral penetration depth into the pedestal, L_{n_n} , by fitting an exponential to the neutral density profile from Ly α spectroscopy for each case. This allows one to constrain the radial decay with a clear parametrization, which is found to apply well in these discharges, as in previous experimental [159] and modeling work [226]. The alternative method of directly computing gradients from n_n profiles is strongly affected by details of the Abel inversion and possible outliers. One drawback of using exponential fits is that the statistical uncertainties obtained for L_{n_n} are unrealistically small, due to the lack of flexibility of this parametrization. We thus set a constant uncertainty for all L_{n_n} estimates. The actual value of such uncertainty is inconsequential for our regressions.

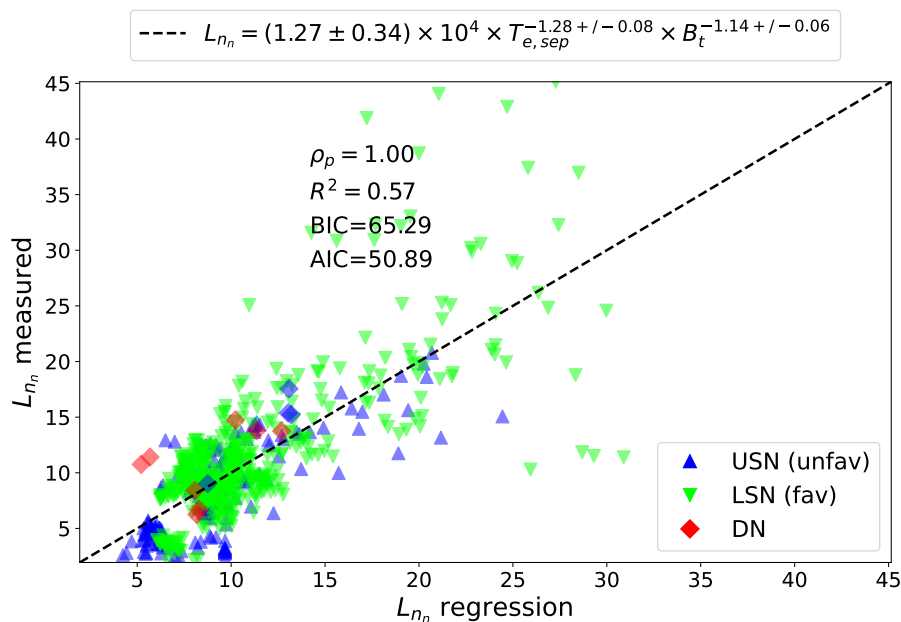


Figure 6-8: Model selected regression for the neutral penetration depth, L_{n_n} , in terms of $T_{e,sep}$, the electron temperature at the separatrix, in L-mode plasmas.

Fig. 6-8 shows the model selected regression on L_{n_n} as evaluated from the database subset containing only L-mode cases. Dependencies on a large set of parameters, divertor pressures, $n_{e,sep}$, q_{95} , $L_{n,e}$, I_p , W_{MHD} , and P_{sol}/A , were found not to be statistically significant. The

optimally chosen parameters are the separatrix electron temperature, $T_{e,sep}$, and the toroidal magnetic field, B_t . The regression

$$L_{n_n} = (1.27 \pm 0.34) \times 10^4 \times T_{e,sep}^{-1.28+/-0.08} \times B_t^{-1.14+/-0.06} \quad (6.10)$$

is found to give a relatively high $R^2 = 0.57$, indicating that L-mode penetration can be appropriately estimated based on $T_{e,sep}$ and B_t . H-mode cases show no visible trends within this database, with most values of L_{n_n} being in the 3-5 mm range.

The simplest interpretation of the observed $T_{e,sep}$ dependence in Eq. 6.10, with a negative exponent of -1.28 , is that higher temperatures cause higher ionization rates and therefore lower neutral penetration. We note that D ionization rates are a steep function of T_e for $T_e < 100$ eV, but they have a weak (decreasing) dependence on T_e for higher temperatures (see Fig. 6-1a). All of the L-mode cases in Fig. 6-8 have $T_{e,sep}$ well below 100 eV (see Fig. 6-5b, which also comprises H-mode discharges with larger $T_{e,sep}$), but T_e typically increases above 100 eV just inside the LCFS. There exists a competing effect from CX that may be expected to contribute to a positive exponent of $T_{e,sep}$ in Eq. 6.10. CX rates are a strong function of T_i , so assuming $T_i \approx T_e$ in the dense conditions of the SOL, one may expect CX to give greater penetration at greater T_e for a fixed ionization rate. Moreover, CX reactions between a cold D neutral and a thermal D ion result in a new neutral whose energy is given by the local T_i , hence suggesting that greater T_e should give rise to greater penetration (larger L_{n_n}). It is worth remarking that the negative exponent in Eq. 6.10 may also be related to hidden variables, which make $T_{e,sep}$ generally correlate with higher performance (including higher plasma density) and lower neutral penetration. The inverse dependence on B_t is indicative of this physically-meaningful trend. We note that an approximate B_t^{-1} dependence suggests a possible $\rho_* = \rho_i/a \propto \sqrt{T_e}/B_t$ scaling, typical of tokamak plasma transport, even though L_{n_n} is a *neutral*, rather than *plasma*, quantity.

6.3.4 Separatrix Electron density

As discussed in Section 2.6, the 2-point model offers a physical constraint to align Thomson scattering data based on a theoretical T_e value at the separatrix. While this offers valuable improvements to the quantification of $n_{e,sep}$, identifying trends in this quantity is notoriously complex. Figs. 6-9a and 6-9b present model-selected regressions of $n_{e,sep}$ for L- and H-modes

in our database, respectively. The quality of the scaling for H-modes is relatively low, with only $R^2 = 0.10$ for this choice of regression parameters. In both cases, we find a similar dependence of $n_{e,sep}$ on midplane wall pressure, approximately scaling as $p_{D2,mid}^{0.3}$. This is reminiscent of the result found by Kallenbach *et al.* from AUG data, where $n_{e,sep}$ was observed to scale with p_0 to the power of 0.31. In this study, however, p_0 represented the *divertor* pressure rather than the midplane one, as a result of AUG being mostly fueled from the divertor [236]. In our database, $n_{e,sep}$ regressions also support the introduction of additional engineering parameters, specifically the toroidal magnetic field, B_t , and the plasma current, I_p , without which the R^2 would be significantly lower. Unfortunately, B_t is mostly fixed at 5.5 T in our H-mode database, hence dependencies on this variable cannot be extracted. Nonetheless, the higher R^2 found in the L-mode regression of Fig. 6-9a suggests that B_t does help to explain some of the variability of $n_{e,sep}$.

6.3.5 Cross-Field Deuterium Pedestal Transport

We next examine the rate of neutral ionizations per surface area at the LCFS. This is equivalent to the total rate of deuterium fueling in the confined plasma, S_D . Let us consider the deuterium continuity equation at steady state

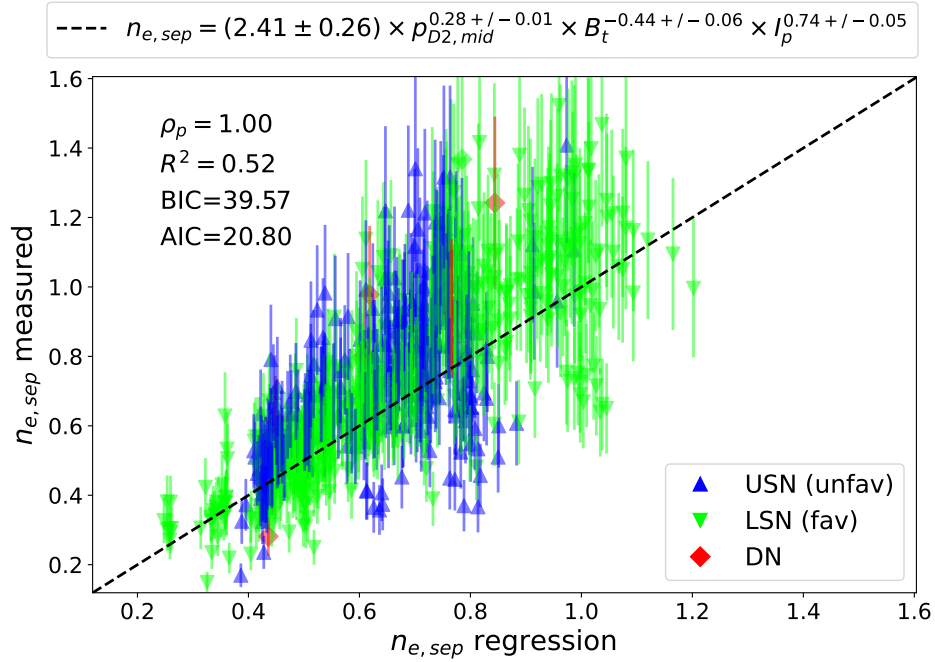
$$\frac{dn_D}{dt} \overset{0}{=} -\nabla \cdot \Gamma_D + S_D \rightarrow \Gamma_D = \int S_D dr, \quad (6.11)$$

where n_D is the deuterium density and the latter relation arises from setting its time derivative to zero. We can expand this expression as


$$\Gamma_D = \int_{r=0}^{r=r_{sep}} n_n n_e S_{ioniz} dr, \quad (6.12)$$

where S_{ioniz} is D neutral ionization rate, in units of $m^{-3}s^{-1}$. Eq. 6.12 involves a radial integration over D neutral densities obtained from the Ly_α signal inversion. This quantity does not depend on the exact n_n profile, but only its integral, which is more resilient to inaccuracies in Abel inversions. Moreover, Γ_D is only weakly dependent on $n_{e,sep}$, which is notoriously hard to determine accurately, since

$$\Gamma_D = \int_{r=0}^{r=r_{sep}} \frac{\varepsilon}{n_e \mathcal{P}} n_e S_{ioniz} dr = \int_{r=0}^{r=r_{sep}} \varepsilon \frac{S_{ioniz}}{\mathcal{P}} dr. \quad (6.13)$$



(a)



(b)

Figure 6-9: Model-selected regressions for the separatrix electron density, $n_{e,sep}$, for (a) L-mode, and (b) H-mode discharges. In (a), dependencies on midplane wall pressure, $p_{D2,mid}$, B_t value, and plasma current, I_p , are identified. In (b), a dependency on B_t is omitted because there is little variation of this variable within the current H-mode database.

We also remark that the ratio of the D ionization rate, S_{ioniz} , and the D Ly_α PEC, \mathcal{P} , is itself only weakly dependent on density. This means that Γ_D estimates are resilient to errors in the local quantification of n_e . In fact, the ratio of S_{ioniz} and \mathcal{P} , often referred to as the “ionization per photon ratio” [114], is nearly constant as a function of electron temperature above approximately 30 eV, as shown in Fig. 6-10 for multiple relevant values of n_e .²

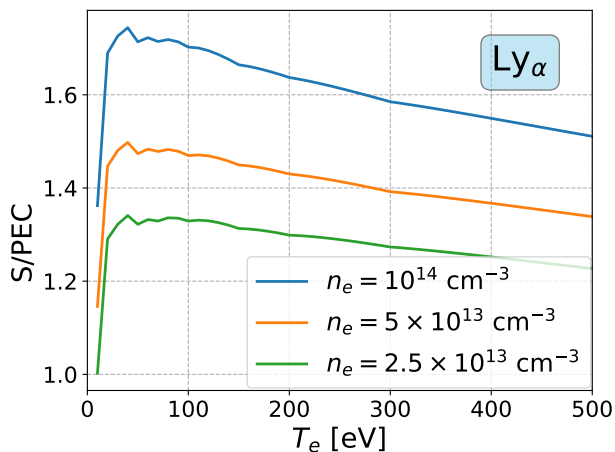


Figure 6-10: Ionization per photon ratio for the Ly_α transition as a function of electron temperature for $n_e = 10^{14} \text{ cm}^{-3}$ (blue), $5 \times 10^{13} \text{ cm}^{-3}$ (orange) and $2.5 \times 10^{13} \text{ cm}^{-3}$ (green).

Fig. 6-10 suggests that, for our purposes, one may take

$$\Gamma_D = \int_{r=0}^{r=r_{\text{sep}}} \varepsilon \frac{S_{\text{ioniz}}}{\mathcal{P}} dr \propto \int_{r=0}^{r=r_{\text{sep}}} \varepsilon dr, \quad (6.14)$$

with a constant of proportionality typically having a value near 1.5 for C-Mod plasmas. Consequently, trends in Γ_D can be studied with remarkable accuracy simply by calculating the radial integral of Ly_α emissivities inside of the LCFS. This makes Γ_D a more robust metric than n_n itself. In practice, in most circumstances it is still preferable to compute the integral in Eq. 6.13 rather than make an approximation about the value of $S_{\text{ioniz}}/\mathcal{P}$.

Figs. 6-11 and 6-12 show the results of model-selected regressions for Γ_D in USN (unfavorable ion ∇B drift direction, open divertor) and LSN (favorable ion ∇B drift direction, closed divertor) discharges, respectively. The midplane wall pressure, $p_{D2,\text{mid}}$, is selected as a regression variable in both cases. In USN, a dependence on B_t appears to be also important, whereas LSN cases show a dependence on q_{95} . In practice, the difference between using B_t and q_{95} in these regressions is relatively small, but we report the result as

²The fact that $S_{\text{ioniz}}/\mathcal{P} \approx 1$ across a wide range of plasma conditions means that for every D atom introduced into a tokamak discharge approximately one photon at the Ly_α energy is emitted (but not necessarily detected).

obtained from model selection via the AIC in each case.




Figure 6-11: Regression of C-Mod deuterium flux through the LCFS for USN (unfavorable ion ∇B drift direction, open divertor) discharges, identifying dependencies on the midplane wall pressure, $p_{D2,mid}$, and toroidal magnetic field, B_t .




Figure 6-12: Regression of C-Mod deuterium flux through the LCFS for LSN (favorable ion ∇B drift direction, closed divertor) discharges, identifying dependencies on the midplane wall pressure, $p_{D2,mid}$, and q_{95} . The latter variable delivers a marginally smaller AIC than B_t , which is instead selected for USN cases (Fig. 6-11).

Comparing the two regressions in Figs. 6-11 and 6-12, we find a clearer trend (higher R^2) in LSN discharges than in USN. The dependence on $p_{D2,mid}$ in the two geometries is effectively the same, with an exponent of approximately 0.75. Comparable magnitudes of Γ_D are found in LSN and USN cases.

6.3.6 Effective Deuterium Transport Coefficients in the Pedestal

Once the deuterium flux through the LCFS is quantified via Ly_α data, effective transport coefficients may be easily computed. The effective deuterium diffusivity in the pedestal, $D_{D,eff}$, can be defined as the ratio of Γ_D and the absolute magnitude of the electron density gradient, $|\nabla n_e|$. Similarly, the effective deuterium convection, $v_{D,eff}$, is given by the ratio of Γ_D and n_e itself. For both these quantities, it may be desirable to use deuterium density, n_D , rather than electron density, n_e , but the former is not directly measured. We remark that our calculation and interpretation of deuterium transport coefficients relies on an assumption of main-chamber recycling, thoroughly demonstrated on Alcator C-Mod, since Ly_α signals at the midplane are taken to be representative of neutral densities at all poloidal angles.




Figure 6-13: Pedestal effective deuterium diffusivity, $D_{D,eff}$, as a function of electron density gradient scale length, L_{n_e} .

Fig. 6-13 shows values of $D_{D,eff}$ at $\rho_p = 0.99$ as a function of the local value of L_{n_e} , the electron density gradient scale length. The observed linear relationship indicates that the deuterium cross-field flux, Γ_D , varies only weakly as a function of density gradients at this location. Analogous observations in previous Alcator C-Mod research have been interpreted as evidence of “critical-gradient” behavior [237], whereby the density is “pinned” to a certain gradient dictated by transport effects. Fig. 6-13 shows this with a remarkably robust trend,

spanning many orders of magnitude. No obvious dependence on plasma current, I_p , shown by data point color, is observed.

Fig. 6-14a shows the same values of $D_{D,eff}$ over L_{n_e} as in Fig. 6-13, but now color-coded to separate L- (black) and H-mode (red) cases. Fig. 6-14b shows an analogous separation of confinement regimes, but now for the effective convection, $v_{D,eff}$, at $\rho_p = 0.99$ as a function of separatrix electron density. For both $D_{D,eff}$ and $v_{D,eff}$ we observe a clear clustering, with H-mode cases displaying much smaller transport coefficients. This agrees with the expectation that in H-mode a particle transport edge barrier forms in the pedestal via turbulence suppression. Future work will examine these transport coefficients in greater detail, attempting to address the question of how fueling and transport affect pedestal structure.




Figure 6-14: Effective pedestal transport coefficients for deuterium, inferred from $\text{Ly}\alpha$ measurements at the outboard midplane, for L-mode (black) and H-mode (red) cases. In (a), we show effective diffusivities as a function of electron density gradient scale lengths, whereas in (b) we show effective convection as a function of the separatrix density.

Previous work on Alcator C-Mod [238, 239] showed a correlation between $D_{D,eff}$ (or $v_{D,eff}$) in the SOL with local collisionality and/or line-averaged density normalized to the density limit, \bar{n}_e/n_G . The latter variable is a proxy for the Greenwald density limit, f_{GW} , defined in Section 6.3.1. In Ref. [239], LaBombard *et al.* showed that a parameter $\alpha_d \sim \lambda_{ei}/q^2 R$ (λ_{ei} being the mean free path for electron-ion collisions), suggested by Electro-Magnetic Fluid Drift Theory (EMFDT), better explains data in the SOL. This is an appealing physical explanation for the observed edge phenomenology in tokamaks and is obviously preferable to power-law regression analysis such as the one discussed here for the

separatrix. Future work will compare the present Ly_α database with theories of pedestal particle transport. The results presented here also form the basis of future comparisons of neutral and deuterium transport in C-Mod and DIII-D, where a similar midplane Ly_α diagnostic has recently been developed [240].

6.3.7 Summary of Ly_α database study

In previous sections, we have presented a new database of C-Mod plasmas comprising Ly_α data at the outboard midplane. This has allowed examination of atomic D neutral densities, n_n , and their relation to engineering and physics parameters, finding robust trends that support the established physical picture of main-chamber recycling on C-Mod. In Sections 6.3.3 and 6.3.4, we have examined the dependencies of pedestal neutral penetration depths inside and separatrix electron density on engineering parameters. In Section 6.3.5 we have then shown that the inference of deuterium cross-field particle fluxes, Γ_D , is robust to possible inaccuracies of kinetic profiles and atomic rates. This makes the comparison of Γ_D between experiment and theory a compelling task for future work. Figs. 6-11 and 6-12 showed some dependencies of Γ_D on wall pressure, B_t , and/or q_{95} , with small differences between USN and LSN in our database. Section 6.3.6 then presented clear evidence for “critical gradient” behavior of fluxes by examining effective transport coefficients. In Fig. 6-14, differences in both $D_{D,eff}$ and $v_{D,eff}$ between L- and H-mode cases have highlighted a strong effect of cross-field transport suppression at the L-H transition.

As discussed in Section 2.7, experimental particle transport research tends to focus on the quantification of transport coefficients to avoid the uncertainties involved in working with *fluxes*. Indeed, the concept of instability “fingerprints” [59], introduced by Kotschenreuther *et al.* to enable effective characterization of pedestal transport, does not assume that particle fluxes may be available from experimental analysis. The calculation of Γ_D discussed in Section 6.3.5 and the demonstration of its robustness to uncertainties in kinetic profiles and atomic rates offers new opportunities in this context. Future work will compare Γ_D profiles over the pedestal region to theory models, such as the neoclassical, gyrofluid and gyrokinetic codes discussed in Chapters 7 and 8.

6.4 Impact of Charge Exchange of Neutrals with Heavy Ions

Previous sections established the validity of the EIRENE neutral model within SOLPS-ITER for three Alcator C-Mod discharges. In Section 6.3 we then presented a much wider database of experimental neutral density measurements, identifying key dependences in Alcator C-Mod. Here, we demonstrate the impact of background neutrals on impurity transport making use of Aurora’s 1.5D forward model for ion transport.

Fig. 6-15 offers an illustration of the importance of background neutrals on charge state balance in the outer part of the plasma radius ($\rho_p > 0.8$), particularly focusing on the I-mode discharge of Fig. 6-3 (#1080416025). Using the SOLPS-ITER results for this case, Fig. 6-15 compares (a) charge state density profiles, and (b) line radiated power, for the highest charge states of Ar, introduced in the Aurora simulations assuming a constant source rate of 10^{20} particles/s. Realistic radial profiles of D and v were applied, with $D = 1$ m²/s and v being small up to the pedestal, where it sharply increases in magnitude to -100 m/s (similar to the profiles in Fig. 7-17). Details of these transport coefficients do not affect the clear observation of Fig. 6-15, where the effect of CX is seen to be significant in both n_z and P_{rad}^{line} . Clearly, any modeling of experimental impurity transport in the pedestal must take into consideration the impact of background neutrals in order to capture the ionization state of any ion. Spectroscopic measurements are also obviously heavily affected by the variation of line radiation from expectations that do not account for CX. Estimates of radiated power, particularly important for considerations on core-edge integration, are similarly impacted. We highlight that the illustrated example is from a C-Mod high-performance discharge, with a density that would be considered very high for any other tokamak. Fig. 6-15 should therefore be seen as a conservative demonstration of the importance of CX on charge state densities and radiation, which would likely be even greater on other devices.

Indeed, an analogous illustration for the DIII-D L-mode discharge (#180533) in Fig. 6-16 shows even more dramatic effects of charge exchange in this lower-density plasma. These Aurora simulations use the same parameters as those in Fig. 6-15 (including the same D and v), except for different neutral and plasma background profiles. One can hardly overstate the importance of including CX as part of the total effective recombination experienced by simulated impurities: charge state densities of the simulated Ar impurity are radically modified by the presence of background neutrals, taken from SOLPS-ITER modeling of this




Figure 6-15: Illustration of the impact of edge neutrals on Alcator C-Mod for identical Aurora simulations in an I-mode discharge (#1080416025) with a constant Ar source of 10^{20} particles/s and realistic D and v profiles. In (a) we show some charge states densities; (b) displays the corresponding line radiated power. Continuous lines indicate the result obtained including charge exchange with background neutrals; dashed lines give the result without charge exchange. Neutral densities are from SOLPS-ITER simulations.




Figure 6-16: Charge state densities and line radiated power for Aurora simulations with and without CX in the DIII-D L-mode discharge (#180533). Except for background neutral (from SOLPS-ITER) and plasma profiles, all parameters are the same as in Fig. 6-15. Continuous lines show the result with CX, dashed lines without it.




Figure 6-17: Radial profiles of (a) atomic D neutral density and (b) its normalization by n_e in the DIII-D diverted negative triangularity discharge #180533.

DIII-D discharge. Fig. 6-17 shows radial profiles (at the LFS, HFS and the FSA) of atomic D neutral density from SOLPS-ITER and its normalization by the electron density. Comparing these values to those from the C-Mod discharges in Fig. 6-3, a much larger n_n/n_e is found to stretch across a wider portion of this DIII-D plasma, clearly indicating that charge exchange must have a larger impact on ionization balance in typical DIII-D L-mode pedestals.

These observations highlight that impurity transport validation efforts aiming to elucidate pedestal physics must carefully quantify neutral densities. Recent work by Dux *et al.* [228] has shown effects of CX on pedestal impurities in AUG that are broadly in agreement with our results. The AUG studies are complementary to those presented here since they offer direct and detailed measurements of impurities, but resort to a simple 1D Monte Carlo model for edge neutrals. For our investigations, the opposite is true: the influence of CX on impurities is only indirectly observed via experimental diagnostics, but neutral densities are obtained via detailed experimental Ly_α measurements and/or advanced neutral modeling, both for Alcator C-Mod and DIII-D.

While SOLPS-ITER results have been directly compared to Alcator C-Mod Ly_α measurements (Section 6.2), a wider effort on DIII-D is underway through the recently-installed D Ly_α array near the device midplane (viewing both the low- and high-field sides). Since data from this DIII-D diagnostic were not available for the DIII-D experiments discussed in this thesis, we relied on results from the SOLPS-ITER and KN1D codes (Appendix F).

6.5 Charge Exchange with NBI Neutrals

In previous sections we considered the impact of CX of ions with thermal neutrals from the plasma edge. On tokamaks that make use of Neutral Beam Injection (NBI), beams constitute an additional source of neutrals in the plasma. Unlike edge neutrals, NBI neutrals have both a fast and a thermal population. The fast population can be separated into three components, typically labelled “full-”, “half-” and “third-” energy. This nomenclature reflects the fact that these neutral populations are created in the NBI beam duct by the application of a single (nominal) accelerating voltage. Such voltage acts on ionized D^+ atoms, D_2^+ molecules or D_3^+ molecules. Each of these is accelerated to the same energy, before being neutralized again so as to be able to penetrate across the tokamak magnetic field. For a fixed energy, each ion has a speed that depends on its mass. Once the molecular ions dissociate in either the neutralizer cell or in the plasma, the resulting individual D neutrals have an energy that corresponds to the original speed from the acceleration process, giving full, half and one third energies for the D, D_2 and D_3 populations, respectively. Upon entering the plasma, the NBI fast D neutrals can undergo CX reactions, causing recombination of D ions into thermal neutrals that can undergo CX once more. Multiple “generations” of CX can thus be produced. The first generation is affected by CX cross sections involving a fast neutral; all following generations only involve thermal particles. Thermal neutrals produced by the NBI are collectively called “halos”, since they surround the fast beam. To compute the spatial distribution of beam neutrals (both fast and thermal), we make use of the FIDASIM Monte Carlo code [241, 242]. FIDASIM is a state-of-the-art numerical tool that takes as inputs the plasma background (kinetic profiles, plasma geometry, etc.) and beam parameters, adopting a comprehensive collisional-radiative model to account for ionization and recombination processes of neutrals. On both DIII-D [177, 178] and AUG [173] FIDASIM is used to accurately model emission from ions that can undergo charge transfer with NBI neutrals.

In this section, we investigate the extent to which NBI neutrals may affect impurity ions, focusing on the DIII-D diverted negative triangularity discharge #180526 that will later be discussed in Chapter 8. Since F impurities were injected via LBO into this discharge, we shall focus on F ions as an example. Similar analysis has been completed for all the DIII-D discharges analyzed in this thesis, looking individually at all beams operating in each case. Fig. 6-18 shows the beam power sequence in shot #180526. Since the LBO injection




Figure 6-18: Time history of NBI power for the DIII-D discharge #180526. Short spikes indicate beam blips used for diagnostic purposes. The labelling of each time trace corresponds to standard nomenclature of DIII-D beams.

analyzed in Chapter 8 for this discharge is at 2.75 s, only neutrals from the beams labelled “30L” and “33L” need to be considered in this case. Fig. 6-19 shows the density of the full-energy component of the 30L beam neutral population over the poloidal cross section. The blue contours represent flux surfaces, with the thicker blue line identifying the LCFS. The 30L beam neutrals are seen to penetrate through the center of the plasma and into the HFS.

Based on Fig. 6-19, one may expect CX between impurities and beam neutrals to be mostly a phenomenon of edge relevance, since the neutral density is clearly higher in the edge. This is however not the case. As discussed in Section 6.4, CX reactions are effectively averaged out over flux surfaces and, since differential volumes between flux surfaces increase as one goes further out in radius, their relevance is greater in the core region. Fig. 6-20 shows the flux surface averaged (FSA) density of full-energy (fdens), half-energy (hdens) and third-energy (tdens) components, as well as the sum over all the halo generations, for the 30L beam in shot #180526. Results for the first 3 excited states of neutrals are displayed. We remark that the $n = 1$ halo density is approximately twice as large as the full-energy component density.

Fig. 6-21 shows the charge exchange recombination rates of an ion of charge $Z = 9$ with each of the beam neutral populations. To compute these, we used rates from Janev & Smith [227] for fast neutrals and the ADAS thermal charge exchange rates for halos (the same used for edge thermal neutrals). For impurity transport simulations, similar calculations are needed for all charge states.




Figure 6-19: Density of the full-energy fast neutral beam component incident from the right hand (low-field) side, displayed over the poloidal cross section of the DIII-D negative triangularity discharge #180526. Contours of the ρ_p coordinate identify selected flux surfaces, with the thick white line showing the LCFS.

Fig. 6-22a compares the total recombination rate from all NBI neutrals from beam 30L (fast and thermal) with those from radiative recombination (RR) and dielectronic recombination (DR). The latter two processes give much larger rates near the edge, but are of similar magnitude to the NBI CXR near the magnetic axis. Since both the 30L and 33L beams were active at the simulated time of interest in this discharge, CXR from both must be considered. For $q = 9$ (e.g. fully-stripped F), this brings the total CXR rate to approximately 10.8 s^{-1} in shot #180526 at 2.75 s, which is approximately a factor of 2 larger than the sum of RR and DR on axis.

However, it is important to compare these values with those of ionization rates, S , shown in Fig. 6-22b. Since core plasmas are in a strongly ionizing regime, the modification of total




Figure 6-20: FSA density components from FIDASIM for the DIII-D discharge of Fig. 6-19, separating full-, half- and third-energy components of fast neutrals and thermal (halo) contributions. Terms arising from the $n = 1$ (ground state) and excited states $n = 2$ and $n = 3$ are separated.




Figure 6-21: Radial profiles of recombination for F^{9+} ions due to CX with high energy and thermal (halo) neutrals from the DIII-D 30L beam in shot #180526. Contributions from excited states are found to be small.

recombination rates due to CX between impurities and NBI neutrals does not have a major impact in forward modeling. As we will discuss in Chapter 8, impurity transport in DIII-D may therefore, to a good approximation, be analyzed without considering the effects of NBI neutrals on the impurity charge state balance. High fidelity modeling like the one discussed in Chapter 8 must however examine NBI CXR on a case-by-case basis, since it hard to generalize how NBI CXR balances with other recombination rates for a given ion of interest.

In Chapter 8, we will also consider discharge #180520, with an LBO injection at 2.5 s, and #180530, with an injection at 3.8 s. In the former case, only the 30L beam was active, while in the latter case 30L, 33L and 33R were all operational.




Figure 6-22: Illustration of the total impact of CX between NBI neutrals and low-Z ions. In (a), we show radial profiles of total CX recombination of F^{3+} , F^{6+} and F^{9+} , summing over all fast and thermal components, ground and excited states 1-3, of beam 30L in the DIII-D discharge #180526. In (b), we compare radiative and dielectronic recombination (RR+DR), CXR with NBI neutrals (NBI CXR), and ionization (S) for F^{8+} . Ionization dominates everywhere in the core, making the impact of NBI CXR relatively small even though it can more than double the total near-axis recombination rate.

6.6 Discussion and Summary

In this chapter, we have presented a number of investigations related to neutral particles, both in the core and edge of the plasma. We have highlighted the critical importance of accurate neutral density measurements and/or predictions in the pedestal in order to correctly estimate ionization balance for ions in the plasma. The strong effects of CX shown in Figs. 6-15 and 6-16 leave no doubt that any investigation of particle transport in the pedestal must resolve the physics of edge neutrals. The impact of NBI neutrals, analyzed in detail using FIDASIM simulations in Section 6.5, has been shown to be much more limited, but may not be negligible in some circumstances.

The comparison of D Ly_α signals from the C-Mod midplane with SOLPS-ITER has demonstrated good agreement with experimental data. Based on this, it is reasonable to take SOLPS-ITER predictions to be valid at all poloidal angles, within the pedestal and in the near SOL. In Chapters 7 and 8 this will enable calculation of flux surface averaged neutral densities for the inclusion of charge exchange recombination (from both the edge and neutral beam injection) in forward modeling of impurity transport with Aurora.

The C-Mod edge neutral density database presented in Section 6.3 offers strong evidence of main-chamber recycling, in agreement with previous work [230, 243, 229, 237]. A wide

range of engineering and physical parameters has allowed the extraction of scaling laws that quantify key dependences of neutral density at the separatrix and neutral penetration depth over the pedestal. In particular, use of Eq. 6.8 as a formula for $n_{n,sep}$ at the LCFS and of Eq. 6.10 for L_{n_n} allow quantification of D neutral densities as a function of radius inside the confined plasma for C-Mod L-mode discharges. Accurate predictions of n_n over H-mode pedestals remain challenging. Future work will expand the database of H-mode discharges with high-quality Ly_α data, possibly also including discharges at a wider variety of B_t values and different ion ∇B drift direction.

In Section 6.3.4 we have also presented regressions of separatrix electron density, $n_{e,sep}$, which are of great interest for integrated modeling and to examine pedestal structure. The obtained scaling laws are only a first attempt at identifying key dependencies of $n_{e,sep}$, but already offer an interesting comparison to scaling laws for the top of the pedestal, for example those presented by Hughes *et al.* in Ref. [224]. In Section 6.3.5 we have also motivated the use of deuterium fluxes inferred from Ly_α signals as a promising new path for the validation of theoretical transport models in the pedestal. The availability of Γ_D profiles in a large number of C-Mod discharges and the demonstrated *robustness* of this measurement with respect to uncertainties in kinetic profiles and atomic rates make this a compelling avenue for future research.

In the next chapter, neutral density predictions from SOLPS-ITER are applied to improve understanding of atomic processes in the pedestal and near SOL in the context of Bayesian inferences of impurity transport for Alcator C-Mod. Analogous inferences, also including the effect of NBI neutrals, will be presented for DIII-D in Chapter 8.

Inference of Impurity Transport on Alcator C-Mod

This chapter presents inferences of radial profiles of cross-field impurity transport on Alcator C-Mod. We begin by describing observations of different particle transport in various regimes without ELMs. We then demonstrate D and v inferences following LBO injections of Ca with progressively more complex spectroscopic constraints. Results are compared to neoclassical, gyrofluid and gyrokinetic modeling, finding partial agreement in diffusion profiles, but significant discrepancies in convection.

Introduction

In this chapter, we explore cross-field impurity transport in 3 regimes on Alcator C-Mod: the low-confinement (L-) mode, the Enhanced D-Alpha (EDA) high-confinement (H-) mode [22], and the intermediate-confinement (I-) mode [23]. We first discuss observations of core electron density peaking and compare these to widely adopted H-mode scalings, finding that Alcator C-Mod data can exhibit significantly different behavior across confinement regimes. These observations serve as a useful starting point for the following sections, which describe inferences of impurity transport. The validation of impurity transport models aims to gain better understanding of strengths and weaknesses of theoretical transport models. We attempt to determine whether modeling tools dealing with coupled atomic and transport physics are sufficiently elaborate to predict the performance of future fusion

devices, including the profiles of fuel and impurity densities. Research presented in this chapter presents the most advanced attempt to date to compare C-Mod core particle transport coefficients with theoretical transport models.

7.1 Particle Transport Across Confinement Regimes

The decoupling of particle and heat transport channels in tokamak plasmas is perhaps nowhere as clear as in comparisons of Alcator C-Mod L-, H- and I-modes. When transitioning from L- to H-mode, a steep pedestal develops in both density and temperature. Due to the stiffness of core transport [73], which causes kinetic profile gradients to be almost fixed, the occurrence of higher pedestal-top values raises the entire core profiles, resulting in higher overall performance. On the other hand, in I-modes a pedestal forms in the temperature channel, but not in the density one. This results in particularly high temperatures across the plasma cross section, while maintaining low particle confinement. This is often considered favorable for reactor performance because it leads to low impurity retention. Moreover, gas fueling is simpler and more affordable than auxiliary heating.¹

Fig. 7-1 illustrates some important features that set I-modes apart from other high-performance regimes. In this figure, we show raw data from the gas-puffing Charge eXchange Recombination Spectroscopy (CXRS) diagnostic of C-Mod across an I-L transition, occurring at 1.3 seconds following a reduction of external heating. On the left, the measured boron density is observed to be unvaried across this transition. On the right, we see that, as the plasma transitions from I-mode to L-mode, the high ion temperature pedestal is lost. This highlights the existence of a heat transport barrier during I-modes, without a particle barrier. Unfortunately, these measurements were not available for the LBO injection studies explored in this thesis. Note that while Fig. 7-1 suggests that there may be no difference in impurity densities in the L- and I-mode phases, it would be incorrect to assume that this observation equally applies to all I-modes and to all impurity species.

The density and temperature differences in L-, I- and H-modes lead to different values of effective plasma collisionality, defined as [244]

$$\nu_{\text{eff}} = \frac{0.1 Z_{\text{eff}} \langle n_e \rangle R_{\text{geo}}}{\langle T_e \rangle^2}. \quad (7.1)$$

¹In practice, each confinement regime tends to have advantages and disadvantages, requiring detailed assessment and trade-offs that are not the focus of our discussion here.




Figure 7-1: Charge exchange recombination measurements from a D gas puffing experiment (shot #1120921014) on Alcator C-Mod, showing (intrinsic) boron density and temperature across an I-L transition occurring at 1.3 seconds.

Previous work on AUG and JET [244] investigated the ν_{eff} dependence of experimental H-mode electron density peaking, taken as $n_e(\rho_\psi = 0.2)/\langle n_e \rangle$, and found the following scaling to apply

$$\eta_{n_e}^{\text{Angioni}} = 1.347 - 0.117 \ln \nu_{\text{eff}} - 4.03\beta, \quad (7.2)$$

typically assumed to apply to D/T fuel density as well. Here, $\langle n_e \rangle$ is the volume-averaged density (in 10^{19} m^{-3}), $\langle T_e \rangle$ is the volume-averaged plasma temperature (in keV), Z_{eff} is the effective charge, and R_{geo} is the geometric plasma radius (in m). The plasma beta is defined as

$$\beta = \frac{4.02 \times 10^{-3} \langle p \rangle}{B_T^2} \quad (7.3)$$

with $\langle p \rangle$ being the volume-averaged plasma pressure (in $\text{keV} \times 10^{19} \text{ m}^{-3}$) and B_T the toroidal magnetic field on axis (in T).

The scaling above was later confirmed by Greenwald *et al.* [77] based on Alcator C-Mod low-collisionality H-modes data, leading to a prediction of density peaking in ITER in the range of 1.4-1.6. In Fig. 7-2 we revisit this scaling to compare it to C-Mod data across confinement regimes. The cases displayed here are from Alcator C-Mod experiments performed across a number of years, with no limitation in collisionality or any other parameter. The ordinate shows experimentally-measured values, the abscissa predictions from the Angioni scaling. The dashed black line shows a 1-to-1 correspondence between

the two. Clearly, experimental data significantly deviates from the scaling predictions, suggesting that either hidden variables are required to align data points, or it is unlikely to find a single regression that applies equally well across confinement regimes. We note that the Angioni scaling was constructed to make predictions in low-dimensionality H-modes, where ITER should operate at its highest performance. Hence, the observed discrepancies seen in Fig. 7-2 may not come as a surprise. We also remark that density peaking in L- and I-mode discharges is not significantly different, although collisionality is lower in I-mode due to the higher temperatures.




Figure 7-2: Scatter plot of electron density peaking $n_e(\psi_n = 0.2)/\langle n_e \rangle$ vs. the peaking prediction by the Angioni scaling for I-, L- and H-mode shots on Alcator C-Mod. Black curves correspond to clustering of the three confinement regimes by a Support Vector Machine using a 3rd order polynomial kernel. The dashed line corresponds to a 1-to-1 correspondence of data with the Angioni scaling, from which the experimental distributions clearly differ substantially.

Fig. 7-2 highlights that density peaking, and particle transport more broadly, display a wealth of collective behavior that is hard to capture via simple scaling laws. Understanding such behavior via detailed physical modeling is paramount in order to reliably predict the performance of future fusion devices. Doing so for impurities, rather than electrons, is even more challenging, since atomic physics comes into play, carrying additional uncertainties.

As discussed in Chapters 2 and 3, the coupling of atomic and transport physics makes the validation of heavy ion transport models a challenging, but rewarding task, since impurities can be introduced into the plasma as trace species and diagnosed individually by spectroscopic means. In the following sections, we describe how Bayesian inferences of radial impurity transport coefficients have been developed within this thesis to provide detailed comparisons to theoretical transport models.

7.2 Inferences of Impurity Transport on C-Mod

7.2.1 Review of Previous Efforts

Inferences of impurity transport have been a subject of research on Alcator C-Mod for a number of years [245, 181, 101, 102, 246]. Pedersen *et al.* [100] attempted to infer D and v values in the pedestal of EDA H-mode discharges using SXR radiation filtered by a thin ($10\ \mu\text{m}$ Be) foil. These measurements were interpreted as deriving from intrinsic fluorine present in the device and suggested the presence of a neoclassical-like edge inward pinch based on variation of D and v parameters over the same region. Howard later developed a laser blow-off (LBO) system [153] to better distinguish the effects of D and v through a time-dependent source of non-intrinsic, non-recycling impurities. Making use of the X-ray Imaging Crystal Spectrometer (XICS) and the Vacuum Ultra-Violet (VUV) diagnostics also employed in this thesis, the inference of core D and v profiles was attempted using the Levenberg-Marquard algorithm to then compare results to gyrokinetic simulations [102]. These studies were limited to L-mode discharges and only made use of the brightest He-like Ca line measured by the XICS spectrum for CaF_2 injections (the Ca resonance line w , see Chapter 3). Some years later, Chilenski analyzed the sensitivity of D and v inferences with the same setup, using Bayesian inference methods with synthetic data [27]. Significant progress was made on the statistical rigor of the approach, leveraging nested sampling for parameter estimation and model selection. It was suggested however that application of similar methods to experimental data was unlikely to succeed, due to the “ill-posed” nature of the problem, where “seemingly small residuals in the synthetic diagnostic signals can mask gross disagreement with the true D and v profiles”. It was then concluded that “a principled approach to model selection can overcome this issue, though only at great computational expense. Because of these results, [...] D and v profiles may not be the best level of the

primacy hierarchy at which to perform model validation.” [184]. This was effectively the starting point of this thesis.

In this chapter, we apply the Bayesian inference methods and the Aurora forward model described in Chapter 5 to infer radial profiles of impurity transport coefficients and compare them to results from neoclassical, gyrofluid, and gyrokinetic modeling. The discharges discussed here were run before the author began this work and presented the same experimental setup as used by Howard and Chilenski. Issues highlighted in previous work have been faced in a number of ways, with significant improvements in both the forward and inverse modeling:

- Aurora has been developed for efficient forward modeling of particle transport, increasing the physics fidelity of experimental modeling with respect to previous C-Mod work (see Section 5.3.1 and Appendix A).
- We have extended the use of nested sampling with more appropriate likelihood (Section 5.3) and prior (Section 5.4) distributions.
- Most importantly, we have leveraged available spectroscopic data to a larger extent than previous research, making use of the entire Ca K_α spectrum rather than just the resonance (w) line, as well as multi-line spectroscopy in the VUV range.

As we will describe in the following sections, these developments overcame many of the previous challenges, opening new paths to transport model validation. Computational bottlenecks have been completely eliminated, thanks to the development of a high-performance computing framework running Aurora. In Chapter 8 we discuss how similar methods were leveraged for DIII-D research.

7.2.2 Experimental Setup

An initial condition for particle transport inferences is provided by an optical fiber viewing the edge of the plasma, near the location of the LBO injection. A bandpass filter limited measured brightness to the 420 ± 10 nm range, providing a proxy measurement of Ca-I emission [102]. This system offers high signal-to-noise ratio (SNR) and sometimes provides evidence of CaF₂ clusters entering the plasma; this often correlates with irregularities in signals on several spectrometers.

For our inferences, a central role is played by C-Mod’s XICS diagnostic [175], which has 3 ms readout time and typical signal integration over 3 ms intervals, for a total time resolution of 6 ms. As described in Section 4.4, the two crystals of this system are normally configured to view 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_ϕ) following tomographic inversion [175, 247]. 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 [106]. On the other hand, v_ϕ profiles are only constrained by XICS data, thus limiting detailed gyrokinetic analysis to the core.

For impurity transport studies, one of the two XICS crystals was substituted with one capable of viewing the Ca^{18+} (He-like) spectrum, typically discretizing the detector 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 mÅ, and past inferences of impurity transport coefficients in Refs. [102, 248, 103] relied mostly on its measurement using existing spectral analysis routines [180].

Finally, the XEUS single-chord VUV spectrometer (Section 4.2), covering the 10-70 Å wavelength range, can be used to view multiple emission lines for Li-like Ca [164]. In the last few years of C-Mod operation, a companion Long-wavelength Extreme Ultraviolet Spectrometer (LoWEUS) spectrometer [165] has allowed emission in the 100-300 Å region to also be analyzed, but these data were not available for the discharges discussed here. Their main use during this thesis work has been in allowing comparison to analogous DIII-D measurements. The spectral range of these spectrometers typically offers valuable observations of Ca^{16+} and Ca^{17+} at $r/a > 0.7$, thus providing a constraint to transport modeling in the near-pedestal region. Unfortunately, the relatively low time resolution (2 ms) of XEUS and LoWEUS does not usually allow clear observation of the signal rise phase following LBO injections.

In this chapter, we describe inferences obtained for 3 discharges: an L-mode (#1101014006), an EDA H-mode (#1101014019) and an I-mode (#1101014030). These discharges are from the same set of C-Mod experiments, but differ significantly in the parameter space that they explore: whereas they all have $B_t \approx 5.5 T$, the L-mode case has $I_p \approx 0.8 MA$, the EDA H-mode $I_p \approx 0.55 MA$, and the I-mode $I_p \approx 1.0 MA$.

Using the neoclassical prediction of Sauter *et al.* [249], Z_{eff} has been estimated to be approximately 1.75, 1.5 and 2.1, respectively, with estimated uncertainty of ± 0.2 in each case, based on the measured loop (toroidal) voltage and the plasma current. These values of Z_{eff} are assumed to be equally applicable across the entire radius, consistent with past visible bremsstrahlung measurements on C-Mod [250]. This leads to a volume-averaged (dimensionless) collisionality, defined as $\nu_{ei} = 0.1 Z_{\text{eff}} \langle n_e \rangle R / \langle T_e \rangle^2$ [244], of 1.5, 1.9 and 0.9 ± 0.2 for the L-, H- and I-mode cases, respectively.




Figure 7-3: Time histories 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 chosen Alcator C-Mod L-mode (green), EDA H-mode (orange) and I-mode (indigo) discharges described in the main text.

In Fig. 7-3, we show time traces of H-minority Ion Cyclotron Resonance Heating (ICRH) power (P_{RF}) with a near-axis resonance, central electron temperature ($T_{e,0}$), normalized core radiated power from bolometry (P_{rad}), and line-averaged density from a Two-Color Interferometer [158] (\bar{n}_e) for the L-mode (green), EDA H-mode (orange) and I-mode (indigo) discharges. Vertical dashed lines in the top panel show the time up to which impurity transport has been simulated for each LBO injection. The I-mode case is seen to have

a sudden drop in RF power approximately 65 ms after the LBO injection, but this is appropriately excluded from the time range of our modeling. As it is often the case in C-Mod, sawteeth modulate core temperatures in all the shots of interest, as shown in the second panel.

Two important observations can be made from Fig. 7-3. First, although all plasma discharges have core MHD activity (sawteeth) modulating T_e near axis, the density and power are constant within the time ranges of interest, i.e. the discharges are satisfactorily described as *quasi-steady*. Secondly, the LBO injections into these discharges appear to be *non-perturbative*, as suggested by the lack of changes in line-averaged density in the fourth panel when the laser ablates. We take this to justify a *trace impurity* approximation, since the injected impurities do not appear to affect background kinetic profiles.




Figure 7-4: Time-dependent kinetic profile fits for the investigated I-mode discharge.

As described in Section 4.1, 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 parametric time-dependent Radial Basis Functions (RBF), as in the DIII-D research discussed in Chapter 8, taking care to avoid smoothing of sawteeth and to preserve pedestal structure. Use of fast Electron Cyclotron Emission (ECE) data allows us to fit very well the T_e modulation, while for n_e we rely on Thomson Scattering data. Fig. 7-4 shows the RBF fits obtained for the I-mode discharge. Similarly to the other two shots, we observe that n_e in Fig. 7-4 is not significantly modulated by sawteeth, to the point that Thomson Scattering cannot quantify changes with its 100 Hz

sampling rate. On the other hand, variations of near-axis T_e can significantly affect photon emissivities from ions near the magnetic axis.




Figure 7-5: Kinetic profiles and normalized gradient scale lengths for the 3 Alcator C-Mod discharges where impurity transport has been inferred.

For theoretical transport modeling, we adopt Gaussian Process Regression (GPR), which minimizes user bias by inferring hyperparameters corresponding to fits of highest probability [154]. In Fig. 7-5, 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 [106]. We quantify GPR uncertainties by propagating diagnostic uncertainties and data scattering in time via the Law of Total Variance and MCMC sampling [163], using a Gibbs kernel with tanh length scale, as suggested in

Ref. [154]. This is a relatively expensive procedure, particularly valuable to estimate uncertainties in gradients, obtained in GPR via simple analytical derivatives [251, 163]. The uncertainty bands shown in each panel of Fig. 7-5 correspond to the 1-99, 10-90, and 25-75 quantiles of the local posterior distributions.

The profiles shown in Fig. 7-5 display typical features of each confinement regime. The L-mode discharge does not show any form of pedestal, either in density or temperature, consistent with its labelling of “low confinement”. The EDA H-mode has a pedestal in both density and temperature (“high confinement”). The I-mode has a strong temperature pedestal, comparable to the EDA H-mode in Fig. 7-5, but a smooth variation of density near the pedestal region (“intermediate confinement”). None of these discharges exhibits Edge-Localized Modes (ELMs), typical of H-mode operation on other devices (and, in specific cases, also on C-Mod). The I-mode case has a weakly-coherent $\sim 100 - 300$ kHz mode near the LCFS, high levels of particle transport and low impurity retention [23, 252]. A Quasi-Coherent Mode (QCM) at ~ 100 kHz appears to play a similar role in the EDA H-mode [22]. We remark that any comparison among the three shots of Fig. 7-5 must necessarily account for the fact that they have different values of plasma current and therefore plasma confinement should not be expected to be the same. As for the validation of SOLPS-ITER in Section 6.2, these studies on impurity transport represent “spot checks” for transport model validation.

Ideally, one would use the same time-dependent kinetic profiles for both experimental and theoretical modeling to provide an accurate validation of the latter. Unfortunately, performing theoretical modeling for each time slice would be a computationally- and time-intensive procedure. We instead make use of time-averaged kinetic profiles such as those in Fig. 7-5, and describe the effects of varying inputs to theoretical models within uncertainties derived from both measurement accuracy and data point scattering over time. One may also wonder whether it would not be most appropriate to then compare time-independent theoretical modeling to similarly time-independent experimental modeling. The reason for not doing this lies in the observation that inaccuracies of experimental modeling at a certain spatio-temporal location can affect the entire spatio-temporal domain. For example, omission of the T_e modulation by sawteeth near the magnetic axis can lead to inference of different D or v at midradius to compensate for the missing fast transport phenomenology of sawteeth. This is to be contrasted with theoretical modeling, which may be meaningfully

treated as time-independent outside of the sawtooth mixing radius (in the absence of ELMs).

The L-, H- and I-mode discharges discussed here are very similar to those discussed in Section 6.2 in the context of a comparison between Ly_α spectroscopy and SOLPS-ITER modeling. Indeed, it was by searching for C-Mod discharges with identical engineering parameters to those of Fig. 7-5 that the shots of Fig. 6-3 were identified. Inferences of impurity transport presented in this chapter can therefore use the SOLPS-ITER flux-surface averaged (FSA) atomic D neutral densities as a background for Aurora simulations, permitting inclusion of charge exchange recombination as part of our forward modeling. This, however, has only become possible in the later phases of this work. Results in Sections 7.2.3 and 7.2.4 do not include charge exchange, whereas this is an important consideration in the most complete inferences presented in Section 7.2.5.

In the next section, we discuss experimental inferences for the I-mode discharge presented above using individual line brightnesses of the resonance (w) and forbidden (z) lines of the Ca K_α spectrum. This was the first demonstration of C-Mod experimental inferences combining more than one XICS spectral line, obtained using a high-performance computing framework and the physical fidelity permitted by Aurora. These results were presented in Ref. [155].² At the end of this section, we also describe the effect of including additional line brightnesses, using synthetic data. This analysis, first presented in Ref. [156], motivates moving to inferences that forward model the entire K_α spectrum, rather than individual line brightness values. This more complete approach is explored in Section 7.2.5 and has recently been submitted for publication [253].

7.2.3 Inferences Using Individual Line Brightnesses

The inference of D and v profiles is strongly dependent on the quality of experimental data, as well as the accuracy and flexibility of the applied “models”. In this section, we first describe spectral fitting methods that were developed as part of this work to fit high-resolution spectral lines from XICS measurements, presented in Refs. [254, 156]. These techniques allow the maximization of accuracy and uncertainty quantification of individual line brightnesses used in inferences of impurity transport. We then illustrate the results of model selection (Section 7.2.3) and parameter estimation (Section 7.2.3) for the I-mode

²In Ref. [155], the modeling tools that later evolved into Aurora were labelled as “pySTRAHL”, since they evolved from STRAHL and presented a Python interface. For simplicity, we simply refer to these codes as Aurora in this section, since this was effectively a first (pre-release) version of Aurora.

discharge previously described, using Ca w and z line brightnesses.

The Bayesian Spectral Fitting Code

Previous impurity transport research on Alcator C-Mod [26, 102, 246] made use of spectral fitting of XICS brightnesses using the THACO [180] suite. THACO covers a number of XICS analysis needs, particularly aiming to effectively implement Doppler tomography [175], used to obtain radially-localized profiles of ion temperature (T_i) and toroidal velocity (v_ϕ). When fitting spectral lines, THACO approximates each line shape as a single Gaussian, assuming dominant Doppler broadening. The Gaussian parametrization however inevitably runs into inaccuracies: even if local spectra were perfectly Gaussian, their line integral can never be exactly Gaussian. A more flexible, efficient, and statistically rigorous spectral fitting method has therefore been developed for this work. The resulting Bayesian Spectral Fitting Code (BSFC), first presented in Ref. [254], is a key component to the XICS spectral analysis presented in this section.

In BSFC, we use a Gauss-Hermite polynomial decomposition to extract moments from line spectra without needing an exact parameterization for the line shape. This flexible truncated expansion is combined with Markov Chain Monte Carlo (MCMC) and nested sampling to perform Bayesian uncertainty analysis on the estimated moments. Fig. 7-6 shows a BSFC fit of the H-like Ca Ly_α lines and nearby satellites, discussed in detail by Rice *et al.* in Ref. [148]. These lines have not been used in inferences of Ca transport and are just shown here for illustrative purposes.




Figure 7-6: BSFC fit of the Ca Ly_α lines and overlapping satellites, each in a different color. In this case, BSFC does not find significant evidence for background Mo and Ar lines, labelled 10d and w12n, respectively. Residuals are displayed in the lower panel. Rest wavelengths are shown by vertical dashed lines.

Details of BSFC methods are described in Appendix B, where we show that BSFC performs well on both synthetic and experimental spectra. Comparing it to THACO, we find significant improvements in both reliability of parameter estimation and uncertainty quantification. A Python implementation of BSFC is publicly available at <https://github.com/Maplenormandy/bsfc>.

Model Selection

As described in Chapter 5, the Bayesian evidence is the central metric in Bayesian inference to assess what level of model complexity is best supported by the 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. [188] for a detailed discussion). In Fig. 7-7, 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.




Figure 7-7: Scaling of the log of the Bayes Factor (BF) with number of free D and 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.

Fig. 7-7 also compares the BF obtained when we set v to scale linearly with Z in the pedestal (taken as $r/a > 0.9$) (blue) and when we set v to be independent of Z (green) in Aurora. 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.

Uncertainties in the 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.

Parameter Estimation

Fig. 7-8 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. 7-7), with 5 (blue), 6 (green) and 7 (red) radial coefficients for D and v . All of these cases use $v(r/a > 0.9) \sim Z$, since this choice is found to give higher BF. In Fig. 7-8 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 (≈ -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, $\tau_{imp} = 30 \pm 2$ ms, while the energy confinement time $\tau_e = 23 \pm 1$ ms.

The vertical magenta dashed line shows the approximate location of the sawteeth inversion radius. We remark that by taking the w and z line signals to be relatively calibrated, we make their fitting significantly more challenging: virtually no discrepancy between simulated and experimental signals would be seen if we allowed lines to be independently normalized.³ The uncertainties shown in Fig. 7-8 are not assumed to be

³This effect is **not** due to inaccurate atomic rates, but to true physical complexity of the experimental data.




Figure 7-8: Inferred profiles of D and v , together with the Rv/D ratio, for I-mode 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 dashed magenta lines show the approximate location of the sawtooth inversion radius. Figure from Ref. [155].

Gaussian and represent the 1-99, 10-90, and 25-75 quantiles of the posterior distribution. The Bayesian weight factors described in Section 5.4.3 (see also Appendix C), with $\nu = 25$, are found to effectively weigh XICS signals (for both the w and z lines) more heavily than VUV lines, with effective weights always less than 1. This suggests that the assumed diagnostic uncertainties are likely too small to have full consistency with the Aurora model. The use of Bayesian weight factors results in an increase of error bars for the profiles in Fig. 7-8 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.

Fig. 7-9 compares the inferred transport coefficients corresponding to the highest BF with the results of selected theoretical transport models. In blue, we show predictions from the NEO [32, 34] neoclassical transport code and in red from the Trapped Gyro-Landau Fluid (TGLF) quasilinear turbulence model [255, 45], both introduced in Chapter 2. In the bottom




Figure 7-9: Core-focused view of inferred profiles of D (on a logarithmic scale) and v , together with the Rv/D ratio, for the model-selected I-mode inference (also shown in Fig. 7-8). 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 Rv/D .

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. $Rv/D_{tot} = R(v_{neo} + v_{turb})/(D_{neo} + D_{turb})$). Scales have been adapted in each panel to complement the visualization of Fig. 7-8 and allow better comparison to transport models.

Since neoclassical predictions are formally limited by near-axis potato orbits and by ion orbit losses near the LCFS [34, 29], we limit our NEO simulations to the range $0.1 \lesssim \rho_\phi \lesssim 0.95$. Fig. 7-9 shows that inferred diffusion values near the magnetic axis are larger, but of the same order of magnitude, as the NEO predictions (obtained for Ca^{18+}). The inferred values are, of course, to some extent dependent on the applied sawtooth model in Aurora simulations, but we have found that D and v profiles of neoclassical character are always inferred inside the sawtooth inversion radius. Neoclassical transport modeling is discussed further in Section 7.3.

At mid-radius, where the effect of sawteeth is reduced, TGLF SAT-1 appears to underestimate both D and v amplitudes, while approximately matching v/D within

experimental uncertainties. Fig. 7-9 also shows a single-point prediction for D and v , including their ratio v/D from a nonlinear, ion-scale CGYRO [37] simulation at $r/a = 0.6$, run at a value of normalized ion temperature gradient scale length, a/L_{T_i} , that was modified with respect to the experimental mean estimate until matching the experimental ion heat flux, Q_i , computed via the TRANSP [256] code. Both TGLF and CGYRO modeling are described in greater detail in Section 7.3.

In Fig. 7-10a we show the self-similar impurity density profiles for the highest ionization stages corresponding to the inferred transport coefficients in Fig. 7-9, 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, well after the LBO injection. Density uncertainties were computed from the inferred D and v uncertainties, i.e. by propagating D and v uncertainties using chains and weights from MultiNest. This gives a clear visualization of how uncertainties in Fig. 7-9 translate to n_z profiles, retaining all degrees of correlation between parameters.

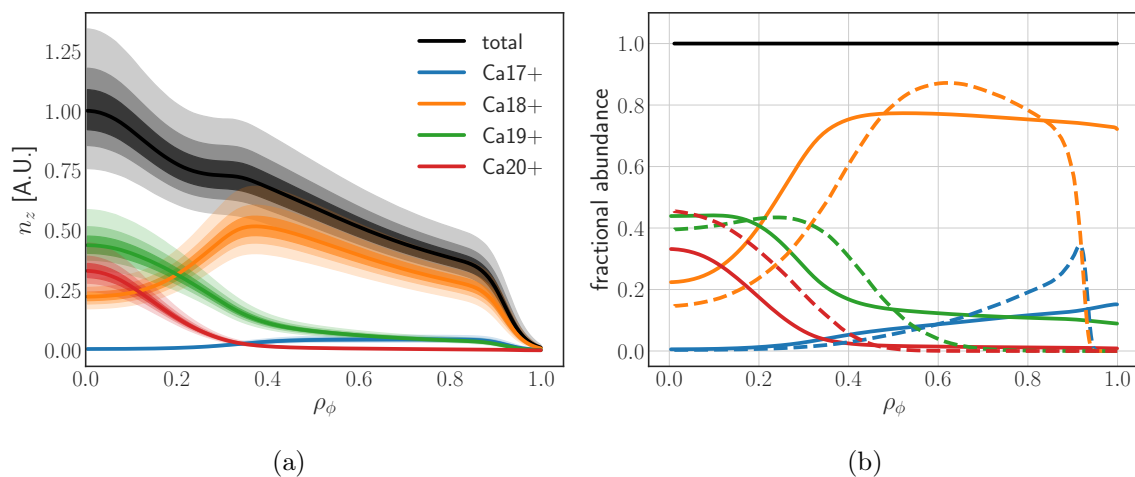


Figure 7-10: (a) Experimentally-inferred self-decaying profile shapes for the highest ionization stages of Ca after an LBO injection into the I-mode discharge. In (b), we compare the median fractional abundances of (a), shown with continuous lines, with those predicted by ionization equilibrium only (no transport), shown with dashed lines. Figure from Ref. [155].

In Fig. 7-10b 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 alone 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 immediately following the LBO injection, since impurity ions must first penetrate into the core to undergo nearly-complete ionization and then be transported outwards; at this stage, the upper level of x-ray transitions can be significantly populated via recombination, as previously described by Rice *et al.* [144]. Such recombination effect contributes to the He-like Ca z line emissivity measured via XICS, but not significantly to the corresponding w line, as described in Section 3.5.1.

The D and v profiles in Fig. 7-9, inferred using Ca brightness measurements, were also applied to predict steady-state profiles of Ar, assuming transport of the two species to be similar. Such predictions can be compared with the 72-chord XICS measurement of He-like Ar brightness, which goes from the bottom to the top of the confined plasma. Except for an evident Ar density excess in the direction opposite to the ion ∇B drift, previously described in C-Mod [98, 107], the Aurora predictions of He-like Ar brightnesses were found to be in general agreement with experimental measurements. However, this does not appear to be a significant validation constraint in the absence of absolute detector calibration, given that variations of D and v coefficients by a factor of 2 don't affect significantly the normalized comparison within error bars.

7.2.4 Constraints from Line Ratios

The inferences discussed in Section 7.2.3 (Ref. [155]) make use of both the w and z line brightnesses of the K_α spectrum. One may easily conjecture, based on atomic physics, that the combination of these two lines constrains transport coefficients better than either line alone. In Ref. [156], multi-line constraints have been discussed in greater detail. In this work, we have explored the opportunity to infer transport coefficients based on time histories of w, x, y, z, k, and j lines, whose wavelengths span the K_α spectrum, described in Section 3.5.1.

To demonstrate the potential of combining multiple lines and considering their brightness ratios, we choose to work with synthetic data that was generated using the NEO and TGLF D and v profiles shown in Fig. 7-9 for this discharge. Transport coefficients from NEO are taken to be relevant near the magnetic axis and in the pedestal, while TGLF results are assumed to be valid in between. D and v profiles have been merged near the

sawtooth inversion radius, $\rho_p \approx 0.4$, and at the top of the pedestal, $\rho_p \approx 0.92$, extrapolating linearly near the plasma center. Here, the radial coordinate ρ_p stands for the square-root of normalized poloidal flux.

In Fig. 7-11 we show radial profiles of the w, z, x, y, and k+j line emissivities that result from charge state densities at ionization equilibrium (i.e. no impurity transport, continuous lines) with those obtained from a forward model evaluation using the NEO-TGLF predictions for D and v (dashed lines). Both sets of predictions refer to 80 ms after the LBO injection, when impurity densities are decaying exponentially while maintaining constant profile shapes (the “self-similarity” property discussed towards the end of Section 7.2.3). The forward model results were normalized such that the He-like Ca density computed by Aurora was the same as the fractional abundance of that charge state at ionization equilibrium. Fig. 7-11 shows only the sum of k and j, rather than their individual emissivities, since their amplitudes are always at a fixed ratio due to their common atomic physics. Fig. 7-11 demonstrates that impurity transport affects the amplitude and shape of all emissivity profiles. These differences are particularly strong in the pedestal region, where He-like ions would not be expected to exist at ionization equilibrium due to the low local electron temperatures; it is only because of transport effects that these ions can emit in this region.

In Ref. [156] we have shown how the addition of each line from the set in Fig. 7-11 helps in reconstructing D and v profiles in synthetic inferences. As expected, uncertainties in inferred results are found to decrease as more data is added. It must however be noted that applying these multi-line inferences in experimental settings is not trivial because of possible line overlaps, particularly between the intercombination (x and y) lines and the satellites (m, s, and t) that exist in the same spectral region (see Fig. 3-4a). This observation strongly suggests that, in order to make the most of multi-line constraints from K_α spectra, one needs to completely abandon spectral fitting, i.e. the attempt at separating lines via an analytical description (in BSFC’s case, using Gauss-Hermite functions), and move towards an approach that forward models the entire spectral range, including all its components.

It is worth highlighting some key differences between these two approaches. In both cases, detailed atomic rates for the lines of interest are needed. However, forward modeling of the entire spectrum ($3.170 \text{ \AA} < \lambda < 3.215 \text{ \AA}$) also requires rates for all other lines present in the spectrum, for example the many satellite lines that BSFC effectively subtracts from the spectral features of interest (e.g. $n \geq 3$ satellites, see Fig. 3-4a). Moreover, this




Figure 7-11: Radial profiles of the w, z, x, y, and k+j relative line emissivities for the I-mode discharge. Continuous lines show predictions from fractional abundances at ionization equilibrium; dashed lines correspond to transport modeling with Aurora. Figure reproduced from Ref. [156].

approach requires assumptions on broadening mechanisms, whereas the computation of spectral moments with BSFC sidesteps this complication. Clearly, it is desirable to abandon the spectral fitting altogether, if possible, and have a model that naturally reproduces the entire broadened spectrum based on complete knowledge of the appropriate rates. The next section demonstrates this approach for inferences of particle transport in tokamaks for the first time.

One additional use of line ratios to constrain impurity transport arises from VUV spectroscopy. The XEUS detector was not absolutely calibrated, but the sensitivity across the detector is expected to be uniform. Previous inferences on C-Mod included multiple VUV lines, but never made use of their relative amplitudes on the detector. In other words, the time history of each line was normalized independently of others, since atomic data were not trusted to a sufficient extent to enable a common normalization. Recently, the release by ADAS of high-quality R-matrix atomic cross sections has made it possible to compute Photon Emissivity Coefficients (PECs) for all the Ca charge states that emit in the spectral range covered by XEUS. While the new cross sections do not seem to be able to match all measured lines, we were able to determine 6 subsets of lines, shown in Table 7.1, that

appear to give emissivities of the right amplitude. These are *groups* of lines, rather than individual lines, deriving from multiple atomic transitions whose wavelengths fall within chosen ranges; spectral resolution of the diagnostic does not permit their separation. It is of course impossible to tell exactly what is the accuracy of these atomic data based on our observations, but we estimate uncertainties to be less than 30% for each line. Table 7.1 specifies the Ca charge states that provide the strongest contributions within each wavelength bin. Charge states lower than Be-like Ca have also been included in our modeling, but they appear to be negligible for the overall signal envelopes.

Label	λ [Å]	Main Contributions
XEUS-0	186.0-188.0	2 Li-like lines
XEUS-1	195.0-197.0	2 Be-like lines, 3 Li-like lines
XEUS-2	197.2-199.0	1 Be-like line, 5 Li-like lines
XEUS-3	202.5-205.2	6 Be-like lines, 2 Li-like lines
XEUS-4	211.0-213.0	3 Be-like lines
XEUS-5	220.5-222.5	1 Be-like line

Table 7.1: VUV line identification used for the analysis of LBO injections of Ca on Alcator C-Mod.

By considering the relative amplitude of VUV lines, i.e. their line ratios, we obtain a stronger constraint on plasma transport in the near-pedestal region than if each line were independently normalized as in previous work (e.g. Ref. [102, 155]). This is particularly true when considering lines that originate from different charge states. In particular, the line labelled “XEUS-0” in Table 7.1 is dominated by emission from Li-like ions, while XEUS-4 and XEUS-5 are entirely from Be-like ions (not directly constrained by the K_α spectrum measured by XICS). The novel use of VUV line ratios enables one to investigate transport in the outer part of the plasma, where these emit most strongly, and leads to unprecedented detail in the inferences of impurity transport presented in the next section.

7.2.5 Inferences using the full K_α spectrum

To make forward modeling of the full K_α spectrum possible within an iterative framework, a significant level of code optimization is required. This involves caching of local Doppler-broadened spectra for all lines in the XICS spectral range, evaluated over wavelength, time, and radial location. During each evaluation of the forward model, the Aurora results for charge state densities are interpolated onto a sufficiently accurate but minimally detailed

set of grids for each dimension of these arrays. Local spectra are then averaged over the time integration bins of the diagnostics and they are line integrated along each XICS line of sight via quadrature weights. In order to reduce the runtime per forward model evaluation to approximately 300 ms, efficient Einstein summation routines in the *numpy* library [257] and superstaging of low charge states (see Section 5.3.1) were adopted. Doppler shifts for each XICS chord and each recorded time were directly extracted from line-integrated BSFC estimates of ion rotation velocity on the brightest (w) line. This inherently assumes that all lines within the spectrum are subject to the same Doppler shift, which is not strictly true for two reasons:

- Doppler shifts are linearly proportional to wavelength. This is negligible in our case, giving errors of less than 1 fm, since the spectral range of interest is rather narrow.
- Different lines may emit differently as a function of radius and thus be subjected to different toroidal rotation (which is taken to be a flux function and therefore only a function of radius). For our application, the brightest lines (w, x, y, and z) all emit from the same approximate regions, so their Doppler shift is expected to be similar. Satellite lines depend on Li-like ions, which are further out in radius and may therefore be subjected to different rotation. This does introduce some error, but certainly within the uncertainties of these lines.

This implementation enables efficient iterations, but makes it impossible to explore corrections to T_i from tomographically inverted Ar measurements using the other XICS detectors. While we expect that fixing n_e , T_e , and T_i in our inferences does not cause significant problems, our approach makes it difficult to assess the impact of uncertainties in these 1D profiles appropriately. Unlike in work described thus far, inferences described in this section also make use of edge D atomic neutral density predictions from SOLPS-ITER. We make use of the EIRENE results presented in Chapter 6 for similar discharges, with identical engineering parameters, for which experimental Ly_α signals at the midplane could be compared to modeling (Section 6.2). Using these neutral densities, we can include the effect of charge exchange (CX) in forward modeling, shown in Section 6.4 to be extremely important for accurate near-edge modeling. In the inferences described next, the SOLPS-ITER neutral density was given as an input to Aurora and an additional free parameter was included in nested sampling as a linear rescaling of the entire neutral density (n_n).

This parameter was set to have a log-normal prior with mean of 1 and $\sigma = 0.25$, enabling exploration of scale factors by approximately a factor of 2. This reflects the uncertainty that was found via the comparison of experimental D Ly $_{\alpha}$ measurements at the midplane with SOLPS-ITER in Section 6.2 (see Fig. 6-3).

In the next sections, we first demonstrate the efficacy of this inference framework using full K $_{\alpha}$ spectra on synthetic data, and then proceed to the final experimental results on this subject.

Inferences from Synthetic K $_{\alpha}$ Spectra

In this section, we verify the ability of the most complete modeling workflow presented in this thesis to infer D and v that were used to create synthetic data. These tests make use of the entire Ca K $_{\alpha}$ spectrum and relative amplitudes of VUV lines, as well as charge exchange with background neutrals, as discussed in the previous section.

In Fig. 7-13, we show the result of running synthetic inferences in the EDA H-mode case described in previous sections (shot #1101014019). Inferences with 1, 2, 3, 4, 5, and 6 free spline coefficients for each of D and v/D are shown. Black dashed lines show the input (“ground truth”) transport coefficients that were used to create synthetic data and that should therefore be recovered by successful inferences. These were obtained, for illustrative purposes, by blending the D and v obtained from NEO and TGLF for this discharge. As described in Section 5.4, an additional Gaussian feature representing the inward pinch in the pedestal region was also always included, with amplitude and width being free parameters in all cases.

While the inferences with 1 and 2 coefficients are visibly poor, due to their lack of flexibility to recover the detail of the “ground truth” D and v , higher-dimensionality inferences are comparably successful. Uncertainties in the inferred profiles are much smaller than one would expect in an experimental scenario, although realistic uncertainties were included for both XICS and VUV synthetic signals. This hints at a problem of over-determination in the forward model, which is found to be mostly related to the very large number of XICS data points ($\sim 50,000$). Inferences using only VUV signals (either synthetic or real) provide much larger uncertainties and mostly constrain the region of $\rho_p > 0.6$. This is in agreement with the observation that one needs to set a low statistical weight for XICS in experimental inferences in order to reduce such over-determination, as a result of




Figure 7-12: Comparison of inferences obtained on synthetic data created from model D and v profiles, labeled as “ground truth”, for the EDA H-mode case. The 5-coefficient case was model selected.

excessive dominance of the likelihood with respect to the prior. While a strong likelihood is generally desirable, Fig. 7-12 clearly shows that in our case we are erring on the side of “over-confidence”. Further discussion of these considerations will be provided in the next section.

In Fig. 7-13 we show the match between a synthetic XICS spectrum from the 5-coefficient inference of Fig. 7-12, which was model-selected by the Bayesian evidence within the set at different dimensionalities. Dark points with error bars (5% minimum relative uncertainties, 1% minimum absolute uncertainties) show the synthetic data. Red continuous lines show the result of running the forward model with the inferred D and v . Given that the D and v coefficients used in the forward model closely approximate the “ground truth” D and v of Fig. 7-12, the red lines match the synthetic data points very well. As discussed in Section 5.5, this is merely a consistency check on the inference framework, rather than evidence that analogous inferences on experimental data would necessarily be successful. In the next section, we demonstrate applications using experimental data.




Figure 7-13: Verification of close match of synthetic XICS signals (black points with 5% error bars) and a reconstruction (red lines) via the full-spectrum inference framework using the H-mode discharge parameters.

Diagnostic Weighting

When combining experimental signals from the full K_α spectrum, plus the 6 groups of VUV lines, the number of data points from XICS becomes much larger than the number from XEUS. If one were not to account for this, inferences would place most of the effort into matching XICS brightnesses, with relatively little weight given to fitting XEUS signals. Tests with synthetic data (Fig. 7-12) suggest that this also causes unrealistically small uncertainties. Ideally, setting a sufficiently low log-evidence tolerance condition in MultiNest would make the algorithm reach an arbitrary level of precision in the posterior exploration. However, in practice one must account for potential discrepancies between the two diagnostics which in our case can cause XICS to overwhelmingly determine the final posterior, at the expense of accurately matching VUV signals. From a physical perspective, this would be unsatisfactory because XICS cannot reasonably constrain transport at $\rho_p > 0.8$: the spectral lines of interest have low emissivity in these regions and the geometry

of XICS lines of sight does not permit significant signal integration there. On the other hand, the VUV lines of Table 7.1 clearly contain information about transport in the near-pedestal range and it is therefore strongly desirable for XEUS signals to be accurately matched in inferences. Previous work used only a few individual line brightnesses from XICS, such that the number of data points from XICS and VUV was similar. When using the full K_α spectrum, however, the numbers of data points from XICS and VUV signals ($\approx 46,850$ and ≈ 180 , respectively) are so different, their ratio being $\approx 180/46,850 \approx 0.004$, that the combination of diagnostics via an uninformative prior appears insufficient. The *Weighted Least Square Regression* (WLSR) method used in recent regression analysis for the international H-mode database [258] to combine data from multiple tokamaks uses the weights

$$w_k^{-1} = 2 + \sqrt{n_k}/4, \quad (7.4)$$

where n_k is the number of data points in dataset k . Applying this simple formula to our datasets gives $(2 + \sqrt{46,850}/4)^{-1} \approx 0.018$ for XICS and $(2 + \sqrt{180}/4)^{-1} \approx 0.187$ for XEUS. Another argument to motivate the use of weights smaller than 1 for each diagnostics is based on the observation that each dataset has an Effective Sample Size (ESS) that is smaller than the actual number of data points, due to spatio-temporal correlations in the measurements. In practice, the determination of appropriate weights is rather arbitrary, so these options were only used as an indication of reasonable weights, on top of which priors for diagnostics weights could be set as in previous analysis. Starting from a value of 0.004, we have run inferences with variable XICS weight within the range $[1, 10^{-4}]$, keeping the VUV weight fixed to 1 for simplicity. We have found a relative weight of 10^{-2} to satisfactorily match both XICS and VUV, while maintaining inferred D and v core features that are clearly attributable to XICS data, and D and v edge features that are clearly attributable to VUV data.

In order to more rigorously enforce the similarity of the final posterior distribution with the posteriors obtained in inferences including only one diagnostic at a time (over radial regions of each diagnostic’s applicability), one could attempt to minimize the Kullback-Leibler divergence between them [163]. In practice, such scheme could be implemented by pre-computing the posteriors obtained with each diagnostic by itself and then adding a penalty term in the combined inference that enforces the applicability of each inference

over a chosen radial region. While such approach has not yet been directly pursued, the simple weighting factor of 10^{-2} adopted here retains a principled, albeit not entirely rigorous, approach to combining diagnostics in our scenario of interest.

Inferences from Experimental K_α Spectra

Fig. 7-14 shows the result of an experimental inference for the I-mode discharge #1101014030, showing the radial distribution of knots from MultiNest samples with highest statistical weight. The distribution of these knots makes it apparent that the knots sampling scheme of Section 5.4.1 is very effective. In all shots presented in this section, a model with 4 radial knots (identical between D and v/D , to permit use of the sampling technique presented in Section 5.4.1) was selected via Bayesian evidence comparisons. We highlight that all knots were sampled in a fully uninformed manner from a unit hyper-triangle; their localization at approximately equally-spaced locations is purely dictated by the data in the case displayed in Fig. 7-14.

Fig. 7-15 shows the same inference result, but now comparing to neoclassical (NEO), quasilinear gyrofluid (TGLF SAT-1) and nonlinear, ion-scale, heat-flux matched, gyrokinetic (CGYRO) results (described in detail in Section 7.3). The values for D , v , and Rv/D from these theoretical models are the same as those shown in Fig. 7-9, although in this figure the experimental inference was derived from a less elaborate forward model (see Section 7.2.3) and the D was displayed on a log scale. In Fig. 7-15 all results are shown on a linear scale to offer a clearer visualization of the close agreement in structure of D , v and v/D at midradius. As in previous figures in this chapter, uncertainty bands represent 1-99, 10-90, and 25-75 quantiles of the local posterior distributions. Both D and v appear to be larger in the experimental inference than in theoretical models at $0.5 \lesssim \rho_p \lesssim 0.95$, but Rv/D appears closely matched. As in Fig. 7-9, the magnitudes of both D and v near and inside of the sawtooth inversion radius are each close to neoclassical levels, but Rv/D strongly differs from the neoclassical result. As discussed in Section 7.2.3, this is likely a result of the time-averaged nature of kinetic profiles given as inputs to NEO; a more accurate comparison near axis would require one to run NEO at multiple time slices during the sawtooth cycle.

Fig. 7-16 shows inferred transport coefficients for the L-mode discharge #1101014006. Very good agreement is found for diffusion within uncertainties, particularly at the location where TGLF, CGYRO and GYRO results were all computed. However, we note that VUV




Figure 7-14: Inference of Ca particle transport coefficients in the C-Mod I-mode discharge #1101014030, showing the distribution of free knots that were sampled with highest posterior probability via MultiNest. Each spline knot sample is shown by a grey circle; the overlap between many samples gives darker points.

signals appear to suggest that D should be lower by approximately a factor of 2 at midradius, in contrast to XICS. The result shown here is therefore sensitive to our choice of using a weight for XICS data of 10^{-2} , unlike in inferences for the other two discharges. The convection profile in Fig. 7-16 is also well matched, except for a region near $\rho_p = 0.4$ where $v > 0$ within uncertainties, in disagreement with NEO and TGLF results. As for diffusion, varying the XICS weight can affect convection at midradius, bringing it closer to $v = 0$ while keeping a fixed Rv/D . The lowest panel, showing Rv/D , is focused on the range $[-10, 10]$, allowing one to recognize the relatively good match between experimental and theoretical modeling at midradius. A flat $Rv/D \approx -2$ at $0.5 < \rho_p < 0.9$ is seen to be possible, within the posterior distribution. Near the magnetic axis, NEO predicts a large




Figure 7-15: Comparison of experimental Ca transport coefficients inferred using the full K_α spectrum in the C-Mod I-mode discharge of Fig. 7-14 with D , v , and Rv/D results from the neoclassical NEO and quasi-linear gyrofluid TGLF SAT-1 models. A single nonlinear, ion-scale, heat-flux matched, gyrokinetic result from the CGYRO code is also shown at midradius.

and negative Rv/D , suggesting very peaked profiles. As discussed above, the mismatch of NEO with experimental results inside the sawtooth inversion radius may be unsurprising due to the time-independent nature of kinetic profiles used for theoretical modeling. We remark that, as suggested already by P_{rad} time traces in Fig. 7-3, core impurity transport in L- and I-mode appears to be similar, even though the discharges discussed here are at different plasma currents ($I_p = 0.8$ and 1.0 MA, respectively).

Finally, Fig. 7-17 shows a comparison of inferred experimental transport coefficients and theoretical modeling for the EDA H-mode discharge #1101014019. Diffusion is found to be in excellent agreement, while significant disagreement is observed in v and Rv/D , except in the outer part of the plasma. The CGYRO prediction for D is very close to the TGLF




Figure 7-16: Analogously to Fig. 7-15, comparison of inferred Ca experimental transport coefficients in the C-Mod L-mode discharge #1101014006 with neoclassical (NEO), quasilinear gyrofluid (TGLF SAT-1), and nonlinear gyrokinetic (CGYRO) results.

SAT-1 one, whereas the CGYRO estimate of v is slightly lower, giving a smaller value of Rv/D which is in agreement with the inference. The inferred values of $Rv/D > 0$ up to $\rho_p \approx 0.6$ correspond to a hollow total impurity density profile, rather than peaked one as suggested by the NEO and TGLF models. As in the other two (higher current) discharges, a good qualitative match in the pedestal inward pinch is observed in this low- I_p EDA H-mode. Both D and v magnitudes inside of $\rho_p \approx 0.95$ are significantly lower than in the L- and I-mode cases.

Fig. 7-18 shows the EDA H-mode density profiles for the five highest charge states of Ca long after the injection, when profile shapes reach steady state and only profile amplitudes vary in time (similarly to Fig. 7-10). The total Ca density (black continuous line) is seen to be weakly hollow (almost flat, within uncertainties), in contrast to the rescaled electron




Figure 7-17: Analogously to Figs. 7-15 and 7-16, comparison of inferred Ca experimental transport coefficients in the C-Mod EDA H-mode discharge #1101014019 with neoclassical (NEO), quasilinear gyrofluid (TGLF SAT-1) and nonlinear gyrokinetic (CGYRO) results.

density (black dashed line).

Fig. 7-19a shows the comparison of XICS experimental signals (chord 0, slightly below the midplane, near the time of peak brightness) and the forward model reconstruction for the I-mode case, using median values from the posterior distribution. We highlight that, unlike Fig. 7-13, black data points in Fig. 7-19a are purely experimental, rather than synthetic/artificial. Most features of the spectrum are well reproduced, except for a systematic underestimation of the region near the intercombination lines (x and y) and a mismatch near the q and r lines. The latter is likely explained by inaccuracies in our background subtraction procedure, which attempts to remove the brightness of $n \geq 4$ He-like Ar satellites from the same spectral region. While this background is observed to be time independent before the LBO injection, our subtraction may in practice be imperfect.





Figure 7-18: Steady-state radial density profiles of the five highest Ca charge states from the experimental EDA H-mode inference of transport coefficients of Fig. 7-17. Dashed black lines show the rescaled electron density profile for comparison with the total Ca density profile.


Overall, the reproduction of XICS spectra to this degree of accuracy suggests that excellent transport constraints are offered by the diagnostic. XICS spectra from the L-mode and EDA H-mode discharge are matched equally well, although signal to noise ratios are slightly lower as a result of T_e being lower in these discharges.

Fig. 7-19b shows an analogous match between the Aurora forward model and experimental VUV data for the L-mode discharge. Here, we show the time history, rather than the spectrum itself, since we don't model the entire spectral range, but only the emission within specific wavelength bins (see details about each line group in Table 7.1), whose variation of amplitude in time is used to constrain transport. We remark that here all lines were normalized to a common maximum. Irregularities in the XEUS time histories (e.g. the bump near 0.02 s in Fig. 7-19b) reflect imperfections of LBO injections, where clusters of particles can take longer to reach the plasma. Our measurement of Ca I using an edge filterscope prevents significant errors in the source from propagating to the inferred transport coefficients. As with XICS, the *simultaneous* good match of all XEUS lines in Fig. 7-19b raises confidence in the inference of transport coefficients in Figs. 7-15, 7-16, and 7-17. VUV signal matches of analogous quality have been obtained for the 3 shots.

In the next section, we describe in more detail how predictions from theoretical models were obtained, what is their sensitivity, and how likely it is that experiment and theory agree within uncertainties in each case.



(a) Comparison of Aurora forward modeling (continuous red lines) and an experimental Ca K_α spectrum (black points with error bars) from the C-Mod I-mode case. The data displayed here is from the spatial channel #0, slightly below the midplane, near the time of peak brightness.



(b) Comparison of Aurora forward modeling (orange lines) and VUV signal time histories (blue points with error bars) for the L-mode case (see Table 7.1 for details of line groups). All signals have been normalized with respect to a common value to obtain additional constraints on transport of multiple Ca charge states.

Figure 7-19

7.3 Comparison of Experimental and Theoretical Modeling

Having described the quantification of experimental transport coefficients for the three discharges of interest, we now proceed to describe in greater depth the theoretical modeling to which these results are compared. The neoclassical and turbulence modeling in this section is in no way informed by the experimental results presented above, thus offering the means for an independent validation effort.

As discussed in Section 2.7, in the *trace limit* the transport coefficients of a given species offer a simple, yet effective parametrization of transport fluxes as linear in the thermodynamic gradients. 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 simulation, 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 vary⁴ 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:

$$\frac{R}{n_z} \begin{pmatrix} \Gamma_1 \\ \Gamma_2 \\ \Gamma_3 \end{pmatrix} = \begin{pmatrix} R/L_{n,1} & R/L_{T,1} & 1 \\ R/L_{n,2} & R/L_{T,2} & 1 \\ R/L_{n,3} & R/L_{T,3} & 1 \end{pmatrix} \begin{pmatrix} D \\ v_T \\ v_p \end{pmatrix} \quad (7.5)$$

By inverting the 3×3 matrix above, one can then solve for the local transport coefficients, D , v_T , and v_p . For comparisons to experimental inferences, one may omit the scan of temperature gradients since thermal convection cannot be separated experimentally. This is indeed how the NEO, TGLF and CGYRO D and v values in Fig. 7-9, 7-15, 7-16, and 7-17 were computed. It is nonetheless interesting to distinguish v_T from v_p in theoretical modeling to compare to predictions of thermodiffusion in different turbulent regimes [259], as we will see in following sections.

Neoclassical NEO modeling

Our comparison of experimentally-inferred D and v with NEO results is inevitably approximate, given that NEO was run with time-averaged kinetic profiles, as opposed to the

⁴The actual percentage variation is inconsequential for trace-level concentrations.

time-dependent Aurora calculations including multiple impurity charge states and sawtooth modeling. The NEO profiles shown in previous sections have been obtained for Ca^{18+} (the dominant charge state throughout most of the plasma radius in all discharges). This is not particularly important in the core, where the difference between Ca^{18+} and adjacent charge states is not significant. On the other hand, a wider effort should be devoted in the future to obtain an accurate comparison in the pedestal, where many charge states co-exist.

NEO scans of electron density and temperature gradients for $\rho_p > 0.92$ have revealed that the amplitude of the neoclassical edge pinch can vary significantly within experimental uncertainties. The results displayed in Figs. 7-15, 7-16, and 7-17 make use of nominal values of all NEO inputs for the computation of transport coefficients for Ca^{18+} . More detailed comparisons in the pedestal should use NEO for each charge state and possibly weigh their relevance to the overall transport by the fractional abundance of each charge state at each radial location. Taking fractional abundances at ionization equilibrium would however give erroneous predictions, since impurities are well out of ionization equilibrium in the pedestal. Using fractional abundances from an experimental inference would make theoretical modeling dependent on experimental modeling, but this may not be an unacceptable compromise.

In its most basic forms, neoclassical theory predicts a linear scaling of v on Z and no dependence of D on Z . Such effect has been observed in experimental studies looking at steady-state pedestal gradients (related to v/D) for multiple impurities [33]. To the author's knowledge, such scaling has never been directly measured or inferred from experimental data for different charge states of a single ion species. Fig. 7-7 showed an attempt at determining whether such a $v \sim Z$ scaling could be confirmed via Bayesian model selection, but our results using C-Mod data were shown to be statistically inconclusive.

Beyond the simple scaling above, a complete numerical model such as the one in NEO shows a more complex set of nonlinear dependencies. Fig. 7-20 shows scans of D and v with charge Z in the pedestal ($\rho_p = 0.97$) for the C-Mod EDA H-mode discharge discussed above. A near-linear dependence of v on Z is consistently found in C-Mod pedestal simulations.

Given the high accuracy of PECs for selected Ca VUV lines used in combination with the full K_α spectrum (Section 7.2.5), we tested whether this extended data set could possibly constrain the Z dependence of transport in the pedestal. Within the C-Mod inferences discussed thus far, we introduced a free parameter η , defined by the relation




Figure 7-20: Scan of Z in NEO simulations for the C-Mod EDA H-mode discharge of Fig. 7-17 at nominal values of all input parameters at $\rho_p = 0.97$.

$V(Z) = V_{\text{ped}}(Z/Z_{\text{max}})^\eta$, where V_{ped} is the amplitude of the Gaussian feature used to parametrize the pedestal inward pinch. Our results indicate, perhaps not surprisingly, that the data do not effectively constrain η . This is likely the result of measurements only constraining the pedestal transport of Be-like and Li-like Ca (and, to a smaller extent, lower charge states), which are close in Z values. Moreover, freedom in the diagnostics' time bases made the task more challenging, resulting in η being different in each shot and sensitive to modeling choices, making its interpretation impossible. Future work may revisit these methods if experimental data from a sufficiently wide range of charge states becomes available.

Quasilinear TGLF modeling

In this section, we explore the sensitivity of the radial profiles of impurity transport coefficients predicted by TGLF, shown in Figs. 7-9, 7-15, 7-16 and 7-17. Fig. 7-21 shows the estimated fractional uncertainties (1 standard deviation) of the normalized gradient scale lengths of n_e , T_e , and T_i for the 3 discharges of interest, computed from the right hand panels of Fig. 7-5. We limit the radial range to $0.4 \leq \rho_p \leq 0.8$ to offer better visualization of the region where temperature gradient uncertainties allow meaningful input parameter scans within turbulence models. We remark that the EDA H-mode discharge has much larger fractional uncertainties in a/L_{n_e} than the other two cases because the density profile is almost flat ($a/L_{n_e} \approx 0$), hence making a 100% change still a small number.

For all TGLF simulations of C-Mod plasmas, we include all electromagnetic components




Figure 7-21: Experimental fractional uncertainties of normalized kinetic gradient scale lengths (for n_e , T_e , and T_i) and the ratio of T_i and T_e for the 3 discharges of interest.

and make use of the SAT-1 “multiscale” saturation rule [49, 46] since this is found to better reproduce kinetic profiles while matching experimental heat fluxes from TRANSP [256] within the TGYRO framework [52]. Fast ions (from ICRH) were excluded from simulations, since currently their effects are not appropriately captured by the models used in this work.

Analysis of the TGLF linear spectrum shows that ion-scale modes with $k_y \rho_s$ up to 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. ρ_s is the ion sound speed gyroradius. Linear spectra display a strong sensitivity to a/L_{T_i} and weaker sensitivity to other typical drift wave drives, suggesting that ITG modes are strongly dominant in all the plasmas considered here, in agreement with past C-Mod simulations [260, 261, 26]. The observation of positive thermodiffusion, $v_T > 0$, also confirms the identification of ITG being dominant [259]. While TGLF and CGYRO turbulence spectra are different, particularly at scales where ITG and TEM co-exist, they agree in the identification of the most intense turbulent mode being at




Figure 7-22: Diffusion and convection coefficients from TGLF 2σ scans around experimental I-mode parameters. Uncertainties are estimated locally at each radial location for each parameter.

$k_y \rho_s \approx 0.4$ in all cases. In Fig. 7-22 we show independent TGLF scans of a/L_{T_i} , $\nu_{ei}/c_s a$, a/L_n , and T_i/T_e within 2 standard deviations, as estimated from experimental data, for the previously described I-mode discharge (cf. Fig. 7-5). Here, c_s is the local sound speed and ν_{ei} is the electron-ion collision frequency, for which we considered an uncertainty of 20% following Ref. [262]. We use 2σ rather than the more conventional 1σ 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 magnitude. A similar conclusion is obtained in TGLF scans for the L- and EDA H-mode cases. We note, however, that a complete assessment of uncertainties and comparison to experimental inference would require varying input parameters simultaneously; correlations across TGLF inputs could be important, but cannot be explored via 1D scans such as those of Fig. 7-22. For this purpose, future work should explore specialized optimization tools such as VITALS [263]. Interestingly, v/D is almost fixed within the radial range that we explore; in other words, changes of D and v correlate linearly 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} . 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 in this region.

Gyrokinetic modeling

We next describe results from nonlinear ion-scale gyrokinetic simulations with CGYRO [37]. Simulations for each shot are focused on the flux tube at $r/a = 0.6$, where experimental profiles are well determined and the effect of sawteeth is expected to be small. We use 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 energies, and 24 poloidal points. We adopt experimental profile inputs, Miller geometry, electromagnetic ($\tilde{\phi}$ and \tilde{A}_{\parallel}) effects, as well as gyrokinetic electrons. We apply the sonic rotation scheme available in CGYRO [38], whose effects with respect to the linear ordering in Mach number are known to account for significant differences of transport for heavy impurities [108]. 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, have been run by increasing each parameter by 50% in turn and checking for significant variations in output. Linear scans of collisionality do not display strong dependencies in any outputs, but do highlight that collisionality 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 [75, 76, 77]. Application of the Sugama operator [264, 39, 37], 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} have been used to find the value at which CGYRO can match the experimental turbulent ion heat flux, Q_i , as determined by power balance via TRANSP [265]. For the I-mode discharge of Fig. 7-15, for example, a Q_i match is found 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 [260], we find the I-mode 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 [266, 267], we observe a clear underestimation of




Figure 7-23: 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. Reproduced from Ref. [155].

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%) see no significant variation of Q_e/Q_i (less than 10%), which remains generally less than 1/2 of the value predicted by TRANSP. This is consistent with expectations from linear CGYRO scans, which show little sensitivity to typical TEM or ETG drives for the 3 cases of interest, suggesting that multiscale simulations [46] 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. 7-23 we show the time history of heat and particle fluxes for two I-mode simulations that differ 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 ± 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.

Fig. 7-24 shows spectra of heat and particle fluxes, for both deuterium ions and electrons, as well as D and v components for Ca trace impurities at the simulated radial location ($r/a = 0.6$). From the left, results are shown for the L-mode, EDA H-mode, and I-mode cases. Spectra of diffusion coefficients show a remarkable similarity to the heat flux coefficients (of both deuterium ions and electrons), whereas particle fluxes show more complex features. Interestingly, in the L-mode and I-mode cases the particle fluxes change sign (direction) at $k_\theta \rho_s \approx 0.45$ and 0.35 , respectively. On the other hand, the EDA H-mode has $\Gamma > 0$ (outward) for both deuterium ions and electrons. Since particle fluxes must remain ambipolar, i.e. $\sum_s q_s \Gamma_s = 0$ with s being a species index, the sum of $q_s \Gamma_s$ for all simulated impurities (a “lumped” Ne impurity, plus Ca trace ions) must be inward overall. Indeed, the sign of total Ca convection is found to be negative in the bottom panel. In agreement with the TGLF scans of Fig. 7-22, thermodiffusion (v_T) always gives $v > 0$ (outwards) contributions, as expected for ITG-dominated regimes [259]. The remaining pure pinch, v_p , is always larger and negative, resulting in predictions of peaked profiles in all 3 shots at this location ($r/a = 0.6$).

Fig. 7-25 shows time histories during the CGYRO simulations for heat and particle fluxes, and D and v/D components. As in Fig. 7-24, left plots refer to the L-mode, those in the center to the EDA H-mode, and those on the right to the I-mode. Significantly stronger and more intermittent turbulence is seen in the L-mode case. As shown in the comparison to the experimental inference in Fig. 7-15, CGYRO appears to under-predict the transport levels in the I-mode case. The fact that TGLF gives larger D and v cannot be the result of it being a “better model” than CGYRO, since it is a quasilinear and gyrofluid approximation to the gyrokinetic model implemented in the latter.

The choice of displaying v/D rather than v in Fig. 7-25 was made to allow a comparison of the nonlinear CGYRO predictions with quasi-linear CGYRO runs, from which an estimate of v/D can be obtained under the assumption that both D and v saturate in a quasi-linear manner. If so, the turbulence saturation amplitude, $\tilde{\Phi}^2$, setting the scale of both D and v , should cancel out (see, for example, Eq. 2.23). Continuous lines in Fig. 7-25 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 $k_y \rho_s \approx 0.4$ in each case. While the dashed line obviously does not match the total D in magnitude (continuous line),

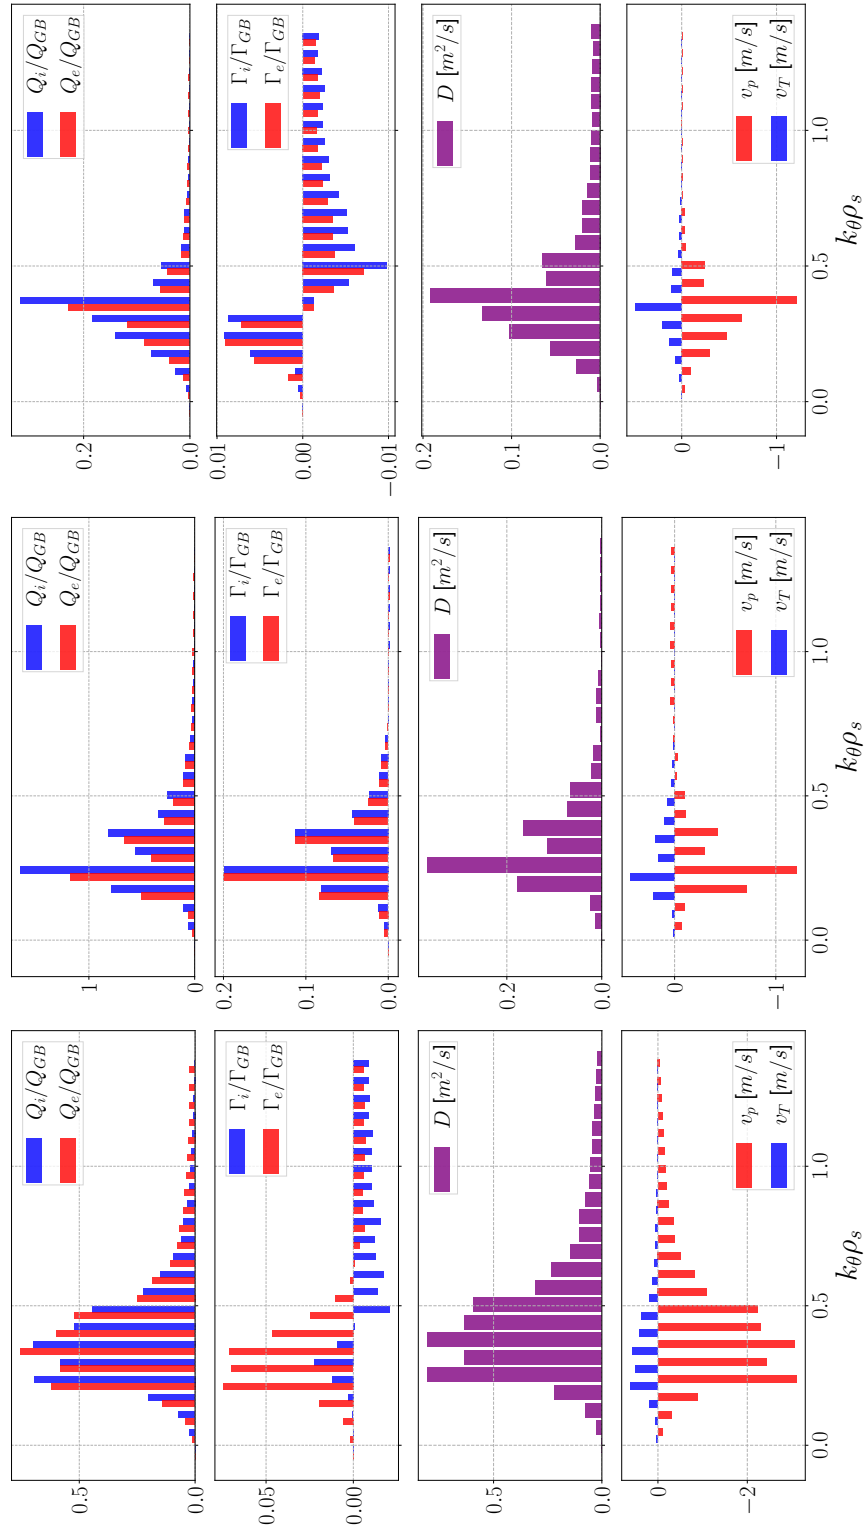


Figure 7-24: CGYRO spectra for L-mode (left), EDA H-mode (center), and I-mode (right) shots.

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. The quasilinear CGYRO results, obtained at $k_y \rho_s = 0.4$, are shown by blue and green crosses for thermodiffusion (v_T/D) and pure convection (v_p/D), respectively, at the end time of the nonlinear simulation. In each case, a good match of v/D components between quasilinear and fully-nonlinear theory is found, suggesting that quasilinear impurity peaking predictions such as those of TGLF may be appropriate.

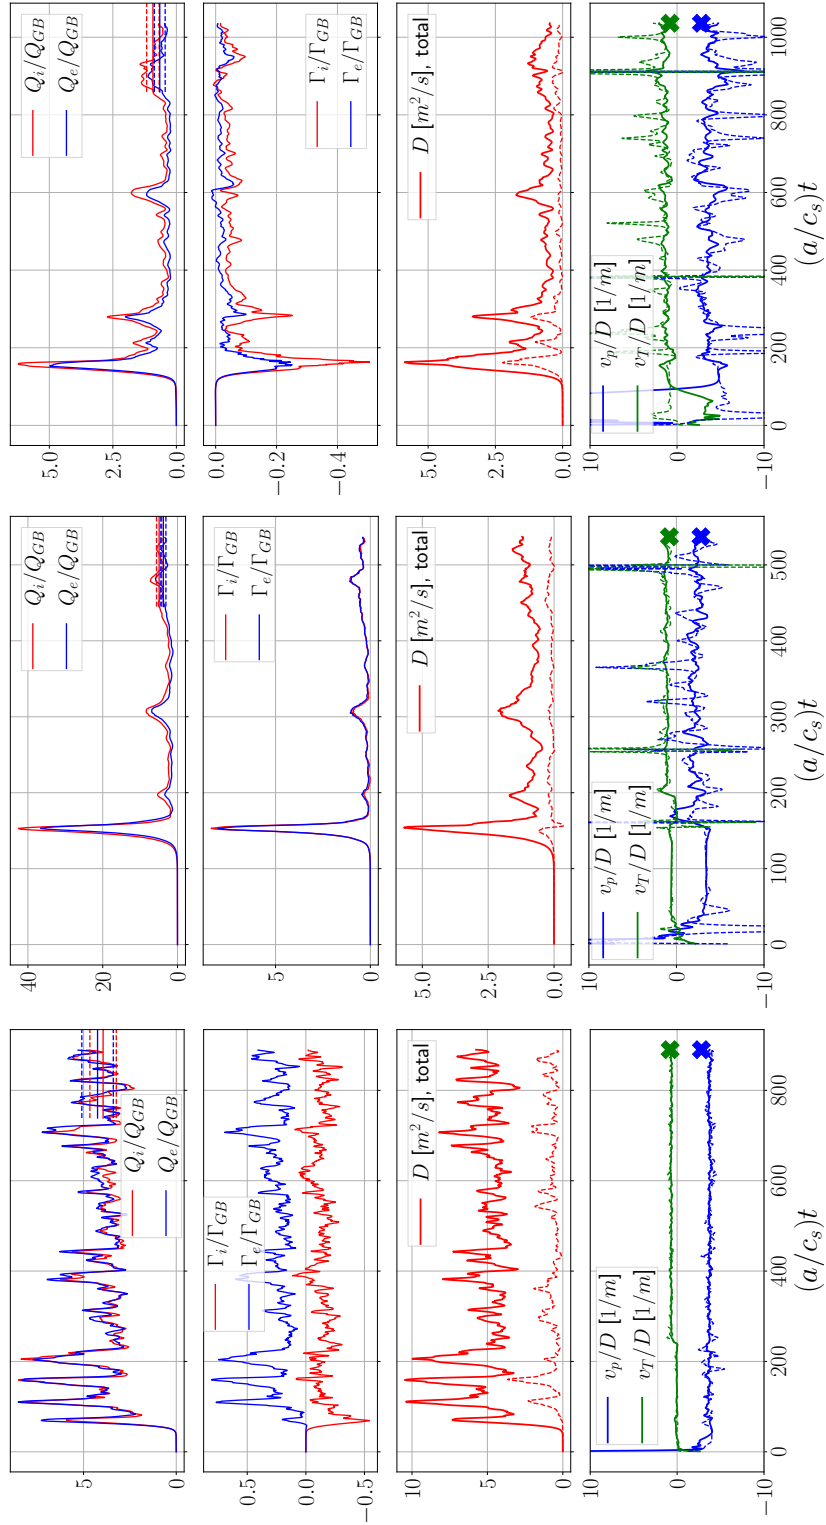


Figure 7-25: Time histories of heat and particle fluxes (for D ions and electrons) and transport coefficients (D and v/D components) from ion-scale heat-flux matched nonlinear CGYRO simulations for the L-mode (left), EDA H-mode (center), and I-mode (right) shots.




Figure 7-26: Comparison of Rv/D components from experimental inferences (crosses) and nonlinear CGYRO simulations, separating pure convection (squares), thermodiffusion (circles) and total Rv/D (stars) for the Alcator C-Mod L-mode (green), EDA H-mode (orange), and I-mode (lilac) discharges. Vertical bars display the interquartile range for experimentally-inferred values.

Finally, In Fig. 7-26 we show a comparison of Rv/D components from experimental inferences (crosses) and nonlinear CGYRO simulations. We separate CGYRO estimates into pure convection (squares), thermodiffusion (circles) and total Rv/D (stars) for the Alcator C-Mod L-mode (green), EDA H-mode (orange), and I-mode (lilac) discharges. Results are plotted as a function of effective collisionality, ν_{eff} (Eq. 7.1). Vertical bars display the interquartile range for experimentally-inferred values. The I-mode, at the lowest collisionality, is predicted by CGYRO to have similar Rv_{tot}/D to the L-mode. On the other hand, the EDA H-mode has more than twice the I-mode value of ν_{eff} and is found to have much smaller peaking, according to both CGYRO and the experimental inference. The collisionality trend observed here matches the general observation that electron density peaking is lower in high-collisionality H-modes than in L- and I-modes [244, 77], as shown in Fig. 7-2. We note that the EDA H-mode case appears to have an excellent match between theory and experiment, but as shown in Fig. 7-17 the nonlinear CGYRO simulation is actually at a radial location where there exist significant variation in Rv_{tot}/D , thus suggesting that this agreement may be fortuitous. In the I-mode case, the experimental inference suggests a larger Rv_{tot}/D than CGYRO does. Finally, a very good match is observed for the L-mode case, with the CGYRO Rv_{tot}/D being well within the experimental interquartile range.

7.4 Discussion and Summary

Previous sections presented the progress made during this thesis on the inference of impurity transport coefficients and their theoretical understanding on Alcator C-Mod. Developments on both forward and inverse modeling have allowed us to make the most of HPC resources and modern computational statistics. We have shifted from constraining D and v based on the w resonance line of He-like Ca to using multi-line constraints and eventually demonstrating an advanced framework utilizing the entire Ca K_α spectrum. The use of VUV spectroscopy has also been widely expanded by making use of high-quality atomic data from ADAS to take into consideration line ratios from emission due to different charge states in the outer part of the plasma.

Detailed inferences of radial profiles of transport coefficients have been presented for L-, EDA H- and I-mode discharges, finding relatively good agreement in both diffusion and convection. However, unexplained discrepancies have been found inside of $\rho_p = 0.6$ in the EDA H-mode discharge, where theoretical predictions suggest that a peaked density profile (i.e. $v/D < 0$) is expected in this region. We note that our simple model for the effects of sawteeth is unlikely to affect inferences significantly, since the sawtooth inversion radius is near $\rho_p = 0.2$ in this discharge and temperature crashes are small and well resolved by ECE measurements. TGLF and CGYRO predictions cannot be easily reconciled with experimental inferences within uncertainties, both considering those in experimental and theoretical work. We note, however, that physical turbulence drives were not scanned with CGYRO, thus our sensitivity assessment is mostly based on TGLF scans like the one shown in Fig. 7-22. Given the relatively good agreement of CGYRO with TGLF in the cases that we have tested, we expect TGLF to give a fair representation of the physics of interest. Based on these observations, at present it appears unlikely that further experimental constraints would resolve the discrepancies found in this EDA H-mode case in the future. These observations suggest that state-of-the-art gyrokinetic models may be missing critical physics related to particle convection. Given the known stabilizing effects of fast ions on ITG turbulence [268, 269], it is possible that future work including fast ions in theoretical modeling may successfully account for these issues.

We highlight that the L- and I-mode cases have been found to have relatively similar transport coefficients, while the EDA H-mode has smaller D and v , according to both theory

and experiment. The similarity between L- and I-mode, which have close values of $I_p = 0.8$ and 1.0 MA, respectively, is in agreement with general observations of particle transport being similar between these two regimes, even though I-modes have a pedestal heat transport barrier [23]. On the other hand, the EDA H-mode case has lower $I_p = 0.55$ MA, which may suggest lower confinement, and therefore larger transport coefficients with respect to the L- and I-mode cases. The small D and v in Fig. 7-17 may seem to contradict this intuition. In fact, global impurity confinement times, τ_p , appear to be dominated by the size of the density pedestal pinch in EDA H-mode. Previous work on Alcator C-Mod has shown a strong positive correlation between impurity confinement time, τ_p , and I_p in EDA H-mode [252], relating it to variations in pedestal density gradients [224]. The EDA H-mode in Fig. 7-17 has slow transport in the core, but a comparable pedestal pinch to the L- (Fig. 7-16) and I-mode (Fig. 7-15) cases, resulting in larger τ_p overall. At higher I_p , core cross-field transport may become slower, but the clear correlation of τ_p with pedestal density gradients suggests that confinement is mostly affected by the linear growth of the density pedestal with I_p .

Research presented in this chapter has made significant progress towards constraining particle transport in the C-Mod pedestal, although our diagnostic means appear inadequate for a complete examination of the subject. As discussed in Chapter 9, the validation of pedestal transport models requires a dedicated effort to develop diagnostics whose spatio-temporal resolution may appropriately resolve the relevant scales. It is also paramount to have clear spectroscopic fingerprints for the dynamics of individual charge states and a reliable calibration of each detector. As part of this work, we attempted to make use of the C-Mod Soft X-Ray (SXR) arrays to better constrain the fast signal rise after the LBO injection. Unfortunately, we found that this SXR system, installed at the very beginning of C-Mod operations, appears to lack accurate relative calibration among spatial chords, at least for the experiments discussed in this chapter. It is also possible that a mis-characterization of the Be filter and the Si photodiode substrates in our analysis makes the atomic data computed for this analysis unable to explain observed signals.

The next chapter will present results from DIII-D impurity transport analysis, again making use of Bayesian inference methods with Aurora. Tokamak regimes and diagnostic setups examined for DIII-D inferences are different from C-Mod ones, thus motivating a generalization of analysis tools for cross-device analysis.

Inference of Impurity Transport on DIII-D

In this chapter we discuss inferences of impurity transport in DIII-D, particularly focusing on RMP ELM-suppressed H-mode and diverted negative triangularity experiments. We present the `ImpRad` module developed for particle transport analysis based on Aurora. As in the previous chapter, we compare inferred D and v with theoretical modeling, again finding discrepancies in the modeling of hollow impurity profiles.

Introduction

The inference of particle transport presents common challenges across tokamaks. One of these challenges is, of course, the universal nature of neoclassical and turbulent phenomena, which is indeed the fundamental subject of this investigation. Another common aspect is the need for detailed spectroscopic measurements, without which the effect of transport cannot be clearly discerned. The generalization of experimental analysis methods to operate on multiple devices constitutes an element of strength and reliability for the validation of universal transport models.

In this thesis, significant effort has been devoted to the development of software to infer impurity transport coefficients for experiments on both Alcator C-Mod and DIII-D, as well as other devices world-wide. This has been greatly facilitated by the collaborative nature of the One Modeling Framework for Integrated Tasks (OMFIT) [198], with which the Aurora

toolbox (Section 5.3.1) is integrated. Previous impurity transport research by Grierson [87] led to the creation of a DIII-D analysis framework for impurity transport based on gas puffing and the STRAHL code [195]. In close collaboration with T. Odstrčil, this has been further developed into the **ImpRad** module in OMFIT.

ImpRad, described in Section 8.1, is widely applicable for the analysis of particle transport in DIII-D scenarios, particularly in combination with Laser-Blow Off (LBO) impurity injections. In Section 8.2, we describe how the **ImpRad** tools are combined with new data analysis strategies, some of them analogous to methods applied for C-Mod research as described in Chapter 7. Inferences of impurity transport are then demonstrated for two DIII-D scenarios.

First, we revisit experiments in ITER-similar plasma shapes that were previously analyzed by Victor *et al.* and described in Ref. [270]. These experiments, discussed in Section 8.3, assessed how impurity transport is modified by the application of perturbative techniques that mitigate the impact of Edge Localized Modes (ELMs) in H-mode discharges. Our inferences of radial profiles of diffusion (D) and convection (v) for one of these experiments are described in Section 8.3.1. Here, we present an analysis that makes use of the same experimental data as in Ref. [270], but completely independent methods.

The second part of this chapter discusses impurity transport in diverted negative triangularity discharges on DIII-D. As discussed in Section 8.4, negative triangularity is a promising concept for reactor operation and is a subject of on-going investigation at DIII-D. Impurity transport plays a central role in motivating interest in this scenario, since low impurity retention has been found in conditions of high energy confinement. Section 8.4.2 presents inferences of D and v profiles using **ImpRad** for three experiments that were executed during this PhD, as part of a wide collaboration to demonstrate the potential of negative triangularity for future devices.

In Section 8.4, we provide a brief overview of research on negative triangularity and motivate our contribution on this topic. Section 8.4.2 describes the inference of impurity transport coefficients in three DIII-D diverted discharges with negative triangularity. In Section 8.4.3 we compare experimental particle transport coefficients with NEO and TGLF modeling, before summarizing in Section 8.5.

8.1 The OMFIT ImpRad Module

The particle transport analysis described in this chapter leveraged Aurora within the `ImpRad` module, developed in collaboration with T. Odstrčil as a new component of the One Modeling Framework for Integrated Tasks [198] (OMFIT). OMFIT is a collaborative Python-based framework that assists in the analysis of fusion experiments and theoretical modeling, with a focus on coupling physical models together. OMFIT offers a set of Graphical User Interfaces (GUIs), many of which have wide applicability within transport analysis, for example in the creation and processing of magnetic equilibria from tokamak experiments via the EFIT code [166]. In this section, we briefly describe the purpose, breadth, and flexibility of `ImpRad` as an expansion of previous OMFIT capabilities to study impurity transport.

The screenshot shows the 'Background' tab of the ImpRad GUI. At the top, there are tabs for 'Setup', 'Namelist', 'Run', 'View Results', and 'Inference'. Below these is a '(Re)Load Default ImpRad Parameters' button. The 'Expert mode' checkbox is checked. A sub-menu at the bottom of the 'Background' tab includes 'Grid', 'Impurity', 'atomdat', 'Edge/Recy.', 'Edge parameters', 'Transport', 'Sawteeth', and 'Elms'. The 'Main ion' is set to 'Deuterium'. The 'Initialize carbon density from experiment' checkbox is checked. A note states 'The CX flag is on -- specify H/D neutral profiles:'. The 'Neutral H/D profiles' are set to 'From SOLPS run'. The 'Scaling factor of neutral density profile' is 0.427923222521167. The 'Path to baserun and SOLPS case output files' is '/projects/toolbox/sciortinof/SOLPS/Hmode_180520_2425/'. The 'Name of specific SOLPS run to load from the path above' is 'test_D+C_NPup'. The 'Path to geqdsk for SOLPS run' is 'tinof/SOLPS/Hmode_180520_2425/baserun/geqdsk_fs_180520_2500'. There are 'Preview SOLPS neutral profiles' and 'Preview NBI CX rates' buttons. A note states 'The NBI CX flag is on -- specify impurity recombination rates due to NBI neutrals:'. The 'Recomb rates due to NBI' are set to 'From FIDASIM module'. There is a 'Plot beam sequence' button. The 'Neutral beams to include' are ['30L']. The 'FIDASIM module' is '[FIDASIM_V1.0] -- OMFIT[FIDASIM_V1.0]'. The 'Include fast neutrals' and 'Include halo neutrals' checkboxes are checked.

Figure 8-1: `ImpRad` tab for the selection of background ion species and neutral densities to compute CX rates. In this case, the main ion species is set to be D, C densities are initialized from fits of CER measurements, edge neutrals are taken from a SOLPS-ITER simulation, and NBI neutrals are from a FIDASIM calculation.

`ImpRad` has offered a development and test bed for Aurora itself. Aurora is used for fast simulations of impurity transport, radiation predictions, development of

synthetic diagnostics and inferences of impurity transport coefficients. `ImpRad` allows for benchmarking of Aurora with the historical STRAHL code in individual runs, permitting an exact comparison when inputs between the two codes are matched. The `ImpRad` GUI, one tab of which is shown in Fig. 8-1, allows the creation of detailed Aurora namelists¹, simple loading of magnetic geometries and kinetic profiles from other OMFIT modules, visualization and post-processing of results. Complex models with time-dependent transport coefficients, possibly arising from sawteeth or ELMs, can be easily constructed via the GUI. A set of workflows for the inference of impurity transport coefficients from experimental data has also been implemented, permitting nonlinear optimization with LMFIT [271], Markov Chain Monte Carlo with the emcee package [189], and nested sampling with MultiNest [193]. These iterative frameworks make use of a set of synthetic diagnostics developed for a range of measurements, including from Extreme Ultra-Violet (EUV) spectroscopy, Soft X-Ray (SXR) arrays, Charge Exchange Recombination (CER), Visible Bremsstrahlung (VB), and bolometry. Inferences of particle transport can be executed on remote clusters to access high-performance computing resources, where Aurora simulations can be run in parallel. The `ImpRad` module has been generalized to minimize device-specific tasks and allow utilization of the framework for multiple experimental facilities in the future.

8.2 Experimental Data and Analysis Methods

In this section, we describe the data and methods adopted for our investigations of RMP ELM-suppressed H-mode and negative triangularity experiments, presented in Sections 8.3 and 8.4, testing the most advanced `ImpRad` capabilities and comparing to neoclassical and turbulent transport models.

LBO Source Functions

Fig. 8-2 shows time traces of the emission measured by an unfiltered fast camera viewing the DIII-D LBO port during the three impurity injections into negative triangularity discharges discussed in Section 8.4. Signals from all pixels viewing a rectangular region surrounding the LBO port have been summed in order to obtain a 1D (temporal) signal, subtracting background from a slightly wider surface area on the image. The high time-resolution of

¹<https://Aurora-fusion.readthedocs.io/en/latest/params.html>

this instrument allows for a quantification of the time history of impurity sources from the LBO injections, similarly to the filterscope Ca I traces used for C-Mod (Section 7.2.2). Fig. 8-2 shows a non-negligible degree of variability of the shape of LBO sources. For example, in the case of #180530 we see that the LBO actually fired slightly before the programmed time of injection. Moreover, in both #180520 and #180530 (the two cases with Al injections) there is evidence of particles reaching the plasma several ms after the laser fired. This is expected to occur when the ablation is incomplete and heavier clusters of particles proceed more slowly than individual particles for a given kinetic energy [20].




Figure 8-2: Effective source functions for LBO injections of F (for shot #180526) and Al (for shots #180520 and #180530) from the spatial integration of fast camera unfiltered signals.

In practice, it is likely that the presence of these clusters is underestimated by the time histories in Fig. 8-2, possibly because clusters may ablate slowly and over a longer time scale. On the other hand, in #180526, where LiF was injected, there is no evidence for such clusters (the time trace decays monotonically). Indeed, this is confirmed by the lack of secondary peaks in signals from low charge states measured via EUV spectroscopy. As discussed in Section 8.4.2, it is important to assess the sensitivity of inferences of transport coefficients to details of source time histories. Analysis on both Alcator C-Mod and DIII-D has found that the underestimation of LBO clusters can have significant impact on inferred transport coefficients in the pedestal. Hence, availability of a dedicated filterscope, as in C-Mod experiments, or a fast camera, as in DIII-D ones, is very important. Some LBO film compositions are found to clearly ablate better and lead to fewer clusters. LiF and CaF₂ appear particularly favorable (in both of these cases, unfiltered fast camera measurements are effectively accounted for by the heavier of the two atomic species in the molecular

compound). Al is typically easily ablated, but in the experiments described here the laser energy appears to have been too low to allow a clean ablation, preventing us from accurately estimating pedestal transport coefficients.

Multi-Impurity CER Analysis

Fig. 8-3 shows the density of F^{9+} measured by all the CER chords available in shot #180526, with the approximate radial location of each chord indicated in the legend. The signal for Al^{13+} measurements in discharges #180520 and #180530 was observed to be of lower intensity, but still high quality. In processing the intensity measured by CER spectrometers, neutral beam attenuation was taken from routines developed by the DIII-D CER group. Photon Emissivity Coefficients (PECs) for the transition of interest were taken from ADAS [114]. The resulting densities are helpful to approximately understand the order of magnitude of particles entering the plasma, but carry too much uncertainty for a detailed comparison to fluxes from the LBO injection, which are themselves only roughly known. We note that the error bands shown in Fig. 8-3 only quantify statistical uncertainty from measured kinetic profiles and CER intensity, with no attempt made at estimating uncertainties in atomic rates. In the impurity transport analysis discussed in this chapter, no direct use has been made of the absolute calibration of these density measurements.

For the first time, we apply a novel scheme to constrain impurity transport inferences from LBO injections making use of simultaneous C measurements from additional CER




Figure 8-3: CER signals from F^{9+} following the LBO injection into shot #180526. The legend shows the approximate radial localization of each channel.

channels. Being intrinsic to DIII-D plasmas, C is almost always measured by the CER diagnostic in order to estimate ion temperature, T_i , and rotation frequency, ω_ϕ . T_i measurements from C^{6+} are also used to constrain fitting of spectra of weaker intensity from non-perturbative LBO injections. The element of innovation brought in by the new scheme adopted here is in making use of C density measurements such as the ones in Fig. 8-11 to constrain steady-state *density* profiles of low- Z impurities, rather than just temperature and rotation values. In practice, C transport is simulated using Aurora by initializing C^{6+} profiles in the plasma from the fitted CER values. We then apply a constant source of neutral C in the SOL and allow it to evolve according to the chosen D , v and background plasma profiles. When comparing C^{6+} profiles between Aurora and CER measurements, we ignore the first 50 ms of physical time after the start of the C injection to avoid an initial phase of profile stabilization. After this short time, comparison to C^{6+} experimental measurements offers valuable constraints on the v/D steady-state profile. We note that background plasma profiles are allowed to vary during this time and C^{6+} may be accumulating or decreasing in density over time, depending on the specific plasma scenario of interest.

The neoclassical transport of C is expected to be different for ions of different Z . Turbulent transport is also likely to have some dependence on Z , but this is likely to be negligible within the achievable accuracy of this experimental study. As shown in Fig. 7-20 for a C-Mod case, the variation of neoclassical diffusion with Z is also generally small, so we take D to be the same for all the low- Z impurities considered here across the entire radius. On the other hand, convection is allowed to differ among ions near the magnetic axis and in the pedestal, but not at midradius. To allow enough freedom in inferences of experimental impurity transport, we set a single spline knot to be inside the $q = 1$ surface and have different values for each impurity (C and the impurity injected via LBO). All other spline values of v are set to be the same among ions across the radial profile. As described in Section 5.4, we make use of a Gaussian feature in the pedestal to represent the sharp changes expected in this region. The amplitude of this Gaussian is set to be different for each ions, with the amplitude for carbon forced to be lower than the amplitude for higher- Z ions, inferred with a prior of $\mathcal{N}(0, 50)$ m/s. Since high-quality steady-state C measurements from CER are typically available in the pedestal, this constrains a minimum v/D for LBO-injected impurities. This multi-impurity inference scheme has been found to be extremely effective and is a key component of the results discussed in this chapter. The combination

of CER measurements of C and other ions injected via LBO is likely to be applied in all future DIII-D work using the `ImpRad` module.

When attempting to infer particle transport coefficients in the pedestal, the question arises of whether the mapping of measurement locations in the edge is sufficiently accurate. As discussed in Chapter 6, for n_e and T_e a constraint on the separatrix location can be set based on a 2-point model prediction of the local T_e value. This allows one to shift n_e values radially since they are measured by the same diagnostics that measure T_e (e.g. Thomson scattering). For DIII-D research shown in this chapter, the $T_{e,sep}$ value was fixed to 80 eV for simplicity. For impurity densities, finding an appropriate correction to the LCFS position is harder. It may be argued that the location of the electron density pedestal top offers a good constraint on where the pedestal top for any impurity is located, but this simplification can be very misleading when impurities, e.g. C, are not fully ionized and only the fully-stripped stage is directly measured via CER. For example, in DIII-D L-mode discharges there may exist significant C^{5+} fractional abundances also inside of $\rho_p = 0.8$. The alignment of n_e and impurity density pedestal tops is likely more accurate in H-mode discharges, where a strong inward pinch and higher T_e cause C to be fully-ionized in the pedestal. Clearly, however, such shifts carry significant uncertainties and they can be counter-productive. For the analysis discussed in this chapter, no shifts of CER data were executed. In practice our results are expected to be only weakly dependent on this feature, which however should be carefully considered in future pedestal impurity transport studies.

EUV Spectroscopy

Data from the core-SPRED (hereafter simply called “SPRED”) EUV spectrometer (Section 4.2) have been used to constrain both Al and F transport in the outer part of the plasma. As in C-Mod analysis, we make use of ADAS PECs produced from cross sections computed via R-matrix calculations. However, we have found it hard to match multiple EUV lines simultaneously, unlike for C-Mod where we could make use of the relative amplitudes of Ca lines measured by XEUS. One explanation for these difficulties is that the F and Al atomic data from ADAS often have atomic energy levels that do not match those given by NIST [136], which makes it hard to accurately add PEC contributions within chosen wavelength bins. Nonetheless, we have attempted to identify by inspection any wavelength shifts for the strongest lines in the spectrum. Such assessments are likely to be inaccurate in

cases where multiple lines overlap within the spectral resolution of the SPRED diagnostic, since each line could in principle have a different wavelength shift.

In view of these complications, the inferences of impurity transport discussed in this chapter normalize each EUV line signal independently, as done in previous impurity transport work on both C-Mod and DIII-D. Future work will continue to examine how to best make use of relative amplitudes between SPRED signals.

Since the SPRED line of sight crosses a beam path near the magnetic axis, it is expected that charge exchange (CX) between impurity ions and beam neutrals may produce some lines within the detectors' spectral range. In Fig. 8-4 we compare the time history of the $F^{8+} n = 10 \rightarrow 9$ line intensity measured by the CER T01 chord with the SPRED signal near 230 \AA , identified as deriving from the $F^{8+} n = 4 \rightarrow 3$ transition. Note that while these transitions are formally from F^{8+} , their emission is indicative of F^{9+} densities, since it is the F^{9+} charge state that recombines via CX and populates the upper level of these transitions. To compare the two measurements, the SPRED intensity was rescaled to match the F^{9+} density measured by CER. The two time histories match very well as a result of the good spatial localization of the individual lines of sight with the NBI beam at a near-axis crossing point.




Figure 8-4: Comparison of SPRED measurements at $\lambda \approx 230 \text{ \AA}$ with the CER signal at the inner-most chord, T01, of discharge #180526.

EUV lines at 150 \AA and 222 \AA observed after LiF injections were also originally interpreted to be due to F^{8+} transitions, but their time histories were found to be impossible to match with Aurora simulations: their long-lived emissivity in the plasma cannot be accounted for by a charge state that is not fully-stripped. It seems likely that this emission

Label	Measured λ [Å]	NIST λ [Å]	ADAS PEC λ [Å] range	Comments
F SPRED-0	464	464.364	464-467	F V
		465.375		F V
		465.975		F V
		466.993		F V
F SPRED-1	656	654.028	654-658	F V
		657.228		F V
		657.332		F V
F SPRED-2	645	647.67	647-648	F V
		647.77		F V
		648.0?		F VIII?
F SPRED-3	128	127.655	126.5-129	F VII
		127.800		F VII
		127.815		F VII
		?		F VI; 127.835 Å (ADAS)
		126.532		F VI
F SPRED-4	882	883.111	876-878	F VII
		885.14		F VII
F SPRED-5	890	885.14	884-885	F VII

Table 8.1: VUV line identification used for the analysis of LBO injections of F on DIII-D.

is in fact mostly the result of second-order diffraction at the SPRED grating of F^{8+} lines near 75 Å and 111 Å whose upper levels are populated via CX. Based on Rydberg scaling of transitions from other ions available from OPEN-ADAS (using Eq. 3.10), we estimate that there exist F^{8+} transitions near 77 Å ($n = 6 \rightarrow 4$), 80 Å ($n = 7 \rightarrow 5$), and 81 Å ($n = 3 \rightarrow 2$). For $Z = 9$, the simple scaling of Ref. [272] for the upper levels that are favored by CX with ground state neutrals gives $n \approx \sqrt{Z^3/(3Z - 2)} = 5.4$ (the more common $Z^{3/4}$ approximation [115] gives ≈ 5.2). This suggests that the CX lines at 77 Å and 80 Å should be very bright. While these wavelengths are out of the SPRED range, their second order (at twice the original wavelength) is likely to give the observed brightness near 150 Å. Similarly, the emission measured near 222 Å could be receiving contributions from a F^{8+} $n = 4 \rightarrow 3$ transition near 108 Å. Given the significant uncertainty involved in modeling these lines, especially in view of other low- n contributions from F^{8+} to the measured brightnesses, we have chosen to exclude these two lines from our inferences of particle transport.

The line groups in Table 8.1 show a selection of atomic lines that was observed on SPRED following LBO injections of F. We compare the measured line center, the identification of relevant lines in the NIST ASD [136] database, and the range of wavelengths within which

Label	Measured λ [Å]	NIST λ [Å]	ADAS PEC λ [Å] range	Comments
Al SPRED-0	282	~ 50 lines	270-288	Al VIII, IX, XII
Al SPRED-1	331	~ 50 lines	320-342.5	Al VII, VIII, IX, X, XI
Al SPRED-2	665	?	663-665	Al X at 664.91 Å (ADAS) Al V and VIII lines in NIST
Al SPRED-3	549	550.03	547-548	Al XI
Al SPRED-4	567	?	564-566	Al X at 565.03 Å (ADAS) Al XI at 565.39 Å (ADAS)

Table 8.2: VUV line identification used for the analysis of LBO injections of Al on DIII-D.

the ADAS PECs were integrated for forward modeling. Many more lines are observed in the SPRED spectral range; the selection above only describes the brightest and most easily identifiable lines that could be used to constrain the density of multiple charge states in our inferences of impurity transport. Table 8.2 offers an analogous list of observations used for Al injections. The emission near 282 Å and 331 Å is found to likely correspond to a large number of lines, according to NIST. ADAS PECs suggest that several charge states contribute to observations within the indicated wavelength ranges, with stages as low as Al⁶⁺ giving non-negligible contributions. Al lines at 665 Å and 567 Å are accounted for in ADAS, but are not listed by NIST. The most reliable measurement of Table 8.2 is likely the one at 549 Å, since this has a single identification (Al¹⁰⁺) according to both NIST and ADAS.

SXR Signals

We use line-integrated SXR brightness, rather than local emissivities from tomographic inversions, to avoid use of multiple SXR arrays, whose relative calibration has been observed to be inaccurate. The data used for the analysis in this chapter were filtered by a ≈ 125 μm -thick Be filter. These filters lay flat in front of photodiode arrays, meaning that different lines of sight go through a slightly different Be thickness. The inaccuracy incurred by this setup is assumed to be acceptable for the results presented here, but this issue should be addressed in future work.

As discussed in Section 7.4, the significant uncertainties that are intrinsic to forward modeling of SXR diagnostics have made analogous SXR modeling on C-Mod an intractable problem. On DIII-D, we choose to make use of only a single array at a time and use data only from the initial signal rise after LBO injections. This choice is motivated by

the observations that other diagnostics with lower spatio-temporal resolution but relying on more validated atomic physics can constrain the decay phase well, while they may not have sufficient resolution to constrain the rise phase. By using SXR signals only from the rise phase, we attempt to reduce over-reliance on the SXR atomic data accuracy.

Background Neutral Density Predictions

As discussed in Section 6.5, the effect of NBI neutrals on impurity transport inferences is relatively small; nonetheless, it is included here using FIDASIM [201] Monte Carlo predictions for both fast and thermal neutrals. The RMP ELM-suppressed discharge of interest, #175674, had only one NBI beam (30L) operating at the time of the Al LBO injection (2.5 s). The negative triangularity discharges had 1, 2, and 3 active beams in shots #180520, #180526 and #180530, respectively (see Fig. 8-10).

Having observed a good match of edge D neutral densities from experimental Ly_α spectroscopy and SOLPS-ITER for 3 C-Mod conditions (Fig. 6-3), we take SOLPS-ITER to be a high-fidelity model to account for pedestal CX within the Aurora forward model. However, at the time of writing SOLPS-ITER results for the RMP ELM-suppressed discharge #175674 are not yet available. In their absence, we resort to the *Kinetic Neutral 1D* (KN1D) code [200], a lower-fidelity model that has been shown in previous work to provide a good description of neutral behavior in the SOL and pedestal of Alcator C-Mod [159]. KN1D is a fast and open-source tool that has been integrated within Aurora. An OMFIT module has also been created to facilitate KN1D modeling on both C-Mod and DIII-D, permitting a simple interface with ImpRad for inferences of impurity transport. Appendix F summarizes features of the code and its application in our analysis. We note that while KN1D appears to be a good description of neutral penetration in a 1D slab, appropriate for C-Mod main-chamber recycling conditions, its model has not been thoroughly validated for DIII-D divertor-recycling conditions [273]. Moreover, KN1D makes no attempt at quantifying the absolute amplitude of neutral densities, but only their radial profiles. In practice, this amplitude is controlled via a code input for the neutral D_2 pressure at the midplane wall.² For the analysis of shot #175674, we set this value to be a free parameter in MultiNest, with a wide prior that ranges between approximately 1 μTorr to

²The D_2 pressure at the wall can be measured on DIII-D via the midplane “ASDEX gauges”. Unfortunately, these were not calibrated to provide useful measurements in the discharges of interest.

100 μTorr with a log-normal distribution of mean 10 μTorr .

For the analysis of the negative triangularity discharges, we make use of SOLPS-ITER results provided by J.D. Lore for discharges #180520 and #180533. The latter shot is not one of those for which we have good LBO injections, but it is a repeat discharge from the same set of experiments. SOLPS-ITER results for #180520 were obtained to match pedestal and SOL measurements from Thomson Scattering at 2.425 s (very close to the time of the LBO injection of interest), whereas those for #180533 were for the condition at 2.3 s. The NBI power ramp in #180533 was identical to those in the L-mode $\delta < 0$ cases analyzed here (#180526 and #180530, *c.f.* Fig. 8-10) and all other parameters were very similar between these shots. Hence, these SOLPS-ITER results are likely relatively accurate for the low-power L-mode condition (#180526), for which the LBO injection was at 2.75 s, but larger discrepancies may exist for the high power case (#180530), since the LBO injection in this case was at 3.8 s (at higher NBI power). As for the KN1D runs of the RMP ELM-suppressed H-mode shot (#175674), we use in all cases a free parameter in our inferences that allows rescaling of the entire atomic D neutral density profile from SOLPS-ITER to best match spectroscopic data, using a log-normal prior with mean of 1 and width of 0.5, permitting corrections of neutral density by more than a factor of 3.

8.3 Impurity Transport in DIII-D Discharges with RMPs

The high-confinement (H-) mode has long been recognized [274] to lead to greater performance than L-mode operation, but also to suffer from a number of important issues. First, H-mode discharges have a strongly unfavorable power scaling of confinement, meaning that performance does not grow significantly as one adds extra heating. Secondly, H-modes tend to confine impurities, sometimes leading to radiative collapse unless measures are taken to counter this effect. Thirdly, the steepening of pressure gradients and bootstrap currents in the pedestal often produces Edge-Localized Modes (ELMs), which threaten the survival of future devices due to their release of large amounts of energy to the wall and divertor surfaces. Resonant Magnetic Perturbations [275] (RMPs) have been used to mitigate or completely eliminate ELMs by coupling non-axisymmetric magnetic fields from external coils with resonant flux surfaces. This has been observed to create stochastic field regions near the edge and reduce pedestal gradients, providing an external actuator that can keep

the plasma away from stability limits related to ELMs. Of course, RMPs achieve this at the cost of sacrificing some of the plasma confinement quality that makes H-modes attractive in the first place. Ideally, such techniques that actively reduce overall confinement would not be needed if one could operate in high-performance regimes that are intrinsically free of ELMs, such as the I-mode and EDA H-mode regimes on C-Mod (Chapter 7) or the Quiescent H-mode (QH-mode) [276, 277] and negative triangularity scenarios (Section 8.4) on DIII-D. However, it remains unclear whether these more favorable regimes will be accessible in future devices such as ITER or SPARC. It is therefore of interest to develop and understand the impact of techniques such as RMP ELM-suppression that could mitigate any unwanted plasma phenomena from negatively impacting operation.

In Ref. [270], Victor *et al.* investigated the effect of RMPs on pedestal impurity transport in DIII-D using LBO injections of W and Al. These experiments clearly showed a reduction of impurity accumulation in high-performance discharges in ITER-similar plasma shapes (at positive triangularity). Using the STRAHL code, this effect was interpreted as the result of a large increase of effective pedestal diffusion due to RMPs. Ref. [270] notes that if transport were affected by RMPs via the creation of magnetic islands, a diffusive-convective description of transport may not be helpful in transport model validation. However, experimental observations on DIII-D have shown that pedestal density fluctuations, \tilde{n}/n , correlate with background density scale lengths, L_{ne} , suggesting that a gradient-driven flux description still applies in scenarios with RMP ELM suppression [278]. In this section, we extend the analysis of Victor *et al.* [270] for a single discharge, #175674, where RMPs successfully suppressed ELMs and trace amounts of Al were introduced via LBO at 2.5 s. Readers are referred to Ref. [270] for a detailed description of these experiments and collected data.

Our comparison to these previous results serves multiple purposes. First, we wish to demonstrate the application of ImpRad on a scenario that has been previously analyzed, before moving on to the more exotic plasmas discussed in Section 8.4; this should not be considered a *benchmark*, since our methods should not necessarily recover previous results but rather improve upon them. Secondly, we compare for the first time TGLF and NEO modeling of this discharge across the plasma radius. Analogous analysis will be demonstrated in Section 8.4.1 for negative triangularity discharges.

8.3.1 Inferences of Impurity Transport

Chapter 6 described the importance of charge exchange (CX) of background neutrals and impurities. Fig. 6-16 shows, based on neutral density predictions from SOLPS-ITER (Fig. 6-17), that this effect can radically modify charge state balance in the pedestal region. In view of this, results from previous work on pedestal impurity transport that did not specifically include CX in transport modeling are likely in need of reconsideration. This includes the work of Ref. [270], which we revisit here.




Figure 8-5: Inferred Al (green) and C (blue) transport coefficients for DIII-D shot #175674. The left column shows diffusion (assumed to be the same for the two species) and convection. The right column shows their normalized ratio, Rv/D , and the normalized cumulative integral of v/D , which for a fully-stripped impurity offers an approximate visualization of source-free steady-state profile shapes.

Fig. 8-6 shows the result of our inference of particle transport coefficients for both C (blue) and Al (green) in the DIII-D discharge #175674. The top-left panel shows D (assumed to be the same for the two species), the lower-left v , the right-top their ratio v/D multiplied by the major radius R , and the bottom-right a prediction for the normalized steady-state profile of total impurity density, obtained via a cumulative radial integral of the v/D profile.

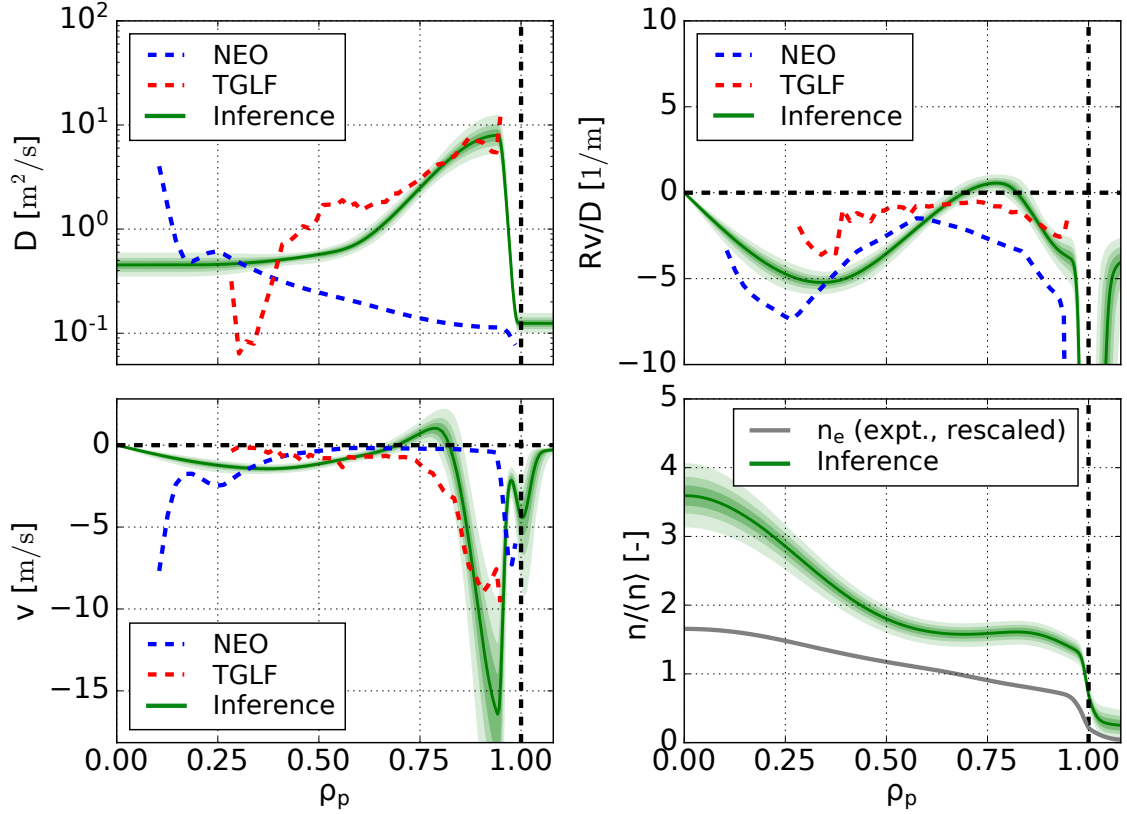


Figure 8-6: Comparison of Al transport coefficients from an experimental inference (green in Fig. 8-5) and from theoretical modeling for DIII-D shot #175674. Results from the NEO (neoclassical) and TGLF (quasi-linear gyrofluid turbulent) codes are shown.

The inferred steady-state profile should only be interpreted as a rough indication of profile shapes for fully-stripped impurities that are free of sources and sinks. Its main purpose is to compare to the experimentally-measured n_e profiles, as we will discuss later. The main difference observed between the C and Al results is that Al appears to accumulate more strongly near the center. We remark that the relation between C and Al values of Rv/D near the plasma center is not enforced by our inference setup: it is the combination of C and Al data themselves that suggest the greater peaking of steady-state Al profiles. This is shown in the top-right panel to correspond to a Rv/D value at $\rho_p \approx 0.3$ that is almost twice as large for Al as for C. This can be interpreted as the result of a neoclassical scaling ($v \propto Z$) of near-axis accumulation. Analogous considerations have been made for all the discharges presented in this chapter. In what comes next, we focus our attention on the Al results, which we compare explicitly to theoretical transport models.

Fig. 8-6 shows the same Al results as Fig. 8-5, now overlapped with neoclassical NEO


and turbulence (quasilinear, gyrofluid) TGLF predictions. The lower-right panel compares the steady-state profile of fully-stripped Al with the normalized electron density profile. The most salient feature in this comparison is that the Al profile presents larger core gradients as a result of the central accumulation already remarked in Fig. 8-5. Section 8.3.2 discusses theoretical predictions for this scenario.

The discharge presented here has excellent CER measurements with respect to most DIII-D datasets for LBO studies. The data discussed here has high signal-to-noise ratios and there is wide diagnostic coverage in the core plasma. Fig. 8-7 shows the match between these data and the result of forward modeling with Aurora using the inferred Maximum A Posteriori (MAP) estimates for the free parameters. Fig. 8-7a shows the match of time histories for Al¹³⁺ density, with different panels distinguishing data from each CER chord, whose radial measurement location is indicated above every signal. Fig. 8-7b shows radial profiles of C⁶⁺ density, with different panels corresponding to binning in time. This plasma scenario exhibited impurity accumulation and therefore the C density in Fig. 8-7b is slowly increasing over time, though this is not visible over the short time range displayed here.


8.3.2 Comparison of Experimental and Theoretical Modeling

Fig. 8-6 shows NEO and TGLF results on top of the experimentally-inferred profiles. NEO results are taken to be mostly relevant in the radial ranges $0.1 \lesssim \rho_p \lesssim 0.25$ and $0.9 \lesssim \rho_p \lesssim 0.99$; TGLF is most applicable in between these regions ($0.25 \lesssim \rho_p \lesssim 0.9$). The inner-most radius for which NEO results are shown is dictated by the size of Al potato orbits near the magnetic axis, calculated to be $\approx 0.05 \rho_p$ units; we take twice this radius to be a limit of applicability for NEO modeling. Similarly, we find that ion orbit losses may be significant outside of $\rho_p \approx 0.99$ and we therefore only consider NEO to be valid inside of this radius.

Theoretical modeling results have been computed using kinetic profiles obtained with the TGYRO code [52]. TGYRO is a framework where a transport model is run iteratively to find T_i , T_e , and n_e (and possibly rotation frequency) values that can match experimental fluxes, modifying the profiles at each step. In this case, TGYRO was run for $0.2 \leq \rho_p \leq 0.9$ with 7 radial knots, and in this region the heat and particle fluxes are mostly the result of turbulence predicted by TGLF, except for the ion heat flux, Q_i , for which neoclassical results from NEO are quite significant. The experimental fluxes that TGYRO attempts to match are estimated via TRANSP [256] using a dedicated OMFIT module [279]. Fig. 8-8



(a)



(b)

Figure 8-7: Match of experimental and modelled CER signals in the DIII-D RMP ELM-suppressed discharge #175674 for (a) time-dependent Al injected via LBO, and (b) intrinsic, time-independent C. In (a), black points with error bar represent measurements and red lines are the final forward model reconstructions. In (b), experimental data is shown as points with error bars, whereas lines represent the forward model result for the displayed time slices.

shows a summary of TGYRO results and comparisons to experimental profiles, displayed as a function of the square root of normalized toroidal flux, ρ_ϕ .




Figure 8-8: Summary of TGYRO results for DIII-D shot #175674, showing kinetic profile reconstructions on the left (for T_i , T_e , and n_e) and heat fluxes on the right (Q_i , Q_e and $Q_{e,conv}$). In each panel, black lines display the experimental target, red lines the converged TGYRO result.

In each of the three figures, kinetic profile reconstructions are shown on the left, with black dashed lines showing the fit to experimental data and red lines showing the result from TGYRO. Right hand side panels display the match between experimental heat fluxes (from TRANSP [256] power balance, black dashed lines) and TGYRO outputs (red). The lower-bottom panel for each discharge shows the convective heat flux, defined as $Q_{e,conv} = 1.5\Gamma_e T_e$. The ability of NEO and TGLF in matching this quantity is indicative of good predictive capabilities for particle transport, at least for electrons. In practice, since TGYRO is run only over a finite number of knots (shown by dots), the actual heat flux targets (black continuous lines) are slightly different from the results of power balance (black dashed lines) due to the linear integration between knots.

Unlike in C-Mod research discussed in Chapter 7, all TGLF simulations discussed in this

chapter use the TGLF SAT-0 saturation rule, rather than the SAT-1 rule, since this has been found to give the best convergence properties in TGYRO, particularly for the negative triangularity cases discussed in Section 8.4. Fig. 8-8 shows that, even with SAT-0, TGYRO is unable to match T_i very well at midradius, while it is more successful in T_e and n_e . We note that profiles inside of $\rho_\phi = 0.2$ are not freely evolved and one should not expect good matches inside of this location due to the presence of MHD activity leading to increased effective heat fluxes.

The comparison of D and v from experimental and theoretical modeling (based on profiles from TGYRO, red in Fig. 8-8) in Fig. 8-6 shows relatively good agreement in convection and v/D profiles, with more discrepancies appearing in the D profile. As discussed in Ref. [270], in this discharge the application of RMPs did not prevent impurity accumulation over time, as evidenced by C^{6+} profiles from CER and in agreement with the peaked steady-state impurity profiles (bottom-right panel in Fig. 8-6) from our inference, corresponding to strongly negative v/D near axis. In the outer part of the plasma, impurity profiles are only weakly peaked, approximately in agreement with TGLF. This is in clear contrast to the v and v/D profiles shown in Fig. 14 of Ref. [270], where hollow Al profiles were observed. In our analysis, the constraint offered by C profiles clearly prevents such hollowness from being inferred. In the pedestal, a strong inward pinch ($v < 0$) is experimentally inferred, bringing the experimental results closer to the theoretical predictions. As one goes from $\rho_p = 0.75$ to $\rho_p \approx 0.95$, TGLF suggests that the impurity profile peaking should increase. The experimental inference is in qualitative agreement with this, although v is inferred to be slightly positive near $\rho_p \approx 0.8$. Between $\rho_p = 0.9$ and $\rho_p = 0.95$, NEO predicts a relatively small inward pinch, with a maximum amplitude of $v \approx 7$ m/s, which combined with the neoclassical prediction of $D \lesssim 0.1$ m²/s gives a large $v/D < -50$ m⁻¹. The corresponding sharp Al density pedestal is found to be inside of the LCFS, but this result could be affected by unwanted inaccuracies in the localization of the separatrix in the fitting of experimental kinetic profiles. Nonetheless, we note the broad agreement in the overall pedestal structure between the experimental inference and theoretical modeling. The most significant disagreement that we find is in the amplitude of D at midradius, where CER data constrain D to be significantly smaller than suggested by TGLF.

8.4 Diverted Negative Triangularity

Section 8.3 described the potential of RMPs as an external actuator to control ELMs in high-performance discharges. The experiments discussed by Victor *et al.* [270] showed that RMPs reduce the impurity accumulation that can be observed in ITER-similar plasma shapes, at positive triangularity. However, as discussed above, it would be desirable to (a) not have ELMs in the first place, rather than actively suppress them, and (b) achieve a clearer decoupling of heat and particle transport, such that high plasma core pressures may be achieved without running into impurity accumulation.

Among the options being explored within the fusion research community, operation at negative triangularity is certainly one of the most interesting. Triangularity, δ , is a measure of the magnetic geometry inboard or outboard bias. It can be calculated as $\delta_u = (R_{\text{geo}} - R_u)/a$ for the upper part of the plasma and similarly for the lower part. The overall triangularity is the mean of the two. Here R_{geo} is the geometric mean of the midplane major radii of a given flux surface on the Low-Field Side (LFS) and High-Field Side (HFS); R_u (R_l) is the major radius of the point with highest Z at the top (bottom) of the plasma. Typically the triangularity is reported for the Last Closed Flux Surface (LCFS), which presents the strongest shaping (non-circular geometry) of any closed flux surface, but δ can in principle be computed anywhere inside the confined plasma region. Together with the elongation, κ , triangularity enters the standard Miller parametrization of flux surface shapes, expressed as

$$\begin{aligned} R(\theta) &= R_0 + r \cos[\theta + \arcsin(\delta) \sin(\theta)] \\ Z(\theta) &= Z_0 + \kappa r \sin(\theta) \end{aligned} \tag{8.1}$$

where θ is the poloidal angle, r is the minor radius and (R_0, Z_0) are the coordinates of each flux surface’s geometric center. In practice, in this work we make use of full Fourier decomposition of flux surface geometries for theoretical modeling, computing δ from the Fourier coefficients themselves.

Fig. 8-9 contrasts “matched” plasma geometries with negative (blue) and positive (red) triangularity. The negative triangularity case corresponds to DIII-D discharge #180526, discussed in detail below, while the other is simply a fictitious reflection along a vertical axis (not a real experimental equilibrium). This diagram illustrates how the inversion of triangularity affects the path of a trapped electron on the LFS, shown by the green lines.

In the $\delta > 0$ case (red), trapped electrons mostly spend time in regions of “bad” magnetic curvature. Here, gradients of pressure are in the opposite direction to gradients of the toroidal magnetic field, the latter decreasing radially as $B \propto 1/R$. On the other hand, in the $\delta < 0$ case, trapped particles may spend relatively more time in regions of “good” curvature. Given that bad-curvature instabilities such as Ion Temperature Gradient (ITG), Electron Temperature Gradient (ETG), and Trapped Electron Modes (TEM) are strongest on the low-field side, Fig. 8-9 suggests a heuristic picture for differences of transport at positive and negative triangularity. Gyrokinetic modeling has previously been used to understand lower fluctuation levels at $\delta < 0$ as being the result of TEM stabilization [280].




Figure 8-9: Heuristic picture of trapped particle motion in negative (blue) and positive (red) triangularity. Illustration inspired by Ref. [281].

Operation at negative triangularity was originally considered in the 1980s to be unfavorable, mostly because of the expectation of worse ballooning stability at the edge and therefore lower confinement. Recently, experiments on the TCV and DIII-D tokamaks have re-examined the negative triangularity concept and found experiments to contradict previous conjectures. On TCV, H-mode grade confinement was obtained in purely electron-heated L-mode discharges that were inner-wall limited (IWL), finding significantly reduced levels of turbulent losses and fluctuations intensity [282]. These results prompted researchers on DIII-D to explore this regime on a larger device where use of Neutral Beam Injection

(NBI) could produce a reactor-relevant T_i/T_e ratio close to 1 [283, 284]. It was found that high confinement with $H_{98y2} = 1.2$ and normalized plasma beta of $\beta_N = 2.7$ could be obtained. Here, H_{98y2} indicates the international H-mode heat confinement scaling (a value greater than 1 indicates performance better than standard H-modes) and $\beta_N = \beta(a B_T)/I_p$ is an important global metric for plasma stability, with $\beta = p/(B^2/2\mu_0)$. A remarkable feature of $\delta < 0$ discharges is that they do not easily transition into H-mode and hence they do not form a pressure pedestal, which has been observed to correlate with lack of impurity accumulation.

On both TCV and DIII-D, confinement improvements in IWL experiments have been interpreted as resulting from greater TEM stability for $\delta < 0$ compared to $\delta > 0$ [280, 285]. DIII-D has also recently been able to produce diverted negative triangularity discharges like the blue one in Fig. 8-9, where the upper triangularity is approximately $\delta_u = -0.4$ and the lower triangularity is close to $\delta_l = 0$, giving an overall (mean) negative triangularity of $\delta \approx -0.2$. Section 8.4.1 describes these experiments in greater detail. Creation of a diverted $\delta < 0$ discharge with both $\delta_u < 0$ and $\delta_l < 0$ is challenging for DIII-D because the device is not appropriately armored on the outer wall to withstand high heat fluxes at the strike points; work is currently on-going to overcome this difficulty.

8.4.1 DIII-D Experiments in Diverted Negative Triangularity

Fig. 8-10 shows time traces of some of the most important quantities describing the three discharges discussed here (#180520, #180526, #180530). While the plasma current and Electron Cyclotron Heating (ECH) power were kept constant throughout most of the discharge at $I_p \approx 0.9$ MA and $P_{ECH} \approx 1.5$ MW, the Neutral Beam Injection (NBI) power was ramped up in discrete steps, reaching $P_{NBI} \approx 13$ MW in two shots. The purpose of these incremental steps was to test the L-H power threshold and β stability in diverted negative triangularity. No β limit was encountered even when reaching $\beta_N \approx 3.1$, in agreement with recent theoretical predictions [286]. Core electron densities and temperatures of up to $5 \times 10^{19} \text{ m}^{-3}$ and 4.5 keV, respectively, were achieved, with $H_{98,y2}$ factors above 1, i.e. exceeding standard H-mode energy confinement. Throughout these experiments, the H-mode transition was only observed in a single case, when one of the upper coils failed and the upper triangularity changed from -0.4 to -0.2 . This appeared to significantly reduce the L-H power threshold and allow the development of a density pedestal. Shot #180520 (red

in Fig. 8-10) is the discharge where this occurred, near 2.2 s, and indeed the core density is seen to rapidly increase at that time and then stabilize at a higher value. Note that density, particularly in shot #180530, is also affected by core fueling via NBI, particularly as P_{NBI} is ramped up. In discharge #180520, the H-mode transition also resulted in the occurrence of small ELMs, although the filterscope trace shown in the top-center panel of Fig. 8-10 does not make this obvious. The filterscope signals generally increase as D_α emissivity near the LCFS rises, indicating a higher degree of recycling from the edge. Finally, we note that the energy confinement time of these discharges was near 70-120 ms (higher for #180520, after the transition to H-mode).

These experiments show great promise for negative triangularity as the operational regime for a fusion reactor: not only did plasmas not meet a “hard” MHD limit, even up to $\beta_N \approx 3.1$, but they demonstrated excellent core properties. The achievement of $H_{98,y2} > 1$ reflects the fact that energy confinement degraded with external heating less strongly than

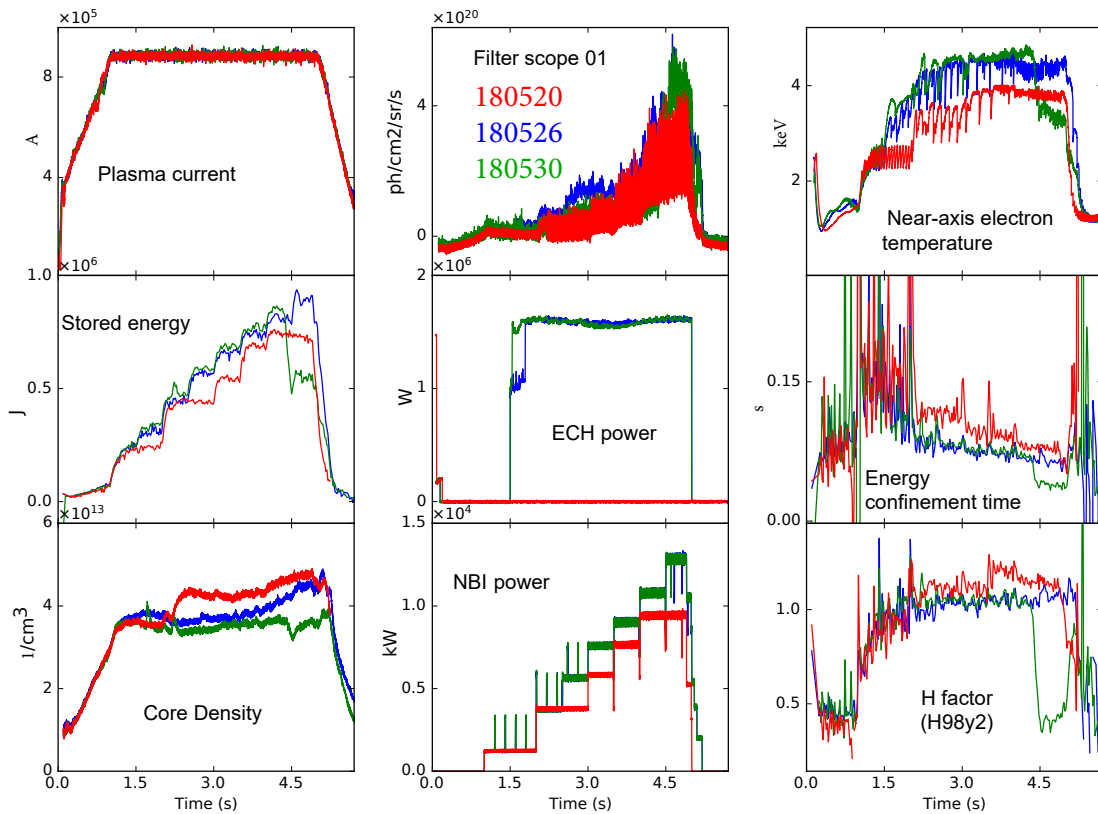


Figure 8-10: Evolution and comparison of the 3 DIII-D diverted negative triangularity discharges discussed in the text.

in standard H-modes described by the IPB98(y,2) scaling. This is all the more remarkable given that plasmas did not actually enter H-mode when δ remained sufficiently negative, thus avoiding the liability of an H-mode pedestal, which normally leads to higher power being deposited at the outer divertor target. At negative triangularity, this outer strike point is naturally located further out in radius with respect to positive triangularity. This constitutes a promising demonstration of the “edge-first” principle, which prioritizes sustainable divertor conditions before attempting to optimize core stability and confinement. In this context, it is obviously also important to assess if diverted negative triangularity discharges have a tendency to accumulate impurities in the core, since that would be a fundamental setback. In order to explore this, Laser Blow-Off (LBO) injections were made in these experiments. In particular, here we examine the following in detail:

- An Al injection at $t = 2.501$ s in #180520, when only the 30L beam was operating, approximately 500 ms after the discharge transitioned into H-mode.
- An LiF injection at 2.75 s in #180526, when both the 30L and 33L beams were operating. Only F ions from this injection were measured by the CER diagnostic and are therefore examined in detail.
- An Al injection at $t = 3.8$ s in #180530, when the 30L, 33L, and 33R beams were all operational and high stored energies were achieved.

Fig. 8-11 shows kinetic profiles for each of these discharges, averaged over ± 50 ms time windows near the times of LBO injections for each of the shots listed above. Individual experimental data points are over-plotted in each panel, except when the profiles are inferred from line-integrated measurements. Comparing to the time traces in Fig. 8-10, we note that different NBI power is applied in each of these 3 time slices and this causes T_e (top-left panel) and T_i (bottom-left) to vary significantly. Electron density (top-center) is highest in the discharge that entered H-mode, #180520. The plasma is seen to have a high rotation frequency (ω_ϕ) in #180530, when the NBI torque is near its maximum. Peaked toroidal rotation and ion temperature could both be contributors to the evidently hollow C profiles (bottom-center) seen in this case. Z_{eff} remains low, near 1.5, in all 3 discharges, suggesting fast (favorable) impurity transport.

During these experiments, some issues with the LBO system prevented consistent injections of the same ions in all discharges and at multiple times, making it difficult




Figure 8-11: Kinetic profiles for the 3 diverted negative triangularity discharges discussed in this chapter, averaged over a short time window near the LBO injection time: 2.501 s for #180520, 2.75 s for #180526, and 3.8 s for #180530.

to clearly assess scalings of confinement and transport coefficients during the NBI power ramps. Nonetheless, it was clearly observed that these discharges had impurity confinement times that were similar to (or shorter than) the energy confinement time. For positive triangularity, this is typically the case in L-mode, but not in H-mode. Hence, we find that negative triangularity results in a decoupling of heat and particle transport, giving high core temperatures and rapid impurity expulsion, also reminiscent of the I-mode regime [23, 252, 287].

Fig. 8-12 shows signals from the CER T06 chord (a) #180520 (measuring Al^{13+}) and (b) #180526 (measuring F^{9+}) in the core. Exponential fits to the decay phase allow quantification of the impurity confinement time, shown in the top-right corner. While these values should not be considered a truly global quantity and are only a rough measure of particle transport, they do offer the clear conclusion that impurity transport is fast and therefore favorable in these discharges. This is best observed by comparing values of $\tau_p \lesssim 50$ ms with the larger energy confinement times in the same discharges, $\tau_e \approx 70\text{--}100$ ms (*c.f.* Fig. 8-10). In the next sections we shall delve more deeply into the matter by inferring radial profiles of transport coefficients for the 3 LBO injections listed above.




Figure 8-12: Signals from the CER T06 chord in (a) #180520 (Al) and (b) #180526 (F), showing exponential fits in the decay phase to quantify impurity confinement times, τ_p , which in both cases are smaller than 50 ms. For comparison, the energy confinement times are ≈ 80 and 120 ms at times corresponding to LBO injections in #180526 and #180520, respectively.

8.4.2 Impurity Transport Inferences

Fig. 8-13 compares the simultaneous inference of C and F transport coefficients near $t = 2.75$ s in discharge #180526, which we shall refer to as the “low-power L-mode case”. Since this discharge had a single sawtooth occurring during the time range of evolution of the LBO injection, at $t \approx 2.784$ s, we also display the sawtooth mixing radius with thick, magenta, dashed lines, having estimated this location using fast Electron Cyclotron Emission (ECE) measurements [162]. Thinner magenta dash-dot lines in each panel represent the sawtooth *inversion* radius, approximated to be the mixing radius divided by $\sqrt{2}$ [63]. As in Fig. 8-5, we show only a single D radial profile (top-left panel) since this was set to be the same for the two impurities. The Rv/D ratio was allowed to vary between the two ions in the region inside of the sawtooth mixing radius, and it is indeed found to be different, even changing sign, between F (green) and C (blue) lines. The inferred pedestal pinches for the two species are not substantially different.

Fig. 8-14 compares the inferred F transport coefficients for this shot (#180526) with the results of TGLF and NEO modelling. While a detailed comparison to theoretical models is left for Section 8.4.3, here we remark that steady-state F profiles (bottom-right panel) have gradients that are approximately similar to the rescaled time-averaged n_e profile (shown in grey), except inside of the sawtooth mixing radius (magenta dashed lines). The C^{6+} density




Figure 8-13: Inferred C (blue) and F (green) transport coefficients for DIII-D shot #180526. The thick, magenta, dashed lines show the sawtooth mixing radius; the thinner dash-dotted lines represent the sawtooth inversion radius.

profile with which this inference was constrained (see Fig. 8-11, bottom-center panel), also shows a profile shape that significantly differs from the n_e one inside of $\rho_p \approx 0.3$.

We note that the CER F^{9+} measurements of this discharge display a surprisingly fast rise near the magnetic axis that cannot be easily matched without invoking some degree of signal line-integration, missing atomic physics or an MHD mode (like a sawtooth crash). No evidence has been found for the latter hypothesis on other diagnostics. Lack of radial localization could arise from a background signal due to charge exchange of F^{9+} with edge neutrals rather than NBI neutrals, but attempts at background subtraction using a beam blip have suggested that this is unlikely to be an appropriate explanation. More likely, the fast rise of F^{9+} density near the magnetic axis arises because of inaccuracies in the beam attenuation calculation, or missing contributions from beam halos. The latter have been observed to be very important on AUG [173], but are believed to be less important at the higher- Z_{eff} conditions of DIII-D. Given that the present negative triangularity

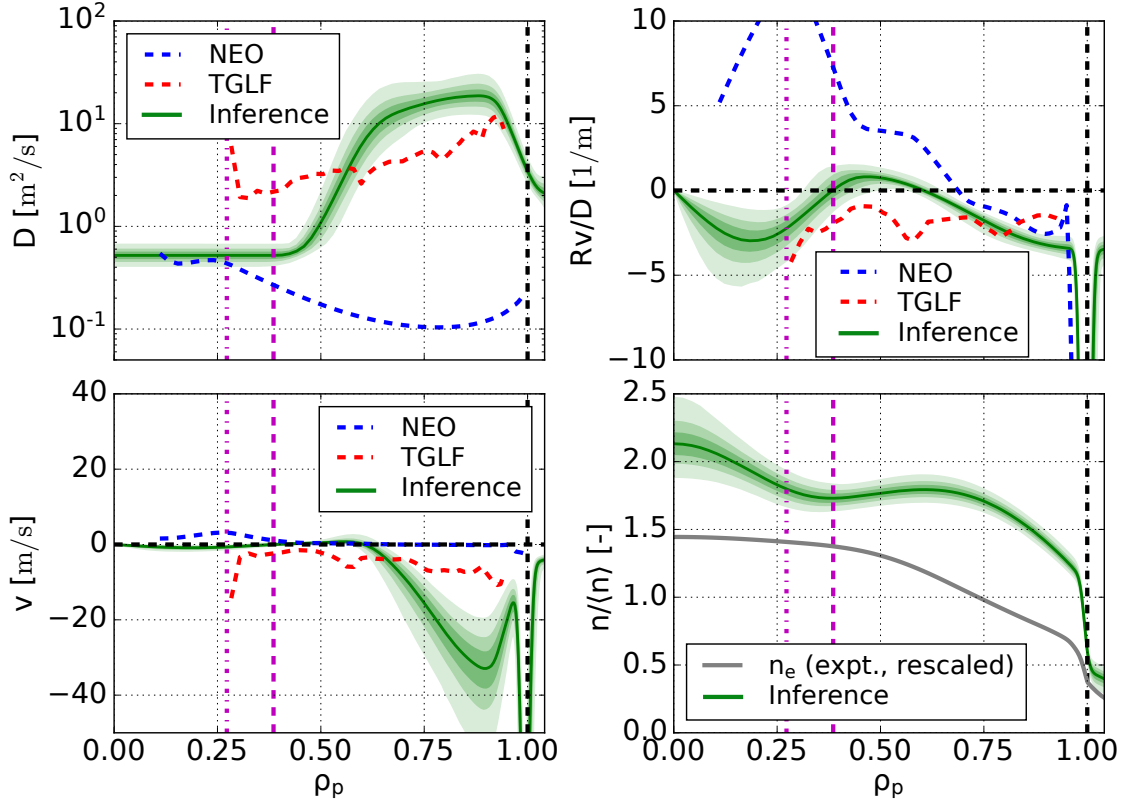


Figure 8-14: Inferred F transport coefficients for DIII-D shot #180526. The left column shows diffusion and convection, the right column their normalized ratio, Rv/D , and the approximate normalized steady-state profile of fully-stripped F. The rescaled time-averaged n_e profile is also shown for comparison. Results from the NEO (neoclassical) and TGLF (quasi-linear gyrofluid turbulent) codes are also shown, with thicker lines representing expected regions of applicability of each model.

discharges had $Z_{eff} \approx 1.5$, the halo $n = 2$ contributions to density estimations may not be negligible. Future work should assess whether FIDASIM simulations such as those discussed in Chapter 6 can account for our observations. At the time of writing, finding a clear resolution to these issues appears unlikely, but does not prevent a close examination of inferred transport coefficients, particularly for $\rho_p > 0.2$.

Fig. 8-15 and Fig. 8-16 show inferred transport coefficients for the H-mode (shot #180520, $t \approx 2.5$ s) and high-power L-mode (shot #180530, $t \approx 3.8$ s) cases, respectively. Both of these inferences are based on LBO injections of Al, rather than F. The presence of an Al^{10+} CX ($n = 14 - 12$) line at 4088 \AA , almost degenerate with the CER Al^{12+} ($n = 12 - 11$) line of interest at 4084 \AA , raises questions about whether signal measured by the CER diagnostic can be reliably taken to relate solely to the Al^{13+} density. Since no

effective subtraction methods for the Al^{10+} contribution were available, data points during the rise phase of the CER signal were ignored in our inferences. Also, having observed a strong dependence of near-edge transport coefficients on details of the source functions, we choose to only focus on profiles inside of $\rho_p \approx 0.8$, particularly since the fast camera data (Fig. 8.2) suggest the presence of Al LBO clusters in these experiments.

Fig. 8-15 and Fig. 8-16 display relatively flat or even hollow impurity density profiles, in agreement with the C profiles of Fig. 8-11. Such profile shape does not appear to be captured by TGLF and NEO, as further discussed in Section 8.4.3.

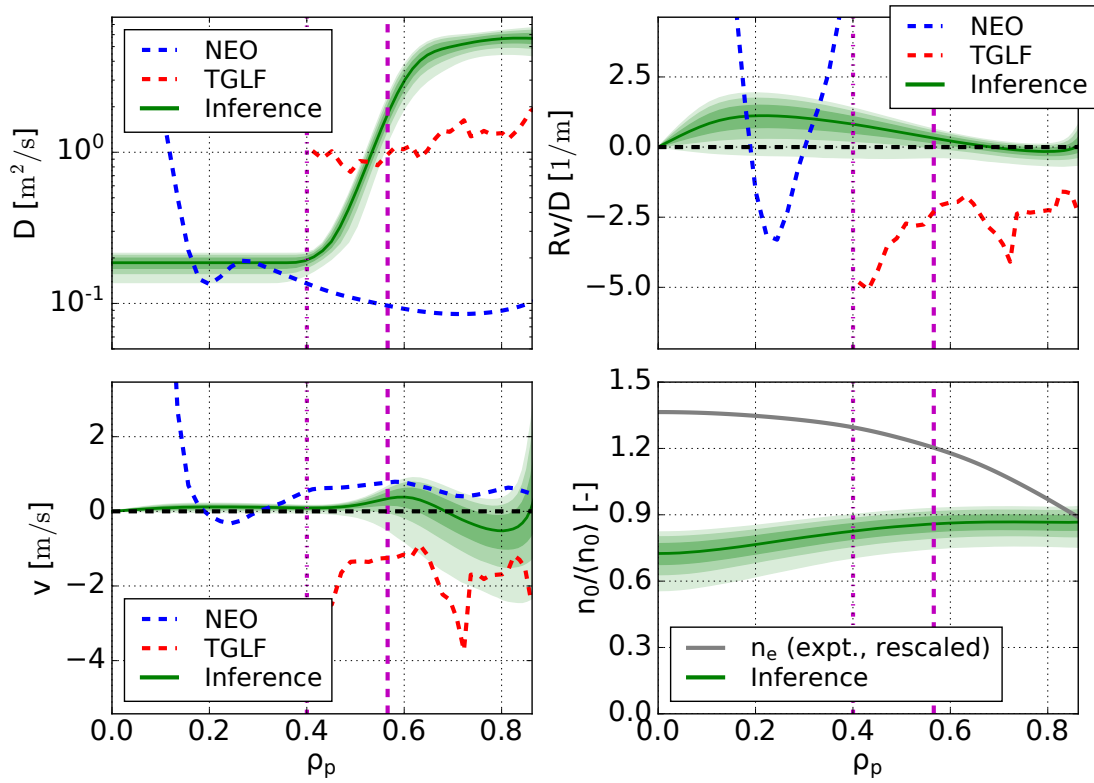


Figure 8-15: Inferred Al transport coefficients for DIII-D shot #180520. NEO and TGLF results are also shown. Profiles are conservatively shown only inside of $\rho_p \approx 0.8$ because outer regions are affected by unclear details of Al LBO sources.

We next display some of the experimental signals used in these inferences to describe the quality of agreement with Aurora forward modeling using inferred MAP parameters. Fig. 8-17a shows time histories for all the CER chords measuring F^{9+} in the low-power L-mode case (discharge #180526, $t \approx 2.75$ s). Here, black points with error bars represent experimental data and the red lines show the result of the synthetic diagnostic. Fig. 8-17b shows the






Figure 8-16: Inferred Al transport coefficients for DIII-D shot #180530. NEO and TGLF results are also shown. As in Fig. 8-15, we only show results inside of $\rho_p \approx 0.8$ since regions further outside are affected by unclear details of Al LBO sources.

match of radial C profiles in the same discharge, with different panels corresponding to different time bins during the LBO injection. C densities are seen to be similar in each panel, as a result of steady-state conditions with no impurity accumulation. Different colors in Fig. 8-17b identify different CER chords. Matches of analogous quality were obtained for Al^{13+} data.

Figs. 8-18 and 8-19 show the match between measurements and synthetic diagnostics for SPRED and SXR chords for the same experiment. Fig. 8-18 compares the experimental data for the six chosen lines (blue) with the time history from the synthetic diagnostic based on Aurora (yellow) and its binning in time over the ranges of signal integration of SPRED. For comparison, black curves in each panel show the source time history obtained from the unfiltered fast camera (*c.f.* Fig. 8-2). The same source time history is also shown in Fig. 8-19 on top of the line-integrated SXR signals. In this figure, the experimental data (with background already subtracted) are plotted in green and the Aurora synthetic diagnostic is



(a)



(b)

Figure 8-17: Match of experimental and modelled of CER signals in the DIII-D RMP ELM-suppressed discharge #180526 for (a) time-dependent F^{9+} injected via LBO, and (b) intrinsic, steady-state C^{6+} .

in blue. A good match is observed for all 6 chords used in this inference.

While SPRED data for the Al lines of Table 8.2 have been matched with comparable quality to the F lines of Fig. 8-18, signal-to-noise ratios for SXR data in shots #180520 and #180530 are too low to effectively constrain transport. CX lines of high quality have been measured for F in #180526 (see Fig. 8-4) and for Al in #180530, whereas the SNR is too low in #180520 to make the signal useful.




Figure 8-18: Match of experimental and modelled signals of SPRED F lines in the diverted negative triangularity discharge #180526.

8.4.3 Comparison of Experimental and Theoretical Modeling

Section 8.4.2 presented inferred radial profiles of transport coefficients. In addition to the experimental results, Figs. 8-14, 8-15, and 8-16 also overplot predictions from NEO and TGLF SAT-0. As in the investigation of the RMP ELM-suppressed discharge #175674 (Section 8.3), theoretical modeling is based on kinetic profiles that have been obtained via TGYRO. Results for each discharge are shown in Fig. 8-20, as a function of the square root of normalized toroidal flux, ρ_ϕ .

In the low-power L-mode case (#180526) we have found good TGYRO convergence up to $\rho_\phi = 0.9$, while in the low-power H-mode and high-power L-mode cases (#180520 and

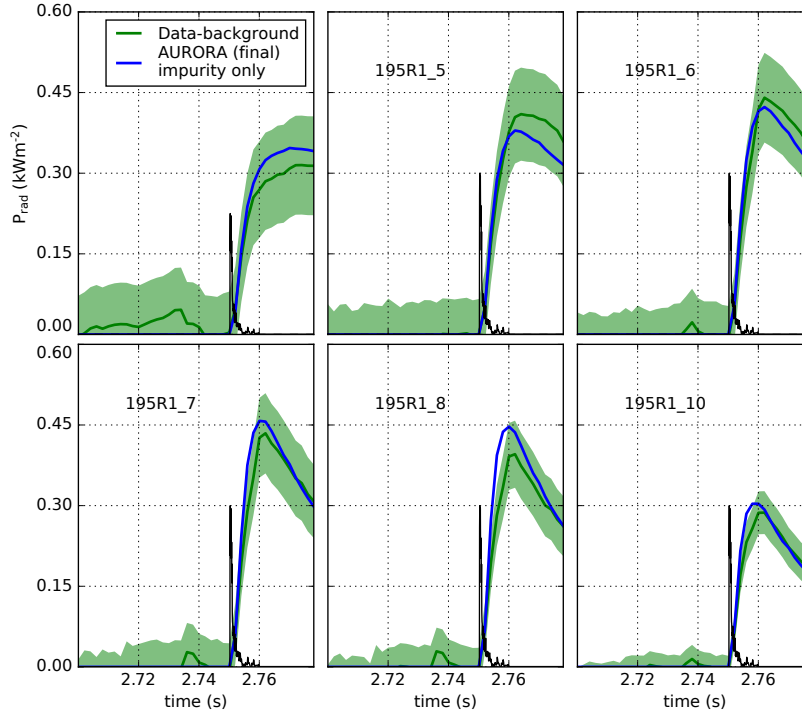
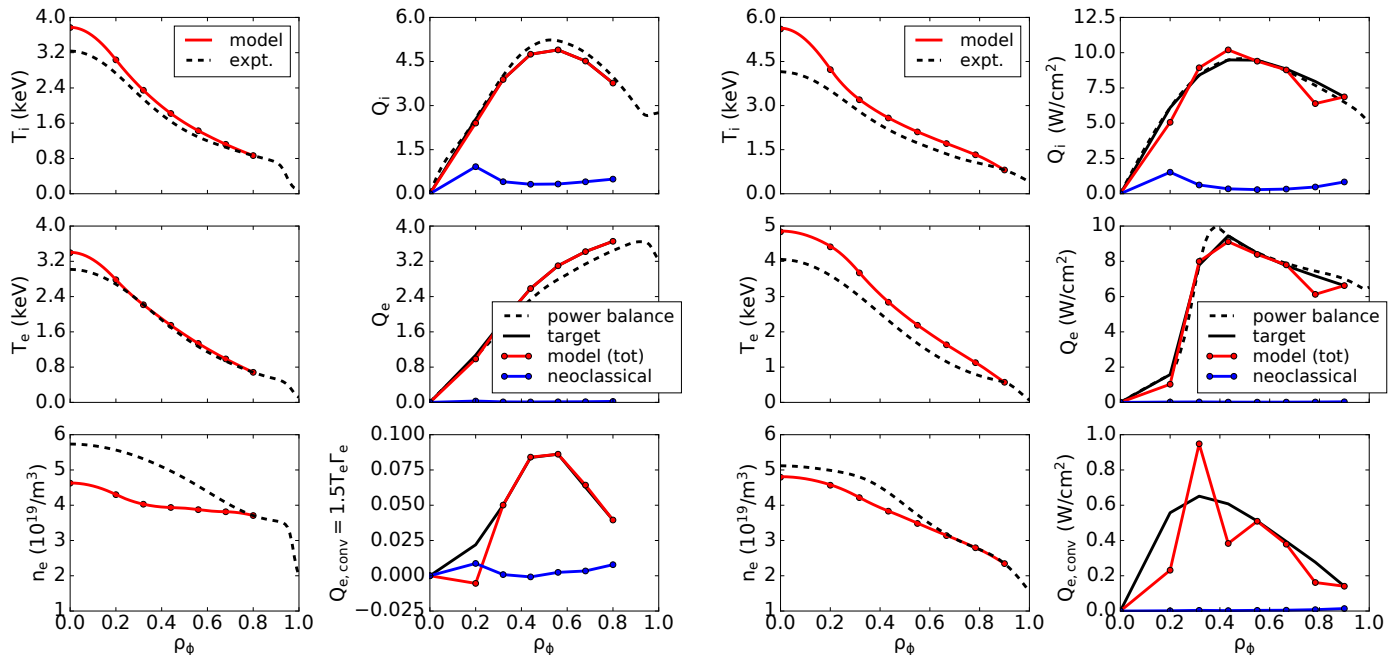


Figure 8-19: Match of experimental and modelled F signals for SXR chords in the diverted negative triangularity discharge #180526.


#180530, respectively) TGLF predicts excessively large gradients outside of $\rho_\phi = 0.8$ and we have therefore limited our modeling to within this radius. A feature worth remarking is that in the low-power H-mode (#180520) there is a significant discrepancy (outside of experimental uncertainties) in the electron density from power balance and from the turbulence model (TGLF). This is a clear indication of TGLF modeling of particle transport being less accurate than in the heat transport channel. Neoclassical contributions from NEO (blue), mostly relevant to the ion heat flux, Q_i , are relatively small compared to turbulent fluxes. The latter have been computed using TGLF SAT-0, rather than SAT-1, since the SAT-0 saturation rule gave better convergence for these cases within the TGYRO workflow, in contrast to the C-Mod analysis presented in Chapter 7. Identifying the reasons for these differences is beyond the scope of this work. However, we remark that the SAT-1 saturation rule is more modern and complete, attempting to include multiscale effects that have previously been observed to be important in recovering measurements across multiple transport channels [50] (see Section 7.3).

Fig. 8-21 shows TGLF spectra (relations between growth rates, γ , and normalized



(a) DIII-D shot #180520

(b) DIII-D shot #180526



(c) DIII-D shot #180530

Figure 8-20: Summaries of TGYRO runs for the three negative triangularity discharges examined in this chapter. In each sub-figure, kinetic profiles are shown on the left and heat fluxes on the right. The convective electron heat flux is defined as $Q_{e,conv} = 1.5 \Gamma_e T_e$ and is therefore related to particle transport.

poloidal wavenumbers, $k_{\theta}\rho_s$) for discharge #180520 at the time of the LBO injection ($t = 2.501$ s). Rather than displaying γ itself, we show its normalization by $k_{\theta}\rho_s$, since previous work has suggested that $\gamma/(k_{\theta}\rho_s)$ allows easier evaluation of the balance between low-k (ion-) and high-k (electron-) modes, as well as the importance of any multiscale turbulence effects [50, 46]. Each panel in the horizontal direction displays scans by $\pm 20\%$ for the turbulent drives (inputs to TGLF) that have been found to affect spectra most prominently: the normalized gradient scale lengths for electron and ion temperature (a/L_{T_e} and a/L_{T_i} , respectively), the normalized electron density gradient scale length (a/L_n), and the ratio of the background ion and electron temperatures (T_i/T_e). Results in Fig. 8-21a (Fig. 8-21b) refer to $\rho_p = 0.6$ ($\rho_p = 0.85$). Blue points refer to parameters at their nominal values, red (green) points refer to an increase (decrease) by 20% of the scanned variable. Points marked by squares (circles) correspond to modes with negative (positive) real frequency, i.e. propagating in the ion (electron) diamagnetic direction. Horizontal red-dashed lines near the bottom of each panel show the estimated $E \times B$ shearing rate, below which turbulence growth rates are expected to be quenched.

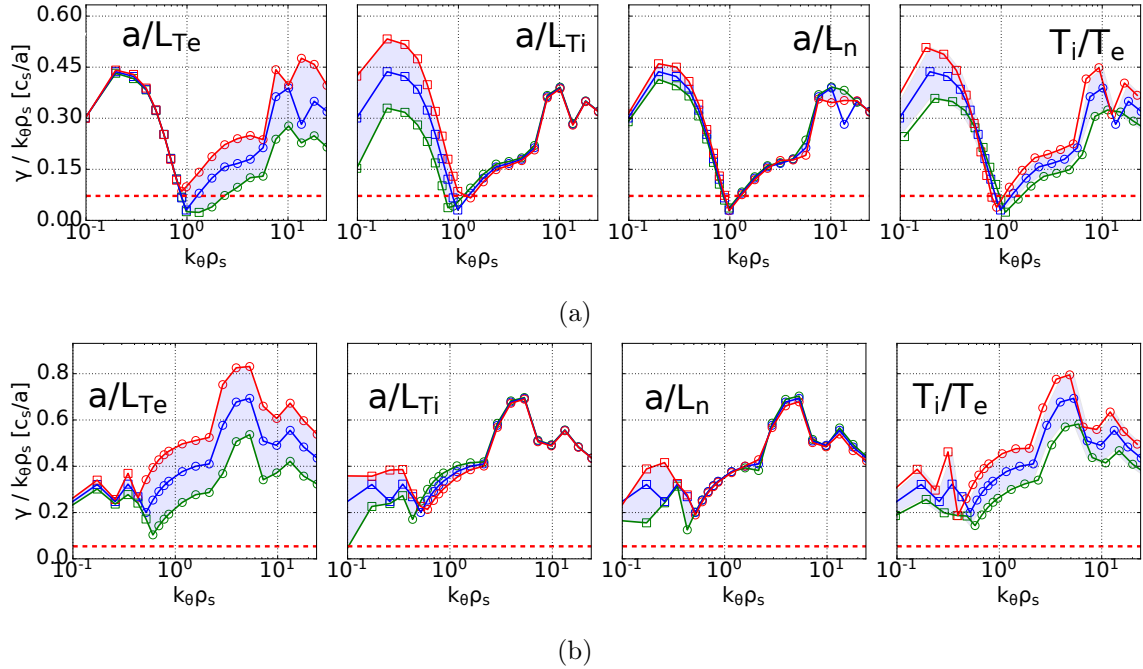


Figure 8-21: TGLF spectra at (a) $\rho_p = 0.6$ and (b) $\rho_p = 0.85$ from scans of major turbulence drives in the diverted negative triangularity discharge #180520. Green points indicate reductions of each input by 20%, red ones indicate increases by 20%; blue points show the baseline (all parameters at nominal values). Horizontal red dashed lines indicate the estimated $E \times B$ shearing rate.

Spectra in Fig. 8-21a for the low-power H-mode case (#180520, $t = 2.501$ s) identify important contributions from both Ion (ITG) and Electron Temperature Gradient (ETG) modes at $\rho_p = 0.6$. TGLF shows only weak dependences on plasma collisionality in all the negative triangularity cases of interest here, suggesting that Trapped Electron Modes (TEM) only play a marginal role. The similar amplitude of $\gamma/k_\theta\rho_s$ at ion ($k_\theta\rho_s < 1$) and electron scales ($k_\theta\rho_s \gtrsim 1$) suggests that multiscale interactions may be important in this H-mode condition. Looking further out in radius, at $\rho_p = 0.85$ (Fig. 8-21b), we find complete dominance of electron-scale modes, with variations in a/L_{Ti} playing only a minor role, and a/L_{Te} and T_i/T_e being most important.




Figure 8-22: TGLF spectra at (a) $\rho_p = 0.6$ and (b) $\rho_p = 0.85$ from scans of major turbulence drives (indicated at the top of each panel) in the diverted negative triangularity discharge #180530. Blue points show results obtained at nominal values of all inputs; green (red) points indicate reductions (increases) of scanned variables by 20%. Horizontal red dashed lines indicate the estimated $E \times B$ shearing rate.

On the other hand, the nature of turbulent transport in the L-mode conditions examined here appears to be very different. At the times of LBO injections in #180526 ($t = 2.75$ s) and #180530 ($t = 3.8$ s), ITG is strongly dominant. Fig. 8-21 shows analogous TGLF scans to those in Fig. 8-21, with the same radial locations ($\rho_p = 0.6$ and 0.85) and percentage changes ($\pm 20\%$) being examined. At both radial locations, ITG is likely to account for all the observed transport. There appears to be no evidence for significant multi-scale interactions.

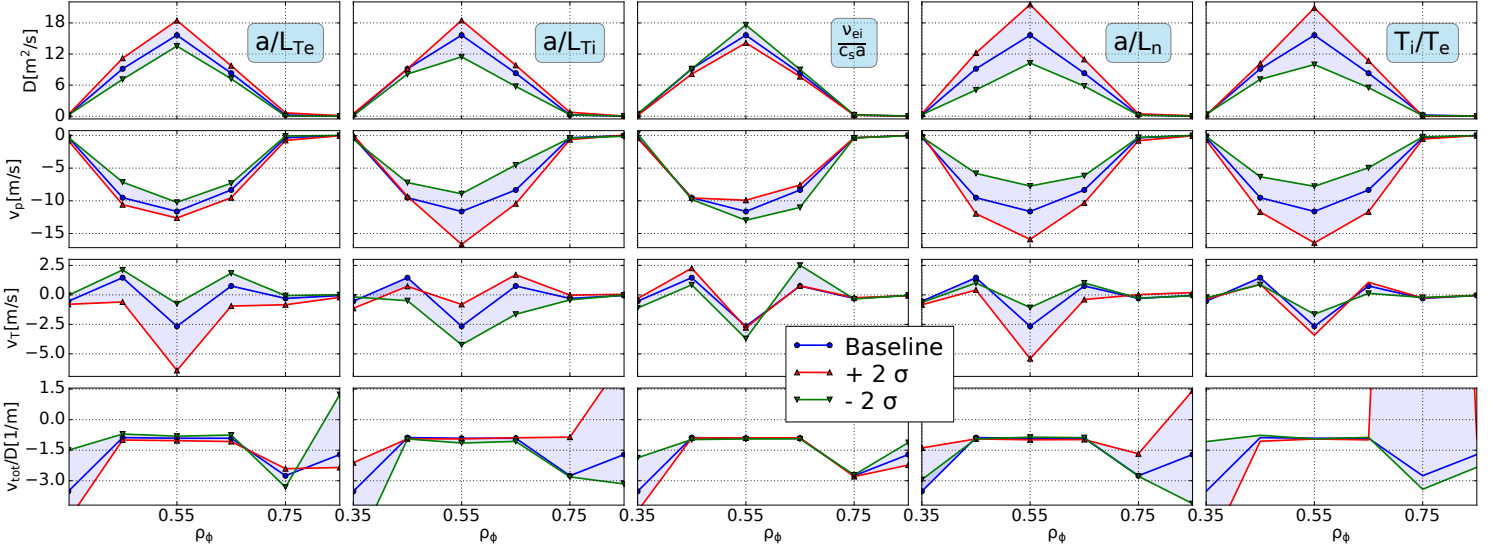


Figure 8-23: Diffusion and convection coefficients from TGLF 2σ scans around experimental parameters for the DIII-D negative triangularity discharge #180526 at $t = 2.75$ s (low-power L-mode phase).

In Fig. 8-23 we show TGLF scans of a/L_{T_e} , a/L_{T_i} , normalized collisionality $\nu_{ei}/c_s a$, a/L_n , and T_i/T_e for the low-power L-mode case (#180526, $t = 2.75$ s), displaying the sensitivity of D and v predictions to the main turbulent drives. As in the C-Mod scans of Fig. 7-22, we find that these 1-dimensional scans suggest an uncertainty in both D and v at midradius of at least 50%, showing that an exact match to experimental inferences should not be expected. Analogous conclusions hold for all the discharges examined in this thesis. In the specific case of Fig. 8-23, limited sensitivity to collisionality is observed for both D and v , compatibly with the negligible effect observed on linear growth rates (not shown). Unlike in the C-Mod I-mode case of Fig. 7-22, some sensitivity of particle transport coefficients on a/L_{T_e} is also observed in this case, as shown in Fig. 8-23.

Figs. 8-14, 8-15, and 8-16 show some qualitative agreement in both D and v , but non-negligible disagreement on important features. While we find a relatively good match of D near $\rho_p = 0.25$ with NEO, D appears to be generally underestimated by TGLF at midradius. TGLF convection estimates at midradius appear to indicate significantly more impurity density profile peaking than we actually observe in experiment. We remark that the value of Rv/D obtained by our inferences is strongly constrained by the CER data for C^{6+} .

In the low-power L-mode case (#180526, $t = 2.75$ s, Fig. 8-14), NEO is in good agreement with inferred coefficients near the magnetic axis. The total impurity density profile is peaked between $\rho_p \approx 0.75$ and $\rho_p \approx 0.95$ in both theoretical modeling (TGLF) and experiment, although the radial variation of Rv/D is significantly different. In the low-power H-mode case (#180520, $t = 2.501$ s, Fig. 8-15), the experimentally-inferred Rv/D is found to be close to 0 across all the examined radial range. D and v have a qualitative match to NEO results near $\rho_p = 0.2$, although inside of this radius NEO appears to be sensitive to details of the rotation profile, resulting in a sudden rise of both D and v inside of $\rho_p \approx 0.15$. The value of D given by TGLF at midradius is a factor of 2-3 smaller than the experimentally-inferred result. We remark a significant disagreement of Rv/D at midradius, where TGLF finds $Rv/D < -2$, whereas the inference (and the CER C^{6+} data) clearly indicates flat impurity profiles. The inference of the high-power L-mode case (#180530, $t = 3.8$ s, Fig. 8-16) yields a close match in D with the TGLF results, but differs significantly from the NEO prediction near the magnetic axis.

A comparison of neoclassical predictions from NEO and experimental inferences for $\rho_p < 0.2$ may be considered to be more meaningful in these DIII-D experiments than it was in the C-Mod analysis of Chapter 7 since there are very few sawteeth in these $\delta < 0$ discharges (none in #180530, and only one in each of #180520 and #180526, for the time ranges of interest). While we do not give much credit to the sudden rise of neoclassical D inside of $\rho_p \approx 0.15$ in the low-power H-mode and high-power L-mode cases (#180520 and #180530, respectively), the NEO prediction of $D \approx 0.9$ at $0.2 \lesssim \rho_p \lesssim 0.3$ for the high-power L-mode is robust and the reason for such a large discrepancy between NEO and the experimental inference in this case is currently not understood.

Comparing experimental and theoretical estimates of transport coefficients is even more challenging in these $\delta < 0$ discharges than in the RMP ELM-suppressed H-mode discharge (#175674) because no CER chords were measuring Al or F in the pedestal in the negative triangularity cases. The constraint provided by C measurements in the pedestal is very useful, but insufficient to separate the effects of D and v in this region. The apparent dependence of pedestal results in shots #180520 and #180530 on details of the Al source time histories prompted us to avoid a close comparison to theoretical transport models in the pedestal for these two discharges. While the Al source is indeed measured in both these experiments, the characterization of LBO injections by spatially-integrated signals over a

rectangular region of the camera is obviously quite approximate. Within our sensitivity tests, small variations of the source function did not appear to affect D and v profiles inside of $\rho_p = 0.8$, thus allowing more meaningful model comparisons in the core.

We highlight that the flat or hollow nature of core impurity density profiles is strongly constrained by C^{6+} measurements. While these are not free of modeling assumptions (e.g. taking $n = 2$ halo contributions to be negligible) and uncertainties (e.g. in the atomic rates and beam attenuation coefficients), the constraint that CER data provide on the C profile shape is likely one of the most robust in our entire analysis. This suggests that the results from NEO and TGLF, rather than the experimental inferences, must be examined most closely to understand the origin of the mismatch. The sensitivity of the TGLF D and v , illustrated in Fig. 8-23, suggests that variations of D and v predictions by a factor of 2 at midradius may not be unreasonable, considering that multiple parameters could be different from our baseline inputs (an effect not captured by our 1-dimensional TGLF scans). However, in all the cases examined here, TGLF consistently predicts similar $v/D < 0$ in the region where the code is most effective (in regions where turbulence is most unstable).

8.5 Discussion and Summary

In this chapter, we have presented radial profiles of particle transport coefficients from experimental inferences based on Al and F LBO injections in DIII-D. We have focused on two operational scenarios: RMP ELM-suppressed H-modes and negative triangularity discharges. In the first case, a single discharge has been analyzed, comparing to previously published results. In the second case, three discharges have been analyzed. All plasmas described here were diverted. This is an important difference from previous inner-wall-limited (IWL) negative triangularity experiments at DIII-D [285], where a high Z_{eff} was observed. The three diverted $\delta < 0$ shots discussed here had much lower impurity confinement, as evidenced by the Z_{eff} profiles in Fig. 8-11, the relatively low density of C, and the rapid decay of CER signals following LBO injections (Figs. 8-12a and 8-12b). On the other hand, the RMP ELM-suppressed case presented evident impurity accumulation, which has been found to derive from $v/D < 0$ well inside of the sawtooth inversion radius and a strong inward pinch in the pedestal region.

The modeling of steady-state C profiles as part of our inferences has offered an excellent

constraint on v/D profiles at midradius. Our novel multi-impurity framework may also be used to compare inferred D and v for C impurities with theoretical models, although this has not been shown here. Such comparisons would allow, for example, for the validation of Z scalings for neoclassical transport near the magnetic axis (where impurity accumulation can be particularly damaging) and in the pedestal.

Inferences of particle transport coefficients have been found to be in relatively good agreement in all cases discussed in this chapter, although some apparent discrepancies have been observed. As in the C-Mod results of Chapter 7, we have found a tendency of TGLF to over-predict the peaking of impurity densities. Diffusion also appears to be underestimated by TGLF in the negative triangularity discharges. All inferences have shown significant deviation of impurity profile shapes with respect to n_e density profiles, suggesting that models that set impurity density profiles as a constant fraction of the n_e profile can run into important inaccuracies. In all cases, we have found diffusion to be strongly reduced from turbulent to neoclassical levels between the sawtooth inversion and mixing radii. The fact that the inversion radius is typically found to be near the $q = 1$ surface suggests that integrated modeling of impurity transport should carefully consider the importance of this rational surface in suppressing transport. A simple method to combine neoclassical and turbulence predictions from codes like NEO and TGLF could then be to smoothly interpolate between their transport coefficients near the $q = 1$ surface (or, if known, the sawtooth inversion radius) rather than at a fixed radial location.

Diverted negative triangularity is a relatively exotic scenario for current tokamaks. Recent experiments described here have paved the way for new DIII-D campaigns to improve the plasma shape further and move the divertor leg on the high-field side to the outer wall, rather than the lower shelf. This is expected to improve performance further and determine whether negative triangularity should be pursued more aggressively in new dedicated devices. Given the high performance and low impurity retention achieved at $\delta < 0$ in L-mode, it appears possible that a high-field device running in L-mode at $\delta < 0$ could offer highly-radiative scenarios, with considerable advantages for divertor operation. The results discussed in this chapter have demonstrated favorable (rapid) impurity transport at $\delta < 0$ in DIII-D and have generally confirmed the applicability of NEO and TGLF in the examined cases, although their predictions have been shown to be relatively inaccurate in some respects. On the path to accurate transport model validation, the analysis of this

chapter has highlighted the difficulty of determining where neoclassical or turbulent models are dominant. Of course, this would not be a concern in future research if both transport scales could be modeled self-consistently within the same code.

Theoretical turbulence modeling of these discharges has so far been relatively limited. TGLF scans have demonstrated dominance of ITG in L-modes and ETG in the H-mode case. Electron-scale transport always becomes dominant near the pedestal-top. Linear and nonlinear CGYRO runs are necessary to determine in greater detail the exact balance of turbulent modes and how they saturate. This analysis is in progress and will be described in an upcoming publication. While the physical picture of Fig. 8-9 in terms of modified trapped electron orbits is appealing, the overall performance of negative triangularity discharges has also been found to be affected by pedestal ballooning stability preventing a transition to H-mode [288], as well as lower core [282] and SOL [289] fluctuations. The impurity transport analysis described in this chapter has clearly shown relatively large D in the outer part of the plasma and weak inward convection across the entire radial profile. In the absence of a well-developed pedestal, these observations in the core clearly explain the main reasons for low impurity retention in this scenario.

Future work on impurity transport at DIII-D should attempt to constrain the Z dependence of pedestal particle transport by including free parameters in the Aurora transport coefficients, as described in the context of C-Mod research in Section 7.3. A separation of Al^{10+} and Al^{12+} emission with the CER diagnostic could offer a particularly compelling method to identify Z scalings. Such effort would likely depend strongly on accurate predictions of edge neutral densities. In Ref. [228], Dux *et al.* showed how CER measurements of multiple Ne charge states on AUG could only be explained when considering the significant role of CX in setting the impurity charge state balance. An extension of this analysis, focusing not only on steady-state ionization balance but also on transport through the pedestal, would greatly benefit from fast CER measurements such as those that have recently become available at AUG [290].

Summary and Conclusions

This thesis has presented contributions across a spectrum of subjects related to particle transport in tokamaks. Our focus has been on increasing the fidelity of impurity transport modeling, expanding spectroscopic constraints, leveraging modern computing resources, and quantifying the impact of charge exchange recombination on impurity charge state balance.

The main outcomes of this work can be summarized as follows:

- **Development of the open-source Aurora package** has advanced capabilities for particle transport simulations in confined fusion plasmas, creating a platform for core-edge integration that interfaces the extensive OMFIT framework and state-of-the-art external simulation suites.
- **Analysis of D Ly $_{\alpha}$ data on Alcator C-Mod** has permitted validation of SOLPS-ITER predictions for atomic neutral densities at the midplane, finding good agreement across the radial profile.
- Based on a **wide database of C-Mod D Ly $_{\alpha}$ signals**, we have established a number of scaling laws for edge neutral densities and separatrix electron densities. We have demonstrated that **cross-field deuterium particle fluxes offer a particularly robust test of pedestal transport**, thus motivating further detailed analysis of our dataset.
- Making use of Monte Carlo FIDASIM simulations, we have shown that **the impact of NBI neutrals on impurity transport is relatively small** for DIII-D plasmas, particularly compared to edge neutrals, consistent with recent work at AUG [228].

- **For the first time, charge exchange has been included in inferences of impurity transport coefficients**, significantly improving the self-consistent assessment of ionization balance near the plasma edge.
- **A novel forward model for the entire Ca K_α spectrum has been developed** for inferences of impurity transport on C-Mod, employing a new compilation of atomic rates that has been shown to match XICS high-resolution measurements in remarkable detail.
- Aurora has permitted efficient **deployment of fully-Bayesian inferences of C-Mod and DIII-D impurity transport to High-Performance Computing (HPC) clusters**. This capability enables parallel execution of nested sampling schemes, providing results in a reasonable time (between 10 minutes and 12 hours, depending on the problem complexity, typically using ≈ 300 CPUs).
- Fully-Bayesian *experimental impurity transport inferences* based on Ca LBO injections on C-Mod have been demonstrated for the first time. **Inferences for L-, EDA H- and I-mode discharges** have been compared to neoclassical and turbulent (quasilinear gyrofluid and nonlinear gyrokinetic) modeling, generally finding good agreement in diffusion across the radius, albeit significant discrepancies have been observed in convection.
- In collaboration with T. Odstrčil, we have developed the **ImpRad analysis framework within OMFIT**, leveraging Aurora capabilities for a wide range of applications on particle transport and radiation in tokamaks. ImpRad is the cornerstone of all DIII-D research discussed in Chapter 8 and is currently being employed by multiple collaborators to push the boundaries of core-edge integration research.
- **Fully-Bayesian inferences of experimental F and Al transport have also been demonstrated on DIII-D**. We have presented results for a **RMP ELM-suppressed discharge** discussed in previous work [270] and for recent **diverted negative triangularity experiments**. In all cases, we find transport coefficients to be in reasonable agreement with neoclassical and gyrofluid modeling, although **flat and hollow impurity density profiles do not seem to be appropriately predicted by NEO and TGLF and turbulent diffusion appears to be often**

under-estimated by TGLF. Results highlight an effective suppression of transport inside the $q = 1$ surface, an effect not typically considered explicitly in integrated modeling of tokamak plasmas.

9.1 Considerations For Effective Particle Transport Inferences

Based on the experience gained by analyzing impurity transport on both C-Mod and DIII-D, a few recommendations are set forth for future experiments:

- One cannot overstate the importance of having a clear characterization of the LBO particle source when attempting to infer radial profiles of impurity transport coefficients. In this thesis, we made use of an edge filterscope for C-Mod and a fast unfiltered camera for DIII-D, which made accurate quantification of core transport coefficients possible, albeit still challenging. In the absence of these diagnostics, the task of quantifying radial profiles of D and v would have been intractable. To reduce the presence of LBO clusters, the LBO concept [20] could be complemented by a new shutter system, placed at the entrance of the vacuum vessel down the LBO duct. By closing the shutter approximately 1 ms after the laser ablation, slower particle clusters could be blocked, resulting in cleaner and more reliable injections with limited hardware modifications.
- When attempting to compare experimental D and v to neoclassical and turbulence models, one should ensure that there is *no significant MHD activity in the plasma*. MHD modes (e.g. sawteeth) would inevitably “pollute” the analysis, since they are not modelled in detail by codes such as Aurora. All the C-Mod experiments of Chapter 7 and half of the DIII-D experiments discussed in Chapter 8 (180520 and 180526) had sawteeth, albeit either quite regular (a condition which helps in modeling) or in small numbers.
- While LBO injections lend themselves to “piggyback” (operation during experiments with a different focus), detailed inferences of radial profiles of impurity transport coefficients can only be successful with excellent diagnostic coverage, which typically requires dedicated experiments. This is unlike in studies where simple impurity transport metrics, e.g. impurity confinement times, are of interest. Therefore, it

is important to set appropriate expectations: piggyback LBO injections are useful, but typically insufficient for detailed transport model validation.

- It is often unrealistic to rigorously examine radial profiles of particle transport coefficients in the core without experimental constraints on pedestal transport. Time histories of spectroscopic signals after LBO injections depend on the cross-field motion of particles going into the core *and* later going out. This means that the pedestal always affects observed time histories in the core, except at the beginning of the signal rise phase. One may argue that the SOL can be more effectively separated due to much faster parallel dynamics. In this thesis, we have attempted to constrain pedestal transport using new spectroscopic analysis methods, improved sampling techniques, and expanded forward models (see Chapter 5).
- The use of relative amplitudes of spectral lines from different charge states offers much stronger constraints on particle transport than individually-normalized emission for each spectral line. Within the context of this thesis, this has become particularly obvious in sensitivity tests with Ca VUV brightnesses on C-Mod. In the absence of accurate photon emissivity coefficients, VUV lines can only provide a constrain on the time history of individual charge states, with only weak constraints on the *relation* between charge states. Future work on pedestal and SOL transport should extend efforts to place constraints on transport based on detailed atomic physics, as in the case of Ca K_α spectra measured via XICS on C-Mod.

9.2 Future Work

This thesis has built on decades of research on impurity transport, developing more reliable analysis tools that are now being adopted across the research community. Aurora is by no means the most advanced suite to model fusion plasmas, but it does offer a unique mix for modern core-edge integration efforts. Aurora’s design is well suited for modern computing resources, compared to pre-existing tools, and it brings to light new opportunities in computational statistics and machine learning. Combined with the ImpRad module presented in Section 8.1, these tools open the door to countless research opportunities in tokamak experimental analysis.

Future work on particle transport should focus on the critical role of pedestals in regulating global properties. Doing so will require facing significant challenges in both measurements and modeling, due to the short spatial scales, fast temporal scales, and the complex influence of SOL (2D) transport. The development of a D Ly $_{\alpha}$ database using existing C-Mod data has contributed to identifying key dependencies of neutral densities in this region. Comparison of this data with analogous Ly $_{\alpha}$ measurements on DIII-D [240] is ongoing at the time of writing.

Overall, both in the context of core and pedestal transport model validation, the key bottleneck is currently in the availability of high-quality experimental data. Both modeling capabilities and computational resources now appear to be adequate for the task, but measurements are often sparse and have insufficient resolution.

To provide a concrete benchmark for the validation of emergent theoretical predictions of pedestal transport, significant effort should be placed on diagnostic development. Indeed, it has been evident over much of this thesis' work that elaborate computational statistics can only go so far to identify transport fingerprints if experimental data is not sufficiently constraining. Integrated analysis of multiple diagnostics, as demonstrated here for both C-Mod and DIII-D, can offer valuable insights, but reliable assessments require single diagnostic systems with high spatial and temporal resolution and relative calibration across many chords. In this context, we foresee a few particularly promising paths forward:

- Fast CER diagnostics such as the one developed on AUG [290] could offer the means for detailed examination of pedestal particle transport. Extended poloidal coverage would also be desirable to assess impurity poloidal asymmetries, which are often important for high-Z ions. X-ray crystal spectroscopy systems like C-Mod's could also be effective if a larger number of chords were set to span the entire plasma (including the pedestal) with higher spatial resolution.
- The analysis of edge D Ly $_{\alpha}$ spectroscopy on C-Mod, discussed in Sections 6.2 and 6.3, has recently identified the quantification of deuterium cross-field fluxes as a robust metric to compare theoretical transport models to experimental measurements. The ability to compare *fluxes* rather than *transport coefficients* in this case is compelling, particularly since these data are available for a large number of discharges. Research on this subject is ongoing.

Currently, fusion reactors are envisioned to operate with puffing of low- and medium-Z impurities in the divertor in order to reduce the extreme heat fluxes near the strike points. Understanding the physics of ion transport, atomic and molecular processes in divertor regions is therefore paramount. At the time of writing, operation and experimentation on the subject are still mostly an empirical effort. With a growing number of sophisticated models becoming available, it has become essential to develop frameworks to *bring together* and *interpret* their results. The work presented in this thesis offers statistical, numerical, and analysis methods that facilitate such efforts. Their applications on Alcator C-Mod and DIII-D have suggested that individual transport models may be sufficiently mature in their individual domains. The challenge of core-edge integration is where fusion research must therefore turn its attention, both from empirical and modeling standpoints.

Appendix A

The Aurora package

“Everything should be made as simple as possible, but no simpler.”

– attributed to A. Einstein

The development philosophy of Aurora, introduced in Section 5.3.1, is closely aligned with the quote above: it is a modern package that aims at providing simple, yet accurate, analysis tools for impurity transport, neutrals and radiation modeling. Aurora was initially developed as a alternative to pre-existing impurity transport codes that did not perform well in high-performance computing environments. Over time, many of Aurora’s capabilities have been generalized and extended to be more broadly applicable. These include the simulation (forward modeling) of particle transport using a diffusive-convective ansatz, the modeling of neutral particles by interfacing other codes of varying complexity, and the processing of atomic rates to perform integrated modeling, create synthetic diagnostics, and estimate radiated power. In this appendix, we describe aspects of the package at the time of writing (version *v2.0.0*), attempting to provide a clear reference that may remain relevant even as Aurora keeps expanding in time.

Aurora’s development strategy is inclusive of efforts made in other collaboration contexts. For example, Aurora is fully integrated with the OMFIT framework [198], leveraging some of the Python classes that are part of OMFIT and have been spun out for independent installation. Aurora builds on the success of the STRAHL code [195] and aims to work together with other modeling suites, such as SOLPS-ITER [203] and the ADAS framework [114], in the context of core-edge tokamak analysis.

At the time of writing, approximately 80% of Aurora is composed by Python3+ code,

12.5% by modern Fortran, and 7% by Julia code, with small fractions of other languages used to support installation and documentation procedures. Aurora is open-source and publicly maintained on GitHub at <https://github.com/fsciortino/Aurora>, with documentation hosted at <https://aurora-fusion.readthedocs.io>.

The objective of this appendix is to provide a high-level view of some features, rather than a substitute to the already extensive code documentation available at the above hyperlink. The Aurora 1.5D forward model of particle transport was already presented in Section 5.3.1. In Section A.1, we briefly describe Aurora’s capabilities related to processing of atomic data and spectral predictions. Section A.2 describes the simple edge recycling model used in Aurora’s simulations. Section A.3 describes the superstaging approximation, available in Aurora to reduce the computational cost of simulations with large numbers of ion stages. Finally, Section A.4 presents a simple demonstration of Aurora’s capabilities to interface with SOLPS-ITER, particularly in the context of atomic D neutral density predictions from EIRENE. Applying these tools to a SOLPS-ITER simulation for the ITER baseline scenario, we show that charge exchange is unlikely to affect the total radiated power inside the Last Closed Flux Surface (LCFS) in ITER. Part of the material in this appendix has been adapted from Ref. [186].

A.1 Atomic Data and Spectral Predictions

Aurora relies on ADAS data for much of its functionality. While it may be convenient to combine usage of Python routines within an ADAS framework with Aurora, it is not necessary to run (or even have access to) an ADAS framework. Aurora can work with atomic data from the OPEN-ADAS website (www.open-adas.ac.uk), fetching and interpreting files automatically through the internet. Most of the Generalized Collisional-Radiative (GCR) formats distributed by ADAS can be processed. While users may indicate specific ADAS files that they wish to use, a set of defaults are conveniently listed within Aurora for most ions of interest in fusion. This offers the side benefit of lowering the “entry barrier” for researchers who are not familiar with ADAS nomenclature and who may be unsure about which files offer the highest data quality. Cooling coefficients, both resolved for each ion stage and weighted by fractional abundances of each charge state at ionization equilibrium, can easily be loaded and processed for use in integrated modeling. For applications in

spectroscopy, the `aurora.read_adf15` function allows parsing and interpolation of Photon Emissivity Coefficients (PECs) in the ADAS ADF15 format, separating components driven by ionization, excitation or recombination processes, as well as metastable-resolved parts whenever these are available in the chosen files. This function has also been used to process files in the ADF15 format from the atomDB database [131], widely used in x-ray astrophysics. Aurora also offers initial support to work with the ColRadPy [120] package, which makes use of ADAS resolved ion data collections in the ADF04 format for GCR model predictions. Specifically for the analysis of neutral H-isotope particles, some atomic data from the formulae by Janev & Smith [227] or Janev, Reiter & Samm [291] are also available.

All the tools discussed in this section are applicable to the Aurora forward model (Section 5.3.1) as well as within other codes and analysis frameworks. Indeed, Aurora has been used for experimental spectroscopic signal analysis [155, 156], as well as to post-process data from neutral codes such as KN1D [200], FIDASIM [201, 202] and EIRENE [203]. Aurora does not include capabilities to work with molecules at the time of writing; the subject will be explored in future work.

A.2 Aurora Recycling Model

As part of its 1.5D particle transport forward model, Aurora has a simple 0D model of the SOL and divertor. By 0D, we mean that it does not have realistic descriptions of the physical geometry, but rather it describes the edge as a set of *particle reservoirs*. A similar approach has been previously adopted in other codes. The Aurora edge model is particularly inspired by the STRAHL one [195], which it attempts to generalize and slightly extend.

When particles in an Aurora simulation reach the edge of the radial grid, they are subject to the boundary condition set by the diffusion coefficient at the last grid point and an input decay length set by the user. The flux $\Gamma_z = D_{wall} \frac{n_{wall}}{\lambda_{wall}}$ is then taken to reach the wall reservoir at every time step.

It has been observed that, in some DIII-D ELMy plasmas, partially-recycling impurities such as Al can be retained at the wall for some time and re-enter the main chamber several ms after the ELM event. To model this effect, Aurora’s recycling model offers a slight extension of the model used in STRAHL, allowing for *temporary* retention of impurities at the wall, before these are recycled. We model such wall retention by a simple exponential

decay, using

$$N_{w,ret} = N_{w,ret}^{t=0} e^{(-t/\tau_{w,ret})} \quad (\text{A.1})$$

The following algorithmic steps are implemented:

1. First of all, a fraction of the particles that were previously retained at the wall now leaves:

$$N_{w,ret} \leftarrow N_{w,ret} + \frac{dN_{w,ret}}{dt} dt = N_{w,ret} \left(1 - \frac{dt}{\tau_{w,ret}} \right). \quad (\text{A.2})$$

Here, left-pointing arrows are used to indicate the change made between time steps, e.g. in the above expression $N_{w,ret}$ is substituted by the expression on the right hand side during each time step.

2. New particles reach the wall:

$$N_{w,ret} \leftarrow N_{w,ret} + r_{w,rcl} dt \quad (\text{A.3})$$

with the wall recycling rate defined as

$$r_{w,rcl} = R_{w-,rcl} (r_{\perp} + r_{\parallel,lim}) \quad (\text{A.4})$$

Here, $R_{w-,rcl}$ is a scalar giving the fraction of particles that hit the wall that is affected by recycling (over some time scale that depends on retention efficiency), r_{\perp} is the rate of cross-field loss between the last 2 radial grid points (i.e. at the wall) and $r_{\parallel,lim}$ is the rate of particle loss to the limiter. Note that a separately-calculated rate of parallel loss to the divertor is computed, but only affects the replenishing of the divertor reservoir, whose relation to the main plasma is not properly referred to as ‘‘recycling’’; hence, this rate does not enter the recycling model directly.

The combination of steps (i) and (ii) gives

$$N_{w,ret} \leftarrow N_{w,ret} \left(1 - \frac{dt}{\tau_{w,ret}} \right) + r_{w-rcl} dt \quad (\text{A.5})$$

3. We then re-define the rate of particles going back to the main plasma as a particle source:

$$r_{w,out} = N_{w,ret}/\tau_{w-ret}. \quad (\text{A.6})$$

Note that the difference between $r_{w,out}$ and $r_{w,rcl}$ is that the latter is the rate at which particles that will be recycled reach the wall; $r_{w,out}$ on the other hand is the rate of particles leaving the wall in the current time step, rather than at later times. The difference between these two rates is introduced here to model particle retention effects observed in DIII-D ELMy plasmas.

4. Finally, the (neutral) particle flux back to the plasma is

$$S = S_0 + r_{w,out} + N_{div}/\tau_{div} \quad (\text{A.7})$$

where S_0 is the neutral source given by the user as a function of time (i.e. the externally injected source), N_{div} is the number of particles in the divertor reservoir and τ_{div} is the divertor clean-out timescale. If part of the externally injected source is set to go to the divertor, e.g. via a divertor puff, then we use

$$S = S_0 (1 - f_{div}) + r_{w,out} + N_{div}/\tau_{div} \quad (\text{A.8})$$

where f_{div} is the fraction of incoming neutrals reaching the divertor rather than the main chamber. This same fraction must also be added to N_{div} when computing the number of particles in the divertor reservoir at each time step.

Note that in order for the recycling model to be numerically stable the adopted time step sizes must always be smaller than any of the time scales imposed by the user (for the divertor clean-out, the wall retention, etc.).

It is not obvious how recycling particles should be radially distributed when re-entering the main plasma reservoir. The trivial solution of setting their birth location at the wall effectively makes them subject to the diffusive-like boundary condition at the last grid point, which is purely numeric and not physical. An alternative is to add recycling particles to the main simulation reservoir with the same radial distribution used for the external impurity injection, which can be computed using ionization rates as a function of radius. This is typically a good option if neutrals are also evolved as part of the Aurora forward model, although the definition of D and v coefficients for the neutral stage does not lend itself to simple interpretation in our context. Yet another option is to simply place recycling impurities at the radial grid point that is closest to the separatrix. This is also physically

unreasonable, but may be numerically more useful since the simple 0D SOL transport model described above is anyway not a sufficiently detailed description of 2D dynamics beyond the last closed flux surface. All of these choices are allowed in Aurora.

A.3 Superstaging

As described in Section A, particle transport can be modelled via a set of coupled continuity equations for all the charge states of a given ion, which for a charge state i can be written as

$$\frac{\partial n^i}{\partial t} = -\nabla \cdot \Gamma^i + Q^i. \quad (\text{A.9})$$

In codes that attempt to solve this system of equations in 2D or 3D geometry, particularly for studies on the edge of fusion plasmas, the computational complexity of impurity ions can be extremely limiting, since each charge state may be seen to represent a different atomic species. For simpler 1.5D modeling in the core, as in Aurora’s forward model, high-Z ions like W make it hard to iterate over some free parameters with sufficient speed. The *superstage approximation* offers the means to reduce such computational complexity with a simple and interpretable method. The technique described here is analogous to the one explored in past JET [197, 208] and ITER [210] modeling. These works rely on the creation of ADAS atomic rates for superstages in a set of dedicated data files. This process has been significantly simplified in Aurora, where transformations of standard GCR coefficients are computed “on the fly”, without need for additional files.

In the superstaging approximation, multiple charge states are “bundled” together to reduce the number of species to evolve in a simulation. Suppose that we wish to group charge states i_0 to i_1 onto a superstage ζ . We then have

$$n^\zeta = \sum_{i=i_0}^{i_1} n^i \quad (\text{A.10})$$

where we have used lower case n for the density of a general atomic species. Superscripts indicate the stage or superstage of interest, in slight departure from the more common notation used in section 5.3.1. In what follows, electron density is simply labeled as n_e . The

continuity equation for a charge state i can be written as

$$\frac{\partial n^i}{\partial t} = -\nabla \cdot \Gamma^i + Q^i. \quad (\text{A.11})$$

and our objective is to obtain an equation of identical form, but with indices i replaced by ζ and each term appropriately adapted to represent superstages. Eventually, after solving for the superstage densities, we wish to split these back into individual charge states densities.

Let us assume that the transport coefficients appearing in the Γ^i term in Eq. A.11 are independent of Z ; we shall come back to this approximation later. Summing particle flux terms of multiple charge states in Eq. A.11 is then trivial, since D and v are the same for all charge states and densities linearly sum up as in Eq. A.10. The same is true for the time derivatives of particle density on the left hand side of Eq. A.11. Let us therefore focus our attention on the sources/sinks term, Q^i , which we write here for a charge state i as

$$Q^i = -(n_e S^{i \rightarrow i+1} + n_e \alpha^{i \rightarrow i-1}) n^i + n_e S^{i-1 \rightarrow i} n^{i-1} + n_e \alpha^{i+1 \rightarrow i} n^{i+1}. \quad (\text{A.12})$$

Here we have adopted a notation that makes it clear which charge states are involved in each reaction, e.g. $S^{i \rightarrow i+1}$ is the process of ionization from charge state i to $i+1$ and $\alpha^{i \rightarrow i-1}$ is recombination from charge state i to $i-1$. Upon summation of the Q^i term for each of the charge states i_0 to i_1 , many terms cancel out. The sum over these charge states can be written as

$$\begin{aligned} Q^\zeta = & + n_e S^{i_0-1 \rightarrow i_0} n^{i_0-1} \\ & - n_e \alpha^{i_0 \rightarrow i_0-1} n^{i_0} \\ & - n_e S^{i_1 \rightarrow i_1+1} n^{i_1} \\ & + n_e \alpha^{i_1+1 \rightarrow i_1} n^{i_1+1}. \end{aligned} \quad (\text{A.13})$$

where we have indicated the meaning of each term.

The purpose of superstaging is to avoid explicit modeling of individual charge states, which must instead be only related to the correspondent (directly modelled) superstages. *In the superstage approximation, we impose quasi-static ionization equilibrium for the charge states i_0 to i_1 within a superstage.* This means that the fractional abundance of each charge state within the superstage is purely set by atomic physics, with no transport involved.

Within a superstage, we therefore take

$$\begin{aligned}
n^{i_0}|_{eq} &= (\alpha^{i_0+1 \rightarrow i_0} / S^{i_0 \rightarrow i_0+1}) n^{i_0+1}|_{eq} \\
n^{i_0+1}|_{eq} &= (\alpha^{i_0+2 \rightarrow i_0+1} / S^{i_0+1 \rightarrow i_0+2}) n^{i_0+2}|_{eq} \\
&\dots \\
n^{i_1-1}|_{eq} &= (\alpha^{i_1 \rightarrow i_1-1} / S^{i_1-1 \rightarrow i_1}) n^{i_1}|_{eq}.
\end{aligned} \tag{A.14}$$

In other words, we wish to evolve n^ζ in space and time, i.e. $dn^\zeta/dt \neq 0$ in a transport code, and we take $Q^\zeta = 0$ when decomposing each superstage density into its charge state components. This obviously incurs an error, which however is small in many circumstances, as we will show later.

In order to express Eq. A.13 solely in terms of superstage properties, we recognize that both charge states i_0 and i_1 are absorbed into superstage ζ , so by applying the changes $i_0 \rightarrow \zeta$ and $i_1 \rightarrow \zeta$ one can write

$$\begin{aligned}
S^{i_0-1 \rightarrow i_0} n^{i_0-1} &\rightarrow S^{\zeta-1 \rightarrow \zeta} n^{\zeta-1} \\
\alpha^{i_0 \rightarrow i_0-1} n^{i_0} &\rightarrow \alpha^{\zeta \rightarrow \zeta-1} n^\zeta \\
S^{i_1 \rightarrow i_1+1} n^{i_1} &\rightarrow S^{\zeta \rightarrow \zeta+1} n^\zeta \\
\alpha^{i_1+1 \rightarrow i_1} n^{i_1+1} &\rightarrow \alpha^{\zeta+1 \rightarrow \zeta} n^{\zeta+1}
\end{aligned} \tag{A.15}$$

which allow us to define effective superstage rates as

$$\alpha^{\zeta \rightarrow \zeta-1} \equiv \alpha^{i_0 \rightarrow i_0-1} \left(\frac{n^{i_0}}{n^\zeta} \right)_{eq}, \tag{A.16}$$

$$S^{\zeta \rightarrow \zeta+1} \equiv S^{i_1 \rightarrow i_1+1} \left(\frac{n^{i_1}}{n^\zeta} \right)_{eq}. \tag{A.17}$$

The terms in brackets are the fractional abundances of the ‘‘edges’’ of the superstage ζ , which we explicitly labelled as being computed from ionization equilibrium. Using the superstage rates of Eqs. A.16 and A.17, one can then express the sources/sinks term for superstages as

$$Q^\zeta = -(n_e S^{\zeta \rightarrow \zeta+1} + n_e \alpha^{\zeta \rightarrow \zeta-1}) n^\zeta + n_e S^{\zeta-1 \rightarrow \zeta} n^{\zeta-1} + n_e \alpha^{\zeta+1 \rightarrow \zeta} n^{\zeta+1}. \tag{A.18}$$

which has the same form as Eq. A.12 and can be used for impurity transport simulations,

including (but not exclusively) within the 1.5D Aurora forward model.

D and v can either be set to be the same for all superstages, as commonly done in simulations with individual charge states, or users may indicate an arbitrary Z dependence. This could, for example, be obtained from a neoclassical or turbulent transport code, in which case Aurora can compute an average of transport coefficients weighted by the fractional abundance of charge states within each superstage. This corresponds to taking

$$D^\zeta = \sum_{i_0}^{i_1} D^i \left(\frac{n^i}{n^\zeta} \right)_{eq}. \quad (\text{A.19})$$

An analogous weighing by fractional abundances at ionization equilibrium allows one to decompose superstages into individual charges after a simulation:

$$n^i = n^\zeta \left(\frac{n^i}{n^\zeta} \right)_{eq}. \quad (\text{A.20})$$

We remark that while n^ζ may be significantly out of ionization equilibrium, resulting from a simulation with finite transport, the right hand side term within brackets in Eq. A.20 is taken purely from ionization equilibrium (no transport).

Once superstage densities have been unstaged, any calculation of radiation terms can proceed as usual using standard collisional-radiative coefficients, with no need for additional tabulation of atomic rates. Superstaging is therefore seen to only require charge state ionization and recombination rates. If charge exchange is considered, the ADAS CCD data are also needed for inclusion within the α rates with appropriate weighting by n_n/n_e , where n_n is the background H-isotope atomic neutral density.

The error incurred by the superstaging approximation is dependent on the intensity of plasma transport, since this technique effectively imposes ionization equilibrium within superstage partitions. For typical tokamak transport levels, there often exist multiple useful and safe partitioning strategies for medium- and high- Z impurities. It is however not possible to specify the “best” partition in a universal manner, since the optimal choice of superstages depends on the application of interest. The top panel of Fig. A-1 shows W fractional abundances of all charge states of W as a function of T_e , using dashed lines. Easily identifiable charge states have been indicated by numbers near the peak fractional abundance of each. A partition of charge states based on the valence electronic shells is also




Figure A-1: Full ionization equilibria of (top) W and (bottom) Ca over electron temperature, T_e , shown by dashed lines. Numbers within the plot identify selected ionization stages. Continuous lines show one bundling scheme for each atomic species, in the case of W based on filled electronic shells, while for Ca only low-Z stages are grouped. Figure from Ref. [186].

shown by continuous lines, corresponding to superstages $Z \in \{0\}, \{1, 2\}, \{3, 8\}, \{9, 18\}, \{19, 32\}, \{33, 50\}, \{51, 72\}, \{72, 74\}$. This partitioning is physically motivated by atomic physics and analogous to the “natural partition” strategy defined in ADAS, whereby charge states are bundled based on variations of ionization potential (I_z) between charge states. Specifically, the natural partition is computed for each species by finding large deviations of the quantity $2(I_{z+1} - I_z)/(I_{z+1} + I_z)$ from a running mean over charge states [208]. To this end, Aurora can make use of ionization energies from NIST [136], collected via ColRadPy [120]. Alternatively, users may also use fractional abundances at ionization equilibrium to determine which charge states have a small range of existence in a given simulation and can therefore be safely bundled. Such “clustering” of superstages can be easily dealt with using the K-means algorithm.

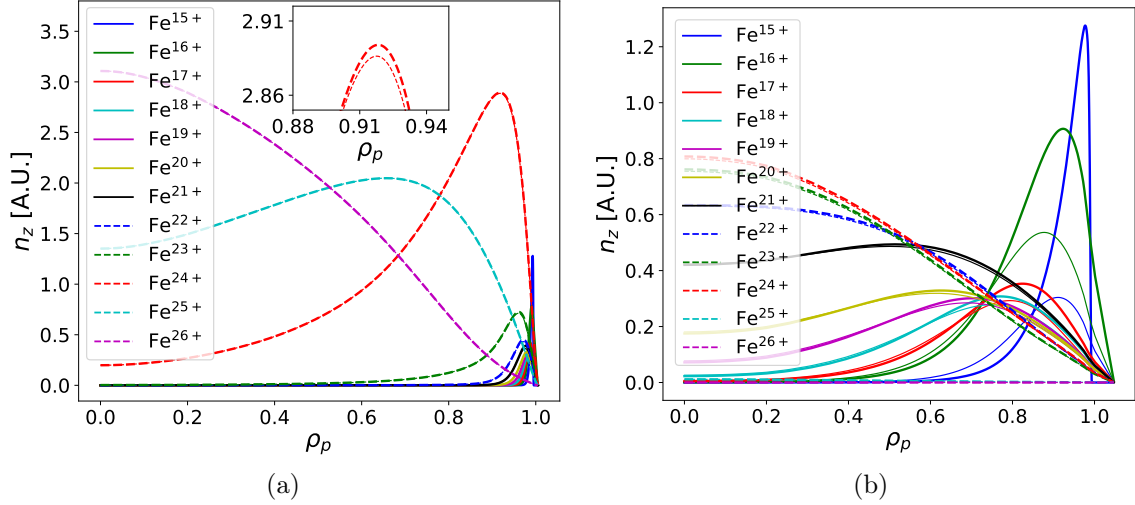


Figure A-2: Comparison of Aurora simulations for Fe transport in (a) the ITER baseline scenario ($T_e \approx 25$ keV and $n_e \approx 1.1 \times 10^{20} \text{ m}^{-3}$ on axis), and (b) an Alcator C-Mod discharge with $T_e \approx 3.5$ keV and $n_e \approx 1.5 \times 10^{20} \text{ m}^{-3}$ on axis. Transport coefficients $D = 10 \text{ m}^2/\text{s}$, flat across the radius, and $v = 0 \text{ m/s}$ were assumed for illustrative purposes. Thin lines show results from a simulation resolving all charge states, thick lines from bundling $Z < 16$ states. Figures adapted from Ref. [186].

However, often the optimal partitioning of charge states may be specific to the application of interest. For example, for edge simulations, charge states with charge greater than $\text{W}20+$ are unlikely to be useful and can always be bundled. Similarly, core simulations may bundle low- and high-lying charge states that are never reached at the temperatures of interest, only simulating charge states that produce observable effects. We note that the 0th charge state is always needed in forward modeling where ion sources are introduced in the neutral state. The bottom panel of Fig. A-1 shows a case where we bundle charge states $\{1, 9\}$ and keep individual charge states with $Z > 9$. This choice is favorable for core impurity transport studies where Ca is injected into tokamak plasmas via Laser Blow Off (LBO) since emission from charge states 1-9 is typically not measured and only higher charge states must be carefully resolved. In tests of the Aurora forward model with this partitioning strategy, we have found that a negligible error in impurity densities and radiation is incurred for any realistic transport levels. Recall that the computational cost of the Aurora forward model is linearly proportional to the number of modelled charge states, so superstaging as in Fig. A-1 gives approximately a factor of 2 speed improvement.

Fig. A-2a shows two Aurora simulations for Fe injection in the ITER baseline scenario, where peak electron temperatures and densities reach $T_e \approx 25$ keV and $n_e \approx 1.1 \times 10^{20} \text{ m}^{-3}$.




Figure A-3: Application of superstaging within an Aurora simulation of W in the ITER baseline scenario. In (a), thin lines show the result of simulating all charge states, while thicker lines show the result of superstaging. Here, only $Z = \{50, 53, 56, 59, 62, 65, 68\}$ were superstaged, using $D = 10 \text{ m}^2/\text{s}$, $v = 0 \text{ m/s}$. In (b), we show the difference between the two cases for cumulative radiated power, integrated from the magnetic axis outwards, in the pedestal region. Figures adapted from Ref. [186].

We compare the standard use of all charge states (thin lines) with a superstage partition that bundles Fe charge states with $Z < 16$ (thick lines), analogously to the Ca strategy of Fig. A-1. Almost no difference is observed in the charge state densities of interest; the inset shows a zoomed in view of the peak of Fe^{24+} near the pedestal, where a small inaccuracy is seen. For this comparison, a flat diffusion coefficient with magnitude $D = 10 \text{ m}^2/\text{s}$ was assumed, setting convection to zero for simplicity. We note that the value of $D = 10 \text{ m}^2/\text{s}$ is relatively high and, since superstaging becomes less accurate at high transport levels, this can be seen as a successful test in challenging, yet reasonable, conditions for the application of this method. We remark however that the same partitioning strategy can be much less successful in other plasma scenarios. Fig. A-2b shows the result of using the same Fe superstaging strategy for realistic Alcator C-Mod kinetic profiles with on-axis values of $T_e \approx 3.5 \text{ keV}$ and $n_e \approx 1.5 \times 10^{20} \text{ m}^{-3}$, much lower than in the ITER baseline. As in Fig. A-2a, thin lines correspond to the simulation resolving all charge states, while thicker lines show the result of bundling all $Z < 16$ states (except the neutral stage, as always). While core profiles are recovered very well, significant discrepancies are found in the outer part of the plasma. Of course, even larger errors would be incurred if a smaller number of superstages were directly modelled, e.g. taking only every other Fe state between $Z = 16$ and $Z = 26$.

Many of the envisioned applications of superstaging are with high-Z ions. Fig. A-3a shows two Aurora simulations of W in the ITER baseline scenario. The displayed charge states were the only ones that were explicitly modelled, i.e. all others were bundled. Except for small inaccuracies near axis, the agreement between thin and thick lines is excellent, particularly considering the large reduction of computational cost of this superstaging partition. Fig. A-3b shows the cumulative radiated power computed by Aurora as a function of plasma radius when setting the total W concentration to be 1.5×10^{-5} on axis [15]. The radial range is focused on the pedestal region in order to better visualize the loss of accuracy due to superstaging, shown by the difference between the full model, simulating all charge states (red) and the reduced superstaged model (blue). Clearly, such simulation strategy could be very advantageous for a number of applications, including integrated modeling where impurity radiation can play a significant role. In these cases, modeling impurities as having charge state densities set by a constant fraction of the electron density can result in significant errors, whereas running Aurora with a small number of superstages, appropriately chosen for the scenario of interest, can be fast and effective.

A.4 Charge Exchange in the ITER Pedestal

Aurora’s routines for the post-processing of SOLPS-ITER [204, 205] results enable consideration of core-edge effects, for example related to the influence of edge neutrals in the pedestal region [253]. Recent work by Dux *et al* [228] demonstrated the large impact of charge exchange on ionization balance in the AUG pedestal, making use of detailed charge exchange recombination spectroscopy measurements of Ne⁸⁺ and Ne¹⁰⁺. A simple 1D Monte Carlo neutral code was used to model edge neutral penetration. An extrapolation to ITER based on these AUG observations suggested that CX could significantly affect the total radiated power emitted within the LCFS, P_{rad} , increasing core Ne radiation by up to a factor of 5. Here, we provide a new assessment of this subject based on EIRENE modeling of neutral penetration in ITER using the SOLPS-ITER code, which is a higher-fidelity model with respect to previous descriptions.

The details of these SOLPS-ITER results, including realistic levels of Ne puffing for divertor heat flux mitigation, have been presented elsewhere as part of ITER divertor studies [292]. Here, we limit ourselves to a demonstration of post-processing and




Figure A-4: Radial profiles of atomic neutral D density, extracted from the SOLPS-ITER results displayed in the inset, showing interpolated slices at the low-field side (LFS) and high-field side (HFS) midplane, as well as a flux surface average. The dashed black line identifies the LCFS and the dashed red line shows the level of $n_n/n_e = 10^{-5}$, below which CX does not significantly affect the local radiated power. Figure from Ref. [186].

interpretation of the SOLPS-ITER data using Aurora. By providing the path to the SOLPS-ITER results directory, including an EFIT gEQDSK file, Aurora can load data on both the EIRENE and B2 grids with just a couple of lines of Python code. An additional call to the `aurora.solps.get_radial_prof` function allows extraction of low- (LFS) and high-field side (HFS) radial profiles at the midplane, as well as flux surface averages. Fig. A-4 shows these atomic D neutral density profiles, normalized by the corresponding n_e radial profiles (LFS, HFS and FSA). We note that FSA neutral profiles are appropriate inputs to Aurora’s 1.5D impurity transport model to examine the effect of neutrals on impurity ionization balance via CX. This is a physical modeling choice everywhere except very close to the LCFS (typically, for $\rho_p < 0.99$) since the parallel motion of impurities on flux surfaces tends to be much faster than their ionization. As a result, the impact of neutrals on impurity ions is effectively averaged out also in the presence of strong poloidal neutral asymmetries, e.g. near divertor x-points or neutral beam injection [228]. Fig. A-4 shows that n_n/n_e is expected to be lower than 10^{-5} , i.e. below the red dashed line. At the realistic Ne densities of the ITER baseline scenario, Aurora predicts the contribution to P_{rad} directly attributable to CX recombination to be only approximately 0.2 MW. The indirect contribution from the modification of the Ne ionization equilibrium due to CX cannot be self-consistently assessed without turning CX on and off in the SOLPS-ITER simulation itself, but at the low values of n_n/n_e this effect is undoubtedly very small. These conclusions can be interpreted as resulting from the large

width of the ITER SOL, which significantly reduces neutral penetration into the confined region. We note that while Alcator C-Mod's highest density shots did reach similar opacity to high-performance ITER projections [80], the C-Mod SOL was significantly thinner. As a result, C-Mod's absolute neutral densities in the pedestal were lower than in any other current tokamak, but still higher than in ITER, whose core fueling is bound to have different global character.

Appendix B

The Bayesian Spectral Fitting Code

In this appendix we describe spectral fitting methods that were developed for high-resolution x-ray spectroscopy and applied in Alcator C-Mod research presented in Chapter 7. Spectral lines are efficiently decomposed using Gauss-Hermite functions. Bayesian techniques discussed in the context of particle transport are also applied here. Applications with both synthetic and real XICS data show improvements with respect to pre-existing methods.

Introduction

Spectroscopy plays a central role in the research presented in this thesis. In this appendix we describe Bayesian spectral fitting methods that were developed to analyse high-resolution spectra from Alcator C-Mod, discussed in Chapter 7. These techniques make use of some of the Bayesian concepts described in Chapter 5 and have offered a test bed for ideas that were later adopted in inferences of particle transport described in Chapters 7 and 8.

Complex atomic spectra are common in plasma physics research, both in astrophysical and laboratory scenarios. As discussed in Chapter 3, local spectra where Doppler broadening is the dominant mechanism determining line shapes may allow one to easily extract valuable physical parameters (emissivity, velocity and temperature) from simple moments of Gaussian line fits. However, local radiation measurements are often not available and a purely Gaussian fit is often not justified. The viewing geometry, causing integration of signals over extended volumes, and the overlap of multiple atomic line shapes make data interpretation often remarkably challenging. In the extreme ultraviolet and for highly-ionized impurities

it is common for natural broadening to be significant [115], giving rise to a Voigt profile, the convolution of the Gaussian and a Lorentzian. Other broadening mechanisms may also contribute to non-Gaussianity, thus motivating the use of more flexible fitting methods.

Here we present the Bayesian Spectral Fitting Code [254] (BSFC) as an efficient means to fit multiple, partially-overlapping, non-Gaussian atomic line shapes in high-resolution spectroscopy. BSFC uses a truncated decomposition into Gauss-Hermite functions and Bayesian sampling techniques to separate lines and estimate spectral moments with the highest accuracy afforded by experimental data. This is particularly useful in the context of fitting highly-resolved K_α spectra, as will be described in section B.3 and later in chapter 7. In the next sections, we first summarize BSFC’s key features and then compare BSFC to pre-existing spectral fitting methods. The successful application of these techniques to Alcator C-Mod spectra has led to investigations on impurity transport presented in Refs. [155] and [156].

B.1 Hermite Polynomial Decomposition

BSFC’s polynomial decomposition exploits the fact that atomic lines, when relatively spectrally resolved, have line shapes that approximately resemble a Gaussian. Such resemblance may be more or less close, depending on the balance of broadening mechanisms, but here we assume that data follow a Gaussian, Lorentzian, Voigt or analogous line shape. It is then convenient to take the Gaussian component of this as a zeroth order description, and subsequently apply correction terms to effectively match the observed line shape. This strategy is realized with an expansion in Hermite polynomials, which form a complete, orthogonal basis under a Gaussian weight function, with

$$\langle He_m, He_n \rangle_G \equiv \int_{-\infty}^{\infty} He_m(x) He_n(x) e^{-\frac{x^2}{2}} dx = \sqrt{2\pi} n! \delta_{nm} \quad (\text{B.1})$$

Here, we have used the He_i notation to indicate a (probabilist’s) Hermite polynomial of order “i”. The meaning of the angle-bracket notation with “G” subscript indicates use of a zero-centered, unit-variance Gaussian factor, shown in the integral expansion over all space. The first three Hermite polynomials are $He_0(x) = 1$, $He_1(x) = x$, and $He_2(x) = x^2 - 1$; they become more complex as one goes to higher order. To better understand BSFC’s strategy,

it is helpful to consider Gauss-Hermite functions, defined as

$$\phi_n(x) = He_n(x) e^{-x^2/2}, \quad (\text{B.2})$$

represented for convenience and to aid intuition in Fig. B-1.




Figure B-1: Gauss-Hermite functions of order $n = 0 - 5$. BSFC expands spectral line shapes into a truncated series of terms composed of progressively more complex features represented by these functions. Figure reproduced from Ref. [156].

Since Hermite polynomials form a complete basis, a function $f(x)$ can be expanded in terms of Gauss-Hermite functions as

$$f(x) = h(x)e^{-x^2/2} = \left(\sum_{j=0}^{\infty} a_j He_j(x) \right) e^{-x^2/2} = \sum_{j=0}^{\infty} a_j \phi_j(x). \quad (\text{B.3})$$

In what follows, we refer to a_j variables as *Hermite coefficients*. Using this expansion, one can easily compute the first few unnormalized moments for each spectral line shape by taking the inner product of Eq. B.1 with x , x^2 , etc., to obtain progressively higher-order moments, e.g.

$$\begin{aligned} m_0 &= \langle 1, h \rangle_G = \sqrt{2\pi} a_0 \\ m_1 &= \langle x, h \rangle_G = \sqrt{2\pi} a_1 \\ m_2 &= \langle x^2, h \rangle_G = \sqrt{2\pi} a_0 + 2\sqrt{2\pi} a_2 \end{aligned} \quad (\text{B.4})$$

Here we used again the angle-bracket notation with “G” subscript to indicate a projection (integration over all space) under a Gaussian weight of zero mean and unit variance. Note that each moment in Eq. B.4 only refers to Hermite coefficients of lower order, e.g. it is never the case that a n th moment requires calculation of a m^{th} order Hermite coefficient, with $m > n$. This observation is key to explain BSFC’s effectiveness.

To examine atomic spectra, it is convenient to define a variable $x = (\lambda - \lambda_c)/s$, a simple shift of the wavelength axis to center it around λ_c and rescale by a factor s . One can then express physical moments of line shapes across a spectrum using

$$\begin{aligned} M_0 &= m_0 s \\ M_1 &= m_1 s^2 + M_0 \lambda_c \\ M_2 &= m_2 s^3 \end{aligned} \tag{B.5}$$

and similarly for higher-order moments. The value of M_0 is related to the emissivity (or brightness, for line-integrated measurements), which is itself linearly proportional to the emitting species’ density, the constant of proportionality being dependent on atomic physics for the transition of interest. M_1 and M_2 are related to the species’ velocity and temperature by

$$v_i = \frac{M_1/M_0}{\lambda_0} c \qquad T_i = \frac{M_2/M_0}{\lambda_0^2} m c^2 \tag{B.6}$$

Finite velocity and temperature are obtained even if $a_1 = a_2 = 0$, but non-zero values of these Hermite coefficients are seen from Eqs. B.5 and B.6 to provide corrections to perfect-Gaussian estimates. Importantly, even higher order Hermite coefficients do not enter the estimation of v_i and T_i at all, suggesting that a truncation to 3 Hermite polynomial terms may be optimal to capture non-Gaussian corrections to velocity and temperature estimates. This was indeed shown to be the case in Ref. [254] using Bayesian model selection (see section B.2).

BSFC’s expansion offers the additional benefit that truncated (i.e. ignored) terms do not affect Hermite coefficients that are kept within the fitting procedure of each line individually. To explain this, let us write the truncation process as

$$f(x) = \sum_{j=0}^{\infty} a_j \phi_j(x) \quad \approx \quad \hat{f}(x) \equiv \sum_{j=0}^N \hat{a}_j \phi_j(x). \tag{B.7}$$

Due to the complete and orthogonal nature of the Hermite polynomial series, we have

$$\langle f(x) - \hat{f}(x), He_j \rangle = 0 \quad \text{for } j = 0, \dots, N \quad (\text{B.8})$$

i.e. any terms that are left out by our truncation are orthogonal to the terms that are kept within the series, meaning that their truncation does not affect fitting parameter estimates. In practice, BSFC fitting of multiple overlapping lines can introduce weak correlations between Hermite coefficients, so this ideal truncation property only holds approximately. Nonetheless, these elements of strength make BSFC’s spectral decomposition robust to truncation of higher order corrections, offering a way to study deviations from perfect-Gaussian line shapes with a minimal number of fitting parameters and maximum accuracy for a given computational effort.

B.2 Sampling Methods

The estimation of Hermite coefficients, as well as scale and shift factors for the spectrum, is dealt with using the Bayesian techniques described in Chapter 5, analogously to the impurity transport research presented in later chapters. We model the likelihood of our fit as a Gaussian, taken as the large-count limit of photon Poisson statistics. We adopt uninformative priors, taking care to set log-uniform priors for scale parameters that can vary over multiple orders of magnitude, like zeroth order Hermite coefficients. BSFC enforces non-negativity of line fits by uniformly sampling from a unit hypercube within parameter simplices; details of this procedure are described in Appendix A of Ref. [254]. By using nested sampling [191] with `MultiNest` [193, 192], BSFC is able to do both parameter estimation and model selection via the marginal likelihood. In the context of BSFC, “different models” are represented by different choices of number of Hermite coefficients, or number of spectral lines within the spectrum. For example, whenever it is not clear whether a certain line is present within an experimental spectrum, comparison of the model evidence between fits with and without this line can offer a quantitative metric to make a data-driven choice.

Typical BSFC spectral fits include 1-4 overlapping lines, requiring 3-4 Hermite coefficients per line as well as 1 line shift and 1 line rescaling value, for a total of 5-20 fitting (free) parameters per case. For such problem dimensionality, we have found the `MultiNest` ellipsoidal NS algorithm [193, 192] to be the most effective. For higher dimensionality, e.g.

for spectra with > 10 overlapped line shapes, we have also tested dynamic slice NS, as implemented in the `dyPolyChord` package [293], which is far more effective than `MultiNest` at very high dimensionality. In practice, in all tests of BSFC we have found `MultiNest` to be the best algorithm for our fits, since it is often reasonable to fit different parts of a spectrum separately if there are no evident line overlaps.

B.3 Applications to XICS Spectra

As described in Section 3.5.1, K_α spectra present a complex overlap of spectral features whose intensity and line shapes reflect the electron density and temperatures (via the amplitude), the ion temperature (from the Doppler broadening), the rotation velocity (from the Doppler shift), as well as the abundance of the Li-like, He-like and H-like charge states of the emitting ion. BSFC was originally developed to effectively fit features arising from line integration of local spectra such as the one shown in Fig. 3-4a. For example, the long-wavelength side of K_α spectra presents a near-degeneracy of the forbidden z line and the j satellite line, whose brightness is constrained via the nearby k satellite line [294]. Fig. B-2 shows an application of BSFC to this spectral region on real XICS data on C-Mod. Not only is the separation of z, j and k lines effectively dealt with by BSFC, but uncertainties (shown by the colored bands) clearly offer an appropriate representation of our confidence in these observations. In the next sections, we first verify BSFC’s implementation using synthetic data and then compare it to pre-existing spectral fitting methods used for XICS analysis of real data from Alcator C-Mod.




Figure B-2: BSFC fit of the He-like Ca z line (green), overlapping with the j satellite line (red), which is related to the nearby k line (blue). The overall fit is shown by the continuous black line; residuals are displayed in the lower panel. Rest wavelengths are shown by vertical black dashed lines. Figure reproduced from Ref. [156].

Synthetic Data Analysis

To validate the results of the fitting code, synthetic tests were performed. To generate experimentally-relevant spectral line shapes, synthetic emission was line-integrated using impurity emissivity, velocity, and temperature profiles inferred using standard tomographic methods. The true moments were calculated exactly using line-integration of the synthetic data. Two cases are presented here: one from an L-mode tokamak plasma in Fig. B-3a, and one from an I-mode with a hollow emissivity profile in Fig. B-3b.




Figure B-3: Comparison between true synthetic line-integrated measurements (blue) and measurements inferred from BSFC (green) for (a) a L-mode plasma and (b) a I-mode plasma with hollow emissivity profile. Figure reproduced from Ref. [254].

These fits were performed on the He-like argon $n = 2$ resonance w line with 3 Hermite coefficients. In both cases, the unresolved dielectronic $n = 4$ satellite lines in our atomic database were removed from the fit, since the Bayesian evidence obtained for a fit without it was found to be higher than with it. This assessment rigorously determines whether the data justify a certain level of complexity in a fitting model. As seen from these plots, the method infers the true moments very well. Moreover, the uncertainty in the inference clearly matches the spread of the data. The non-Gaussian shape of the line resulting from higher-order Hermite coefficients typically accounted for a correction of approximately 10%

to the temperature, and 20% to the velocity.

Comparison to Pre-existing Methods

Results from BSFC were also compared to results inferred from THACO [180], the standard code used to analyze XICS data in Alcator C-Mod. THACO estimates line moments by using a nonlinear χ^2 minimization to remove satellite lines, followed by a direct calculation of line moments by integration. We remark that this method is typically only applied on non-overlapping spectral lines to avoid complications that are directly addressed by BSFC. To illustrate the advantages of BSFC in these cases, we compare its results to those of THACO in two cases with overlapping lines: in the first case, one has a steady-state argon He-like w resonance line spectrum as seen in Fig. B-4a; in the second, we have a rapidly-evolving calcium He-like w resonance line spectrum from a laser-blow-off impurity injection Fig. B-4b.

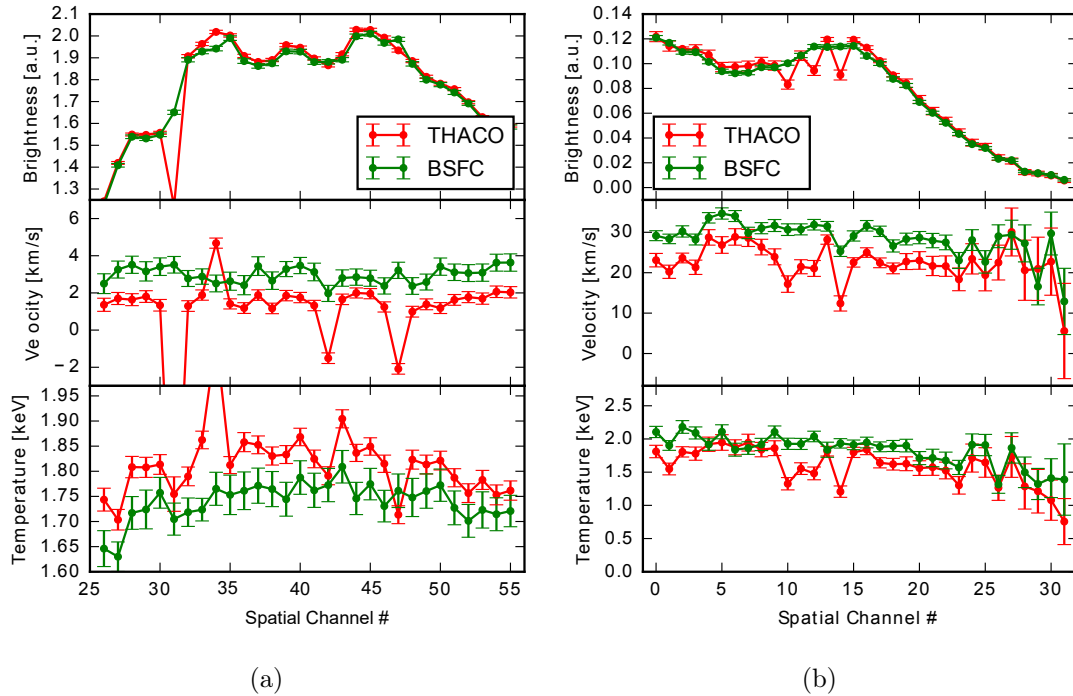


Figure B-4: Comparison between line-integrated measurements inferred from BSFC (green) and from THACO (red) for (a) a subset of XICS channels following Ar injection into an L-mode plasma, and (b) a Ca injection into an I-mode plasma with hollow brightness profiles. Figure reproduced from Ref. [254].

In both cases, we see that traditional Gaussian line fitting methods can produce outliers. The BSFC method provides more consistent results and the error bars also more realistically

reflect the spread of data. Moreover, we note that BSFC finds a systematic shift of velocity and temperature measurements compared to the THACO inference. The former inference was validated against synthetic spectra with unresolved lines and shown to be more accurate.

B.4 Summary

The Bayesian Spectral Fitting Code offers novel methods that are well adapted to the task of inferring line-integrated spectral moments from spatially non-uniform emission, as in the case of highly-resolved XICS spectra analyzed in this thesis. BSFC does not depend on prior knowledge of line ratios or precise determination of plasma geometry, offering an unbiased method to quantify spectral moments and their uncertainties using Bayesian nested sampling techniques. In previous sections, we have shown that BSFC performs well on synthetic data and clearly outperforms the pre-existing analysis methods used for XICS spectral fitting on Alcator C-Mod. The adopted sampling techniques permit greater interpretability of results than nonlinear optimizers, allowing correlations between inferred parameters to be explored. BSFC also enables model selection to best identify an appropriate complexity of the fit, in terms of number of lines and number of Hermite coefficients that should be used to describe each line shape. In Chapter 7, BSFC is applied to LBO injections of Ca to isolate individual lines, before exploring a different approach based on forward modeling the entire K_α spectrum.

Appendix C

Bayesian combination of experimental datasets

Experimental inferences often rely on multiple pieces of instrumentation, each providing a different measurement of fundamentally the same phenomenon. For example, this occurs when multiple detectors look at the dynamics of an object at different times/locations, all of these related to the same motion that one wants to reconstruct. It also occurs in astrophysical environments, where multiple telescopes may investigate the same object, but over different spectral ranges. The combination of different datasets is also important in fusion energy research, where diagnosing the full dynamics of a thermonuclear plasma is challenging from an experimental perspective. In this field, one often must rely on multiple measurements each telling a part of the full story, and none being a complete representation of the physics of interest.

When solving inverse problems, one often combines diagnostic measurements assuming that they are equally valid and important. This means that no instrument is considered more important than others, except to the extent suggested by the uncertainty quantification (UQ) of signals. UQ can be challenging and, although one may resolve to work on a complex analysis, random errors can be left unaccounted for, and systematic errors are always looming. The most common approach is simply to ignore this, once appropriate efforts have been undertaken to correct signals as much as possible.

Let us assume that uncertainties can be modelled as normally distributed around experimental signals. Given a set of signals and their uncertainties, the standard approach

involves maximizing the likelihood probability, or equivalently minimizing the chi-squared (χ^2). For a single dataset with n data points, one can write the likelihood as

$$P(D|\theta) = \frac{1}{(2\pi)^{n/2}|V|^{1/2}} \exp\left(-\frac{1}{2}\chi^2\right) \quad (\text{C.1})$$

where

$$\chi^2 = (D - f(\theta))^T V^{-1} (D - f(\theta)). \quad (\text{C.2})$$

with the function $f(\theta)$ representing the forward model, evaluated on some set of parameters (note: we omit vector notation for simplicity). Here, D represents the experimental data points, and V is the covariance matrix. While we made the dependence of the forward model on θ explicit, we are going to assume for simplicity that the covariance matrix does not depend on θ (a good approximation under most circumstances).

Maximizing this likelihood corresponds to minimizing the χ^2 . If multiple datasets are taken into consideration (e.g. from different fusion diagnostics), and we label these with an index k , then we can write

$$P(D|\theta, H_0) = \prod_{k=1}^N \frac{1}{(2\pi)^{n_k/2}|V_k|^{1/2}} \exp\left(-\frac{1}{2}\chi_k^2\right), \quad (\text{C.3})$$

where we indicated a dependence on a model labelled H_0 here, to distinguish it from different options to be described later in the text.

It is often convenient to work with logs of likelihoods and ignore constant factors in an optimization. In view of this, one often minimizes

$$\chi_t^2 = \sum_k \chi_k^2. \quad (\text{C.4})$$

This is effectively stating that uncertainties have been quantified to the best possible extent and that all diagnostic signals have equal worth. However, it is also common to modify this expression to read

$$\chi_t^2 = \sum_k \alpha_k \chi_k^2 \quad (\text{C.5})$$

where we included factors α_k that can be arbitrarily set to more or less than 1 in order to bias an inference towards better fitting one experimental signal or another. This is common practice in fusion as in many other fields. Of course, this approach can be useful, although

obviously unsatisfactory. One may also choose a scheme to set weights based on the number of data points in each dataset, as discussed in Section 7.2.5 in the context of Alcator C-Mod impurity transport inferences.

In the following sections, we discuss a generalization of the practice of fixing weights α_k *a priori*, before the data is examined. We present a Bayesian approach for this task, expanding on the methods described in Refs. [218, 219, 220].

C.1 Bayesian Priors for Diagnostic Weights

In Bayesian analysis, we always want to assign a probability distribution to parameters that we do not know. The appropriate values of α_k are unknown and therefore we must choose a prior distribution for them. We can do this on statistical grounds, while remaining “unbiased” towards one dataset or another. The appropriate distribution should of course be normalized to 1, and we would also like its expectation value to be 1, so that the values of α_k will be 1 on average, in a statistical sense. Following the recipe by Jaynes, Shannon and others, we maximize the entropy functional

$$S[P(\alpha)] = - \int_0^\infty P(\alpha_k) \ln P(\alpha_k) d\alpha_k \quad (\text{C.6})$$

under the constraint of normalization

$$\int_0^\infty P(\alpha_k) d\alpha_k = 1 \quad (\text{C.7})$$

and unit expectation

$$E[\alpha_k] = \int_0^\infty \alpha_k P(\alpha_k) = 1. \quad (\text{C.8})$$

This constrained optimization can be solved analytically using Lagrange multipliers and gives an exponential solution

$$P(\alpha_k) = \exp(-\alpha_k), \quad (\text{C.9})$$

which is therefore the appropriate prior representing our state of knowledge about α_k . Note that this exponential distribution has a decay parameter of 1, corresponding to an expectation value of 1. As pointed out in Ref. [219], if one were interested in adding a constraint on the variance of α_k , for example because the range of acceptable values of α_k

is limited to numbers close to 1, the same maximum entropy recipe could be followed while placing an additional constraint. This can only be analytically solved when allowing values of α_k to be negative. In practice, such nonphysical condition could be reasonable if the prior support for $\alpha_k < 0$ is minimal. By extending the domain of α_k to $[-\infty, +\infty]$, one can easily find the familiar Gaussian distribution. While this form of prior may be useful, in practice we observe that the analytical integrals described in the following section become complex and the computation may lead to overflows. In the following section, we will simply consider the exponential prior in Eq. C.9, but this will be generalized in Section C.2.

Modification of the Likelihood function

As stated previously, we want to introduce α_k factors in front of each χ_k^2 , in order to account for diagnostics whose data are more precious. We want to allow the data themselves and the forward model to define what this really means. We can write the weighing of χ^2 's as

$$\ln P(D_k|\theta, \alpha_k) = \alpha_k \ln P(D_k|\theta, \alpha_k = 1) - \ln Z_k(\theta, \alpha_k). \quad (\text{C.10})$$

This corresponds to

$$P(D_k|\theta, \alpha_k) = \frac{[P(D_k|\theta, \alpha_k = 1)]^{\alpha_k}}{Z_k(\theta, \alpha_k)} \quad (\text{C.11})$$

where Z_k is a normalization factor that we must carefully treat for Bayesian model comparison and selection. At this stage, the α_k values are arbitrary. At this point, let us define a likelihood function that is marginalized over the possible values of α_k , taking into consideration the rigorously-defined prior distribution from the previous section:

$$P(D_k|\theta) = \int_0^\infty \frac{[P(D_k|\theta, \alpha_k = 1)]^{\alpha_k}}{Z_k(\theta, \alpha_k)} e^{-\alpha_k} d\alpha_k. \quad (\text{C.12})$$

Note that this is still a likelihood function, although it was formally integrated over one of the variables (α_k). One might call it the “ α -marginalized likelihood”, to distinguish it from the full “marginalized likelihood”, which is often used as a synonym for “Bayesian evidence”.

Two models at comparison

Let us compare 2 models: H_0 , which is our standard model without variable diagnostic weights, and H_1 , which includes weights that are marginalized over in the likelihood as seen in the previous section. The total likelihood function for H_0 can be written as in Eq. C.3, which we rewrite here for convenience:

$$P(D|\theta, H_0) = \prod_{k=1}^N \frac{1}{(2\pi)^{n_k/2} |V_k|^{1/2}} \exp\left(-\frac{1}{2}\chi_k^2\right), \quad (\text{C.13})$$

On the other hand, the likelihood for H_1 is

$$P(D|\theta, H_1) = \prod_{k=1}^N \frac{1}{(2\pi)^{n_k/2} |V_k|^{1/2}} \alpha_k^{n_k/2} \exp\left(-\frac{1}{2}\alpha_k \chi_k^2\right). \quad (\text{C.14})$$

Following the recipe of Eq. C.12, we multiply by the exponential prior of α_k and integrate over α_k :

$$P(D|\theta, H_1) = \prod_{k=1}^N \frac{1}{(2\pi)^{n_k/2} |V_k|^{1/2}} \int_0^\infty \alpha_k^{n_k/2} \exp\left(-\alpha_k\left(\frac{1}{2}\chi_k^2 + 1\right)\right) d\alpha. \quad (\text{C.15})$$

Using the definition of the Gamma function $\Gamma(n) = \int_0^\infty x^{n-1} e^{-x} dx$, one obtains

$$P(D|\theta, H_1) = \prod_{k=1}^N \frac{2\Gamma\left(\frac{n_k}{2} + 1\right)}{\pi^{n_k/2} |V_k|^{1/2}} (\chi_k^2 + 2)^{-\left(\frac{n_k}{2} + 1\right)}. \quad (\text{C.16})$$

Now, compare the form of Eq. C.13 and Eq. C.16. Marginalizing over the exponential of α_k has changed the exponential dependence of the likelihood on the χ^2 for each diagnostic into a power-law dependence. Moreover, we note that the normalizing factors in the front are, of course, very different in the two cases.

For simplicity, we compare the logs of these two likelihood functions. For H_0 , we have the familiar form

$$-2 \ln P(\theta|D, H_0) = \sum_{k=1}^N \chi_k^2 + c \quad (\text{C.17})$$

while for H_1

$$-2 \ln P(\theta|D, H_1) = \sum_{k=1}^N (n_k + 2) \ln(\chi_k^2 + 2) + c'. \quad (\text{C.18})$$

Here, c and c' are constants with respect to θ and are therefore not involved in the maximization of the likelihood function for parameter estimation. However, in order to compare the Bayesian evidence for model H_0 and H_1 , we need to take into consideration some of the terms that are currently hidden in these constants. Note that model comparison relies only on the ratio of the Bayesian evidences (the so-called ‘‘Bayes factors’’), so any multiplicative factors that are not functions of θ in the log-likelihoods that appear both in H_0 and H_1 can be ignored. In other words, if we write such multiplicative factors as β , the Bayes factor B can be written as

$$B = \frac{\mathcal{Z}(\mathcal{D}|\mathcal{H}_l)}{\mathcal{Z}(\mathcal{D}|\mathcal{H}_\infty)} = \frac{\int \beta P'(\theta|D, H_0)P(\theta|H_0)d\theta}{\int \beta P'(\theta|D, H_1)P(\theta|H_1)d\theta} = \frac{\int P'(\theta|D, H_0)P(\theta|H_0)d\theta}{\int P'(\theta|D, H_1)P(\theta|H_1)d\theta}.$$

Hence, we need to compute likelihood functions for the models H_0 and H_1 such that the same factors β are present in both and can be cancelled, while all other normalization factors must be taken into consideration for an appropriate evidence comparison. Writing the log-likelihood for H_0 as

$$\ln P(D|\theta, H_0) = \sum_{k=1}^N \left[-\ln\left(\pi^{n_k/2}|V_k|^{1/2}\right) - \ln\left(2^{n_k/2}\right) - \frac{1}{2}\chi^2 \right] \quad (\text{C.19})$$

we see that the first 2 terms are just constants in our case (here, n_k is the number of data points for diagnostic k). Let us define then

$$\beta_k = -\ln\left(\pi^{n_k/2}|V_k|^{1/2}\right) - \ln\left(2^{n_k/2}\right) \quad (\text{C.20})$$

Assuming that all diagnostic measurements are independent, i.e. that the likelihoods of each are multiplicative, we can cancel this term from the correspondent partially-marginalized log-likelihood for H_1 . After some algebra, one obtains

$$\ln P(D|\theta, H_1) = \sum_{k=1}^N \left[\ln \Gamma\left(\frac{n_k}{2} + 1\right) - \left(\frac{n_k}{2} + 1\right) \ln\left(\frac{\chi_k^2}{2} + 1\right) \right]. \quad (\text{C.21})$$

Note that this form has been written in a way to facilitate the computation of the Γ function of a large number by taking its log. Programming languages often have routines to approximate the value of the log of a *Gamma* function to avoid overflow. To conclude this section, we remark that the log-evidence from an inference that maximizes this log-likelihood

is comparable to one that maximizes the log-likelihood

$$\ln P(\theta|D, H_0) = -\frac{1}{2} \sum_{k=1}^N \chi_k^2 \quad (\text{C.22})$$

for model H_0 .

Effective weights

The maximization of the log-likelihood function for H_1 in Eq. C.21 leads to optimal values of θ , which we denote by $\hat{\theta}$. However, the process of parameter estimation does not provide single values of α_k for each diagnostic. Nonetheless, one can obtain effective weights by differentiating the log-likelihood for H_1 with respect to α_k and setting the result to 0. This gives

$$\alpha_k^{eff}(\theta) = \frac{n_k}{\chi_k^2(\theta)}. \quad (\text{C.23})$$

This expression is most interesting and meaningful when estimated at the end of an inference, with $\theta = \hat{\theta}$. The final values are expected, in a statistical sense, to be close to 1, but they may differ significantly depending on the problem at hand.

We note in concluding that there are multiple possible interpretations to what these α -weights mean in practice. Of course, they can be interpreted as giving more or less importance to one diagnostic vs. the others, but from Eq. C.14 they can also be seen to be factors that can increase or reduce experimental uncertainties. Therefore, by allowing for an inference of α_k^{eff} , we are inferring modifications to the given uncertainties in the data. While these interpretations are equally valid, understanding the mathematical formalism that led to these results is likely the clearest way of understanding the meaning of this technique.

C.2 Generalization of Constrained Weights Prior

The analysis above deals with the simplest prior that obeys our rigorous mathematical constraints, the latter taking the form of an exponential function. However, as mentioned in Section C.1, one may want (or need) to constrain the width (variance) of the prior on α_k around reasonable values, which are usually taken to be 1. The addition of this constrain to the maximum entropy argument leads to a Gaussian prior, after extending the integration

domain to nonphysical values of $\alpha_k < 0$. However, obtaining an analytical expression for the partially-marginalized likelihood, as in the previous section, is challenging. Here we present a simpler alternative that leads to further generalization.

Recall that the exponential distribution is a special instance of the Gamma distribution, given by

$$p(x|a, b) = \frac{1}{\Gamma(a)b^a} x^{a-1} e^{-x/b}. \quad (\text{C.24})$$

In this notation, a and b are the shape and scale parameters of the distribution, respectively. Note that for $a = b = 1$, one recovers the exponential prior of Eq. C.9. However, the range of distributions that can be obtained with this form of prior enables one to impose more physical constraints, avoiding the inclusion of zero or infinite values of α_k in the prior, while maintaining analytical feasibility.

Substituting the Gamma distribution above in Eq. C.12 instead of the exponential distribution, gives

$$P(D|\theta, H_2) = \prod_{k=1}^N \frac{1}{(2\pi)^{n_k/2} |V_k|^{1/2} \Gamma(a) b^a} \int_0^\infty \alpha_k^{n_k/2} e^{-\frac{1}{2}\alpha_k \chi_k^2} e^{-\alpha/b} d\alpha \quad (\text{C.25})$$

and after some algebra:

$$\ln P(D|\theta, H_2) = \sum_{k=1}^N \left[\ln \Gamma\left(\frac{n_k}{2} + a\right) - \left(\frac{n_k}{2} + a\right) \ln\left(\frac{\chi_k^2}{2} + 1/b\right) - \ln \Gamma(a) - a \ln b \right]. \quad (\text{C.26})$$

Here, we denoted this new “model” by H_2 . Note that, as expected, this expression reduces to Eq. C.21 in the $a = b = 1$ case, and we dropped the same factors of β_k to enable effective model comparison via numerical computation of the Bayesian evidence.

To choose a prior on α_k , it is desirable to maintain an expectation of 1, $E[\alpha_k] = 1$. This property was indeed obeyed by the exponential prior used in previous sections. For the Gamma distribution, we remark the following properties:

- The mean is $a \times b$;
- For $a > 1$, the mode is $(a - 1) \times b$;
- The variance is $a \times b^2$;
- The skewness is $2/\sqrt{a}$.

Therefore, in order to keep the mean equal to 1, one can set $a = 1/b$. This reduces the number of parameters of our Gamma prior to 1. From the exact expressions above, we see that when $a = 1/b$ we have variance scaling as $1/a$ and skewness scaling as $2/\sqrt{a}$.




Figure C-1: Gamma priors over diagnostic weights for some values of the ν parameter. All of these choices have a distribution mean of 1, indicated by the vertical dotted black line. Quantiles of each distribution are shown in different colors: 25-75 (red), 10-90 (blue), 1-99 (green). $\nu \equiv a = 1/b = 25$ was used in the inferences described Ref. [155], from which this figure is reproduced.

Fig. C-1 shows the resulting distribution for a number of choices of $\nu \equiv a = 1/b$, where the colors show the 1-99 quantile range (green), overlapped with the the 10-90 quantile range (blue) and the 25-75 quantile range (red). We see that this resembles a skewed Gaussian centered at 1; the standard deviation of is 0.1 and the skewness is 0.2. This choice of distribution appears to be an appropriate choice for the prior of α_k . Note that the skewness of the distribution is not a feature that is necessarily problematic, but is inevitably connected with the use of Gamma priors.

Setting $a = 1/b$, one can rewrite Eq. C.26 in terms of ν as

$$\ln P(D|\theta, H_3) = \sum_{k=1}^N \left[\ln \Gamma \left(\frac{n_k}{2} + \nu \right) - \left(\frac{n_k}{2} + \nu \right) \ln \left(\frac{\chi_k^2}{2} + \nu \right) - \ln \Gamma(\nu) + \nu \ln \nu \right]. \quad (\text{C.27})$$

which we denote as yet a new model H_3 , although this is really just a special case of model H_2 .

Finally, differentiating Eq. C.25 with respect to α and setting the result to zero, one

finds that the “effective” weight for each diagnostic can be estimated as

$$\alpha_k^{eff} = \frac{(n_k + 2a - 2)b}{b\chi^2 + 2} \quad (\text{C.28})$$

in analogy with Eq. C.23 for the exponential prior case.

C.3 Summary

We obtained new forms of the log-likelihood for a Bayesian inverse problem. This was derived by considering weight factors for different datasets/instruments, and assigning a prior to all possible values of these factors, α_k . We were able to analytically integrate over the full likelihood function for these new “models” and this led to different dependencies of the likelihood function on the χ^2 for each diagnostic. We initially demonstrated this for an exponential prior of expectation value of 1 (model H_1) and later for a more general Gamma distribution (model H_2), describing a useful special case of practical interest (H_3). We also showed that model comparison between the model H_0 (which does not include diagnostic weights) and H_1 , H_2 and H_3 (which marginalize over weights) is possible by carefully taking normalization factors into consideration. This led to the substitution of Eq. C.21, C.26 or C.27 to Eq. C.22, with **no additional dimensions (free parameters) for our learning process**. Estimated optimal parameters have been shown to provide some intuition for the weights that were effectively inferred for each diagnostic (from Eq. C.23).

This procedure reduces the arbitrary process of specifying weights for each diagnostic solely based on physical knowledge of the diagnostics and separates the inference development in two components: the estimation of experimental values and uncertainties to the best possible degree for the experimentalist, and the statistical estimation of correction factors to account for the possible mismatch of different diagnostic systems studying the same phenomenon. These factors are learnt from the data following reasonable prior knowledge, which may be confined, as in model H_2 to a specific region of the α_k parameter space.

Appendix D

Application of Nested Sampling

As described in Chapter 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. [193, 192] 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` [194], 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 [194], 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 a 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.

Appendix E

Neutral dynamics in Alcator C-Mod

This appendix goes into greater detail on the analysis of D neutrals on Alcator C-Mod, complementing the main results presented in Chapter 6. We expand here on the penetration of neutrals in Alcator C-Mod plasmas, comparing expectations from pure ionization equilibrium and from more detailed SOLPS-ITER simulations. We then complement the description of C-Mod D Ly α signals given in Section 6.2, particularly for the three discharges that have been compared to SOLPS-ITER. Finally, we present recent work on x-ray measurements of the high-n Rydberg series of He-like Ar.

E.1 Penetration of Edge Neutrals

In this section, we compare the degree of atomic D neutral penetration that would be expected purely from ionization equilibrium with the results of the high-fidelity SOLPS-ITER code using the EIRENE Monte Carlo neutral model.

In Fig. E-1, the ratio of the FSA atomic D neutral density to electron density from SOLPS-ITER (red) is compared with the fractional abundance of the neutral stage in a pure plasma at ionization equilibrium (IE, blue). The IE results were computed with AURORA using H effective ionization and recombination rates at local values of n_e and T_e . SOLPS-ITER results only extended from the edge as far as the vertical black dashed lines, inside of which an exponential extrapolation was used all the way to the magnetic axis, setting the IE values as a minimum level of n_n/n_e . The difference between the red and blue lines is due to the existence of multiple generations of CX reactions. When a cold D neutral exchanges an electron with a D ion, the ion becomes neutral itself, but typically has higher




Figure E-1: Fractional abundance of the neutral stage of D, n_n , over the local electron density, n_e , in typical Alcator C-Mod (a) L-mode and (b) H-mode discharges.

energy. Each time a CX reaction occurs, neutrals are free to move across magnetic field lines and the more energy they have, the faster they may penetrate before they get ionized, undergo charge transfer again, or recombine. The blue lines representing IE are only the result of a local calculation that does not consider such multi-generational CX reactions in any way. In other words, the IE ratios correspond to zero-energy neutrals that cannot charge exchange. Clearly, the comparison of Fig. E-1 shows that neutral penetration into the edge can be very significant and the minimum n_n/n_e ratios at IE cannot be expected to be representative of true fractional abundances of the neutral D stage in H-mode (right) except in the edge. However, in the L-mode case (left) neutrals are seen to penetrate much further into the plasma, due to the absence of a density pedestal. It should therefore come to no surprise that DIII-D plasmas, typically having lower density by nearly 1 order of magnitude with respect to C-Mod, exhibit larger neutral penetration and therefore a larger impact of neutrals on plasma dynamics, as shown in Section 6.4.

E.2 Experimental Analysis of D Ly $_{\alpha}$ Signals

In this section, we expand on the data analysis procedures used to process D Ly $_{\alpha}$ data in Chapter 6. As discussed in Section 6.2, the Ly $_{\alpha}$ array of interest has 20 spatial chords at the outer midplane. An Abel inversion is used to compute emissivities, ε , from brightness measurements. Photon emissivity coefficients (PECs) for the excitation- and recombination-driven components of the D Ly $_{\alpha}$ transition have been taken from ADAS [114] using Aurora.

Neutral densities are then computed via

$$n_n = \frac{\varepsilon}{E_{ij}(n_e \mathcal{P}_{Ly_\alpha}^{exc} + n_i \mathcal{P}_{Ly_\alpha}^{rec})}. \quad (\text{E.1})$$

Fig. E-2 shows the experimental data and workflow used to compute neutral densities and D ion source rates from Ly_α emissivities for the three discharges presented in Fig. 6-3, where experimentally-inferred neutral densities were compared to SOLPS-ITER results. The left column in Fig. E-2 shows the L-mode shot #1100308004 ($t = 0.9 - 1.1$ s), the middle column the EDA H-mode shot #1100305023 ($t = 0.85 - 1.3$ s), and the right column the I-mode shot #1080416025 ($t = 0.8 - 1.0$ s). The top panels show Ly_α emissivities, averaged within the chosen (quasi-steady) time windows; the second panels from the top display electron densities (left axis, blue) and temperatures (right axis, red); the third panels show the log-10 of inferred atomic D neutral density; finally, the bottom panels give the corresponding D ion source rate from neutral particles ionizations. Fitted data is displayed as shaded bands whose color intensity is proportional to the Gaussian probability distribution function of the measurement. This visualization is equivalent to plotting multiple standard deviations near the Gaussian mean, but offers the benefit of clearly displaying what is the assigned probability to any value away from the center of the fit. All n_e and T_e data are from Thomson scattering, with all data points within the chosen time windows shown explicitly in the second panels from the top. Points shown in the bottom two panels correspond to processing of the Ly_α emissivity using individual local Thomson scattering measurements, which give both n_e and T_e values at each time and location.

Fig. E-2 shows detailed Thomson scattering coverage over the radial extent where we make use of Ly_α data, extending over the pedestal and near-SOL. While in the EDA H-mode case there exist Thomson scattering data up to much further out, we have not made use of deep-SOL results in this work. Future work will make use of selected discharges where Ly_α and kinetic profile data are available also further out in radius, also relying on reciprocating Langmuir probe data. While probe data were included in numerous shots that are part of the Ly_α database presented in Section 6.3, they were not available for the three discharges shown in Fig. E-2.

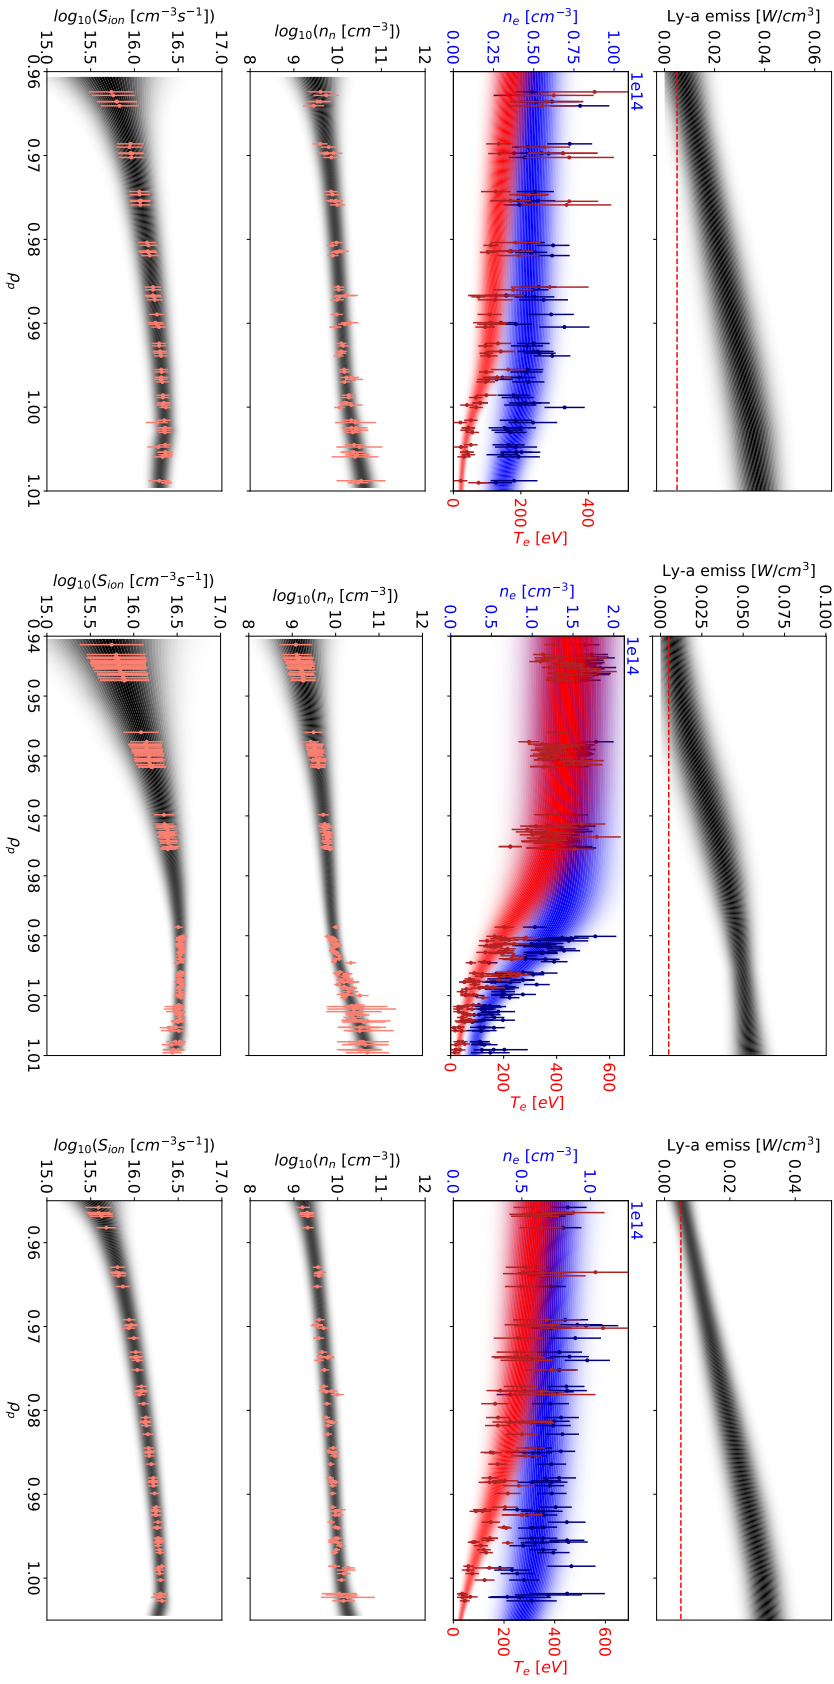


Figure E-2: Ly α analysis of the Alcator C-Mod discharges #1100308004, #1100305023, and #1080416025. The top panel shows local emissivity, the second panel the electron density (left, blue) and temperature (right, red), the third panel the log-10 of inferred atomic D neutral density, and the bottom panel the log-10 of the corresponding D ion source rate from neutral particle ionizations.

E.3 High- n CX X-Ray Measurements

While neutral particles have a strong effect on plasma dynamics and are therefore important to tokamak operations, they also offer opportunities to test atomic theory and obtain indirect measurements on plasma properties. In this section, we briefly describe recent work, presented in detail in Ref. [272], on measurements of very high- n CX transitions in the He-like Ar Rydberg series on C-Mod. These measurements, collected via several x-ray imaging crystal spectrometers (including the XICS system described in Section 4.4) over the years, can differentiate the emission from CX between H-like Ar and the excited neutral states of hydrogen isotopes (hereafter, simply referred as “H” for simplicity). CX tends to occur into shells with principal quantum numbers (n) that scale linearly with the n value of D neutrals excited states. In Ref. [272], the following simple formula is derived

$$n_f = n_H^* \sqrt{Z^3 / (3Z - 2)} \quad (\text{E.2})$$

as an alternative to the widely adopted $Z^{3/4}$ scaling [115]. Based on this, one may estimate that CX of ground state ($n_H^* = 1$) H neutrals with $Z=18$ ions mostly populate the $n = 10-11$ impurity states, the first excited states ($n_H^* = 2$) populate $n = 18 - 20$, and the second excited state ($n_H^* = 3$) populate $n = 27 - 30$. One may expect the population of $n_H^* = 3$ H neutrals in a fusion plasma to be negligibly small, but in Ref. [272] we show experimental observations proving that this is not the case. Such finding can be understood as a result of the scaling of CX cross sections with the approximate size of the recombining ion, which scales as $n_c^4 a_0^2 / Z^2 \sim Z^2 a_0^2 (n_H^*)^4$, where a_0 is the Bohr radius. Due to this strong scaling (to the fourth power of n_H^*), the Rydberg series of He-like Ar (Ar¹⁶⁺) is in practice extended all the way to $n = 27 - 30$, near the ionization limit.

Fig. E-3 compares spectra between $n = 11$ and the ionization limit for plasmas with high (solid red line) and low (dot-dashed green line) D neutral density, normalized to have matching $n = 11$ brightness. While neutral density is not directly measured, it is expected that strongly differing electron density and temperature cause different degrees of neutral penetration from the plasma edge. The low neutral density case in Fig. E-3 corresponds to conditions of $\langle n_e \rangle = 0.8 \times 10^{20} \text{ m}^{-3}$ and $T_e(r = 0) = 3.2 \text{ keV}$, whereas the high neutral density case has $\langle n_e \rangle = 1.4 \times 10^{20} \text{ m}^{-3}$ and $T_e(r = 0) = 1.2 \text{ keV}$. Vertical magenta bands indicate the spectral regions where the brightness is expected to derive from CX populating

$n = 18 - 20$ and $n = 26 - 31$ states. We note that excitation processes are unlikely to populate such high principal quantum numbers, hence the difference between red and green lines can only be interpreted as the result of CX only.




Figure E-3: Spectra between $n = 11$ and the ionization limit (dotted vertical line) for plasmas with $\langle n_e \rangle = 0.8 \times 10^{20} \text{ m}^{-3}$ and $T_e(r = 0) = 3.2 \text{ keV}$ (solid red line) and for plasmas with $\langle n_e \rangle = 1.4 \times 10^{20} \text{ m}^{-3}$ and $T_e(r = 0) = 1.2 \text{ keV}$ (dot-dash green line). Vertical bands indicate the regions for $n = 18 - 20$ (from $n^* = 2$ excited neutrals) and $n = 26 - 31$ (from $n^* = 3$ excited neutrals). Reproduced from Ref. [272].

In the absence of CX, the oscillator strengths (linearly proportional to transitional probabilities, or Einstein A coefficients) of H-like Rydberg series transitions are theoretically expected to scale as $1/n^3$ [112]. The same scaling has also been found for the Ar^{16+} Rydberg series using the Flexible Atomic Code [116] (FAC). Experimental data presented in Ref. [272] from spectrometer chords with tangency radii near the magnetic axis have found a scaling closer to $1/n^{3.25}$. In the edge region, where the neutral density is higher, the emissivity of $n = 9$ and 10 levels is significantly greater than suggested by this scaling, due to the additional CX recombination process. The same is true for transitions near $n = 30$, whereas transitions with $n \approx 20$ often do not appear to have this “excess” brightness. Since population of $n = 20$ states is attributed to $n_H^* = 2$, this suggests that the radial distribution of neutral deuterium may vary for different excited states. Calculation of excited state fractions based on rates from the collisional-radiative model in the DEGAS-2 code [295] suggest that fractional abundances of excited states of deuterium are approximately 0.004, 0.002 and 0.001 for $n_H^* = 2, 3$, and 4, respectively, with only a weak radial dependence across most of

	LSN	USN
∇B drift up	n_n higher bottom n_I higher bottom	n_0 higher top n_I higher bottom
∇B drift down	n_n higher bottom n_I higher top	n_n higher top n_I higher top

Table E.1: Up/down neutral (n_n) and impurity density (n_I) asymmetries relative to X-point location and ion $\mathbf{B} \times \nabla \mathbf{B}$ drift direction.

the plasmas radius. It appears challenging to explain the lack of $n = 20$ brightness in the x-ray measurements based on such radial profiles of D neutral excited states.

Focusing on a spectrometer chord with tangency radius near the magnetic axis, brightness ratios for transitions above $n = 21$ and $n = 7$ have been found to increase with electron temperature, decrease with electron density, and to be independent of ion temperature. Observations have been limited to L- and I-mode discharges, since H-mode cases present much lower neutral penetration (and therefore lower CX rates) due to the presence of a density pedestal (see Fig. E-1).

In Ref. [272], x-ray measurements have also been shown to highlight a strong up/down asymmetry in the emissivity of transitions induced by CX, suggesting a higher neutral density near active x-points. By looking at USN and LSN discharges, with normal and reversed toroidal field direction, it has been possible to exclude the possibility that the observed up/down brightness asymmetries are solely attributable to asymmetries of Ar density. These are known to be due to an unrelated parallel transport effect, leading to an excess of impurity densities in the direction opposite to the ion ∇B drift [98, 107]. On the other hand, the up/down asymmetry of emissivities induced by CX has been found to relate to the proximity to an active X-point. This supports the mental picture described by Table E.1.

While state-of-the-art atomic models such as in the Flexible Atomic Code (FAC) [116] do not seem to be capable yet of handling CX into the very high-n states that were experimentally observed in this work, we have examined whether neutral densities from SOLPS-ITER could appropriately account for the up/down asymmetries of neutral D density that are expected from spectroscopic data. Fig. E-4 shows the neutral (atomic) D density over the poloidal cross section, obtained in an L-mode discharge where particle and heat diffusivities were varied in SOLPS-ITER to match edge Thomson scattering data, as

previously described, following the procedure in Ref. [226]. Due to the large range of neutral densities in this figure, any poloidal asymmetries do not appear with sufficient clarity in this 2D plot. Fig. E-5 displays n_n profiles along the poloidal direction for various radial locations inside of the LCFS. The left-most poloidal angle, θ , in this figure corresponds to the X-point location, whereas other vertical dashed lines show -90° , 0° , $+90^\circ$ and 180° for ease of visualization. Neutral densities are found to vary significantly as a function of θ , but do not display an up/down asymmetry, nor a neutral density excess near the X-point. In fact, neutral densities seem to be greatest at the midplane ($\theta \approx 0^\circ$ and $\theta \approx 180^\circ$). It is yet unclear whether this result is due to choices made in SOLPS-ITER modeling or a more fundamental disagreement. Future research aiming at understanding the up/down asymmetry of neutral density inferred from x-ray measurements should investigate the possible influence of gas puffing, drifts and impurities on these modeling results.




Figure E-4: Distribution of atomic D neutrals across the poloidal cross section in the L-mode shot that was compared to Ar high-n Rydberg series measurements. Reproduced from Ref. [272].




Figure E-5: Poloidal distributions of neutral density at multiple radial locations inside of the separatrix.

Appendix F

Kinetic Neutral 1D Modeling

In this appendix we briefly describe the KN1D code, its potential use within Aurora and the OMFIT module that has been created to facilitate KN1D modelling across devices. KN1D was used in Section 8.3 to model atomic D neutral edge densities in a DIII-D discharge.

The inference of impurity transport coefficients in the DIII-D RMP ELM-suppressed discharge 175674, discussed in Section 8.3, makes use of edge neutral density profiles from the *Kinetic Neutral 1D* (KN1D) code [200]. Originally developed by LaBombard at Alcator C-Mod, KN1D offers a faster, albeit lower-fidelity, model than the 2D neutral Monte Carlo codes like EIRENE, integrated within the SOLPS-ITER code. KN1D computes neutral density distributions for both D₂ and D as a function of radial distance from the wall, in a 1D slab geometry, accounting for ionization, charge exchange (CX) and elastic scattering on both ions and neutrals [159]. Inputs to the code include background profiles of electron density, n_e , electron temperature, T_e , ion temperature, T_i , and a few geometric parameters describing the location of the last closed flux surface (LCFS) and the tokamak limiter. Impurities are not included, and therefore the D ion density is taken to be equal to n_e . An input D₂ pressure at the wall boundary is used to set the amplitude of neutral densities in the simulation. On Alcator C-Mod, this quantity can be taken from the midplane pressure measurement widely used in the Ly $_{\alpha}$ database of Section 6.3, even though in practice this must typically be tuned to match experimental Ly $_{\alpha}$ brightness.

Fig. F-1 shows a comparison of two KN1D simulations for the DIII-D discharge 175674 for two values of the D₂ wall pressure: 0.1 mTorr (blue) and 0.01 mTorr (green). The top two panels show the input n_e and T_e profiles, with T_i taken to be equal to T_e . The neutral

(D) temperature, also shown in the second panel, is an output of KN1D. In the third panel, the D atomic neutral density, n_D is shown together with the D₂ neutral and ion density. The bottom panel finally shows the atomic ionization rate (source of D ions). Note that the radial coordinate is in units of meters, taking the zero at the wall. The comparison of the blue and green cases in this figure make it clear that neutral density profiles are effectively scaled linearly with the D₂ wall pressure value given as an input. This justifies the use of KN1D results in impurity transport inferences that take the absolute magnitude of atomic D neutral density as a free parameter. Of course, constraining the absolute neutral density is no easy task and, in the context of our analysis of discharge 175674, this free parameter is mostly used to assess the degree of uncertainty that derives by our lack of knowledge for this quantity.

KN1D has been coupled to Aurora for convenient inclusion of CX in impurity transport forward modeling. An OMFIT KN1D module has also been developed to enable simple creation of KN1D namelists and visualization of results for multiple devices. It is expected that these capabilities will facilitate comparison of 1D kinetic modeling with the Lyman Alpha Measurement Apparatus (LLAMA) diagnostic recently installed on DIII-D [240]. A wide comparison of KN1D and SOLPS-ITER for both Alcator C-Mod and DIII-D is also in progress to examine the applicability of KN1D on devices where main-chamber recycling is not the dominant source of neutrals as on C-Mod.




Figure F-1: Comparison of KN1D results for DIII-D discharge 175674 using two values of edge D_2 pressure: 0.1 mTorr (blue) and 0.01 mTorr (green).

Bibliography

- [1] P. Dunn. An aggressive market-driven model for US fusion power development. <https://news.mit.edu/2021/aggressive-market-driven-model-us-fusion-power-development-0224>.
- [2] J. P. Freidberg. *Plasma Physics and Fusion Energy*. Cambridge University Press, 2007.
- [3] Committee on the Key Goals and Innovation Needed for a U.S. Fusion Pilot Plant. Bringing fusion to the U.S. grid: a consensus study report of the National Academies of Sciences, Engineering and Medicine. Technical report, February 2021.
- [4] International Atomic Energy Agency. *Fusion Physics*. Non-serial Publications. International Atomic Energy Agency, Vienna, 2012.
- [5] H. S. Bosch and G. M. Hale. Improved formulas for fusion cross-sections and thermal reactivities. *Nuclear Fusion*, 32(4):611–631, 1992.
- [6] J. D. Lawson. Some criteria for a power producing thermonuclear reactor. *Proceedings of the Physical Society. Section B*, 70(1):6–10, 1957.
- [7] J. Wesson. *Tokamaks*. Oxford University Press, 4 edition, 2011.
- [8] R. Lourie. *Sakharov, A Biography*. Brandeis University Press, 2 edition, 2002.
- [9] E. J. Doyle, W. A. Houlberg, Y. Kamada, V. Mukhovatov, T. H. Osborne, A. Polevoi, G. Bateman, J. W. Connor, J. G. Cordey, T. Fujita, X. Garbet, T. S. Hahm, L. D. Horton, A. E. Hubbard, F. Imbeaux, F. Jenko, J. E. Kinsey, Y. Kishimoto, J. Li, T. C. Luce, Y. Martin, M. Ossipenko, V. Parail, A. Peeters, T. L. Rhodes, J. E. Rice, C. M. Roach, V. Rozhansky, F. Ryter, G. Saibene, R. Sartori, A. C.C. Sips, J. A. Snipes, M. Sugihara, E. J. Synakowski, H. Takenaga, T. Takizuka, K. Thomsen, M. R. Wade, and H. R. Wilson. Chapter 2: Plasma confinement and transport. *Nuclear Fusion*, 47(6):18–127, 2007.
- [10] ITER Physics Expert Group on Confinement and Transport. Chapter 2: Plasma confinement and transport Interval estimation of global H-mode energy confinement in ITER. *Nuclear Fusion*, 39(2175), 1999.
- [11] O. Kardaun, K. Thomsen, and A. Chudnovskiy. Corrections to a sequence of papers in Nuclear Fusion. *Nuclear Fusion*, 48(099801), 2008.
- [12] A. J. Creely, M. J. Greenwald, S. B. Ballinger, D. Brunner, J. Canik, J. Doody, T. Fülöp, D. T. Garnier, R. Granetz, T. K. Gray, C. Holland, N. T. Howard, J. W.

- Hughes, J. H. Irby, V. A. Izzo, G. J. Kramer, A. Q. Kuang, B. Labombard, Y. Lin, B. Lipschultz, N. C. Logan, J. D. Lore, E. S. Marmor, K. Montes, R. T. Mumgaard, C. Paz-Soldan, C. Rea, M. L. Reinke, P. Rodriguez-Fernandez, K. Särkimäki, F. Sciortino, S. D. Scott, A. Snicker, P. B. Snyder, B. N. Sorbom, R. Sweeney, R. A. Tinguely, E. A. Tolman, M. Umansky, O. Vallhagen, J. Varje, D. G. Whyte, J. C. Wright, S. J. Wukitch, and J. Zhu. Overview of the SPARC tokamak. *Journal of Plasma Physics*, 86(5), 2020.
- [13] P. Rodriguez-Fernandez, N. T. Howard, M. J. Greenwald, A. J. Creely, J. W. Hughes, J. C. Wright, C. Holland, Y. Lin, and F. Sciortino. Predictions of core plasma performance for the SPARC tokamak. *Journal of Plasma Physics*, 2020.
- [14] General Atomics. DIII-D National Fusion Facility. <https://www.ga.com/magnetic-fusion/diii-d>. Accessed on March 8, 2021.
- [15] T. Pütterich, E. Fable, R. Dux, M. O’Mullane, R. Neu, and M. Siccino. Determination of the tolerable impurity concentrations in a fusion reactor using a consistent set of cooling factors. *Nuclear Fusion*, 59(5):ab0384, 2019.
- [16] T. Eich, A. W. Leonard, R. A. Pitts, W. Fundamenski, R. J. Goldston, T. K. Gray, A. Herrmann, A. Kirk, A. Kallenbach, O. Kardaun, A. S. Kukushkin, B. Labombard, R. Maingi, M. A. Makowski, A. Scarabosio, B. Sieglin, J. Terry, and A. Thornton. Scaling of the tokamak near the scrape-off layer H-mode power width and implications for ITER. *Nuclear Fusion*, 53(9), 2013.
- [17] M. L. Reinke. Heat flux mitigation by impurity seeding in high-field tokamaks. *Nuclear Fusion*, 57(3):034004, 3 2017.
- [18] M. G. Dunne. Impact of impurity seeding and divertor conditions on transitions, pedestal structure and ELMs. *Nuclear Fusion*, 57(2):25002, 2017.
- [19] A. Kallenbach, R. Neu, R. Dux, and and others. Tokamak operation with high-Z plasma facing components Tokamak operation with high- Z plasma facing. *Plasma Physics and Controlled Fusion*, 47:B207–B222, 2005.
- [20] E S Marmor, J L Cecchi, and S A Cohen. System for rapid injection of metal atoms into plasmas. *Review of Scientific Instruments*, 46(9):1149–1154, 1975.
- [21] F. Wagner, G. Becker, K. Behringer, D. Campbell, and and others. Regime of Improved Confinement and High Beta in Neutral-Beam-Heated Divertor Discharges of the ASDEX Tokamak. *Physical Review Letters*, 49(19):1179–1068, 1982.
- [22] M. Greenwald, R. Boivin, P. Bonoli, C. Fiore, J. Goetz, R. Granetz, A. Hubbard, I. Hutchinson, J. Irby, Y. Lin, E. Marmor, A. Mazurenko, D. Mossessian, T. Sunn Pedersen, J. Rice, J. Snipes, G. Schilling, G. Taylor, J. Terry, S. Wolfe, and S. Wukitch. Studies of EDA H-mode in Alcator C-Mod. *Plasma Physics and Controlled Fusion*, 42(SUPPL. 5A):A263–A269, 2000.
- [23] D. G. Whyte, A. E. Hubbard, J. W. Hughes, B. Lipschultz, J. E. Rice, E. S. Marmor, M. Greenwald, I. Cziegler, A. Dominguez, T. Golfopoulos, N. Howard, L. Lin, R. M. McDermottb, M. Porkolab, M. L. Reinke, J. Terry, N. Tsujii, S. Wolfe, S. Wukitch, and Y. Lin. I-mode: An H-mode energy confinement regime with L-mode particle transport in Alcator C-Mod. *Nuclear Fusion*, 50(10):105005, 2010.

- [24] C. M. Greenfield, M. Murakami, J. R. Ferron, M. R. Wade, T. C. Luce, C. C. Petty, J. E. Menard, T. W. Petrie, S. L. Allen, K. H. Burrell, T. A. Casper, J. C. DeBoo, E. J. Doyle, A. M. Garofalo, I. A. Gorelov, R. J. Groebner, J. Hobirk, A. W. Hyatt, R. J. Jayakumar, C. E. Kessel, R. J. La Haye, G. L. Jackson, J. Lohr, M. A. Makowski, R. I. Pinsky, P. A. Politzer, R. Prater, E. J. Strait, T. S. Taylor, and W. P. West. High performance advanced tokamak regimes in DIII-D for next-step experiments. *Physics of Plasmas*, 11(5 PART 2):2616–2626, 2004.
- [25] C. C. Petty, J. E. Kinsey, C. T. Holcomb, J. C. Deboo, E. J. Doyle, J. R. Ferron, A. M. Garofalo, A. W. Hyatt, G. L. Jackson, T. C. Luce, M. Murakami, P. A. Politzer, and H. Reimerdes. High-beta, steady-state hybrid scenario on DIII-D. *Nuclear Fusion*, 56(1), 2016.
- [26] N. T. Howard. *Experimental and Gyrokinetic Studies of Impurity Transport in the Core of Alcator C-Mod Plasmas*. PhD thesis, 2012.
- [27] M. A. Chilenski. *Experimental Data Analysis Techniques for Validation of Tokamak Impurity Transport Simulations*. PhD thesis, MIT, 2017.
- [28] M. Kardar. *Statistical Physics of Particles*. Cambridge University Press, 2007.
- [29] P. Helander and D. J. Sigmar. *Collisional transport in magnetized plasmas*. Cambridge University Press, 2010.
- [30] S. I. Braginskii. Transport Processes in a Plasma. *Reviews of Plasma Physics*, 1:205, 1965.
- [31] R. D. Hazeltine. Recursive derivation of drift-kinetic equation. *Plasma Physics*, 15(1):77–80, 1973.
- [32] E. A. Belli and J. Candy. Kinetic calculation of neoclassical transport including self-consistent electron and impurity dynamics. *Plasma Physics and Controlled Fusion*, 50(9):0–32, 2008.
- [33] T. Pütterich, R. Dux, M. A. Janzer, and R. M. McDermott. ELM flushing and impurity transport in the H-mode edge barrier in ASDEX Upgrade. In *Journal of Nuclear Materials*, volume 415, pages 334–339, 2011.
- [34] E. A. Belli and J. Candy. Full linearized Fokker-Planck collisions in neoclassical transport simulations. *Plasma Physics and Controlled Fusion*, 54(1):015015, 1 2012.
- [35] P. J. Catto. Linearized gyro-kinetics. *Plasma Physics*, 20(719), 1978.
- [36] G. G. Howes, S. C. Cowley, W. Dorland, G. W. Hammett, E. Quataert, and A. A. Schekochihin. Astrophysical Gyrokinetics: Basic Equations and Linear Theory. *The Astrophysical Journal*, 651(1):590–614, 2006.
- [37] J. Candy, E. A. Belli, and R. V. Bravenec. A high-accuracy Eulerian gyrokinetic solver for collisional plasmas. *Journal of Computational Physics*, 324:73–93, 2016.
- [38] E. A. Belli and J. Candy. Impact of centrifugal drifts on ion turbulent transport. *Phys. Plasmas*, 032301(December 2017), 2018.

- [39] E. A. Belli and J. Candy. Implications of advanced collision operators for gyrokinetic simulation. *Plasma Physics and Controlled Fusion*, 59(4), 2017.
- [40] G. W. Hammett and F. W. Perkins. Fluid moment models for Landau damping with application to the ion-temperature-gradient instability. *Physical Review Letters*, 64(25):3019–3022, 1990.
- [41] G. W. Hammett, W. Dorland, and F. W. Perkins. Fluid models of phase mixing, landau damping, and nonlinear gyrokinetic dynamics. *Physics of Fluids B*, 4(7):2052–2061, 1992.
- [42] G M Staebler, J E Kinsey, and R E Waltz. Gyro-Landau fluid equations for trapped and passing particles. *Physics of Plasmas*, 12(10):1–24, 2005.
- [43] P. M. Bellan. *Fundamentals of Plasma Physics*. Cambridge University Press, 2006.
- [44] T. Neiser, O. Meneghini, S. Smith, M. Fasciana, G.M. Staebler, and J. Candy. Big Data Validation of the TGLF Transport Model. In *APS Division of Plasma Physics Meeting Abstracts*, volume 2020 of *APS Meeting Abstracts*, page ZO04.007, January 2020.
- [45] G. M. Staebler, N. T. Howard, J. Candy, and C. Holland. A model of the saturation of coupled electron and ion scale gyrokinetic turbulence. *Nuclear Fusion*, 57(6):066046, 2017.
- [46] N. T. Howard, C. Holland, A. E. White, M. Greenwald, J. Candy, and A. J. Creely. Multi-scale gyrokinetic simulations: Comparison with experiment and implications for predicting turbulence and transport. *Physics of Plasmas*, 23(5):056109, 2016.
- [47] A. M. Dimits, G. Bateman, M. A. Beer, B. I. Cohen, W. Dorland, G. W. Hammett, C. Kim, J. E. Kinsey, M. Kotschenreuther, A. H. Kritz, L. L. Lao, J. Mandrekas, W. M. Nevins, S. E. Parker, A. J. Redd, D. E. Shumaker, R. Sydora, and J. Weiland. Comparisons and physics basis of tokamak transport models and turbulence simulations, 3 2000.
- [48] P. H. Diamond, S. I. Itoh, K. Itoh, and T. S. Hahm. Zonal flows in plasma - A review. *Plasma Physics and Controlled Fusion*, 47(5), 2005.
- [49] G. M. Staebler, J. Candy, N. T. Howard, and C. Holland. The role of zonal flows in the saturation of multi-scale gyrokinetic turbulence. *Physics of Plasmas*, 23(6), 2016.
- [50] A. J. Creely, P. Rodriguez-Fernandez, G. D. Conway, S. J. Freethy, N. T. Howard, and A. E. White. Criteria for the importance of multi-scale interactions in turbulent transport simulations. *Plasma Physics and Controlled Fusion*, 61(8):1–6, 2019.
- [51] P. Rodriguez-Fernandez, A. E. White, N. T. Howard, B. A. Grierson, G. M. Staebler, J. E. Rice, X. Yuan, N. M. Cao, A. J. Creely, M. J. Greenwald, A. E. Hubbard, J. W. Hughes, J. H. Irby, and F. Sciortino. Explaining Cold-Pulse Dynamics in Tokamak Plasmas Using Local Turbulent Transport Models. *Physical Review Letters*, 120(7):075001, 2 2018.

- [52] J. Candy, C. Holland, R. E. Waltz, M. R. Fahey, and E. Belli. Tokamak profile prediction using direct gyrokinetic and neoclassical simulation. *Physics of Plasmas*, 16(6), 2009.
- [53] G. De Dominicis, G Fuhr, C Bourdelle, P Beyer, X Garbet, Y. Sarazin, and G. L. Falchetto. Quasilinear nature of L-mode edge turbulent transport in fluid nonlinear simulations. *Nuclear Fusion*, 59(12):126019, 2019.
- [54] P. H. Diamond and K. Itoh. *Modern Plasma Physics Vol . 1 Physical Kinetics of Turbulent Plasmas*, volume 1.
- [55] R.J. Goldston and P.H. Rutherford. *Introduction to Plasma Physics*. Institute of Physics Pub., 1st editio edition, 1995.
- [56] W. Horton, D. I. In Choi, and W. M. Tang. Toroidal drift modes driven by ion pressure gradients. *Physics of Fluids*, 24(6):1077–1085, 6 1981.
- [57] S. Brunner. Waves and instabilities in inhomogenous plasmas. EPFL Doctoral School Lectures, 2014.
- [58] S. Brunner. *Global approach to the spectral problem of microinstabilities in a cylindrical plasma using a gyrokinetic model*. PhD thesis, EPFL, 1997.
- [59] M. Kotschenreuther, X. Liu, D. R. Hatch, S. Mahajan, L. Zheng, A. Diallo, R. Groebner, J. C. Hillesheim, C. F. Maggi, C. Giroud, F. Koechl, V. Parail, S. Saarelma, E. Solano, and A. Chankin. Gyrokinetic analysis and simulation of pedestals to identify the culprits for energy losses using 'fingerprints'. *Nuclear Fusion*, 59(9), 2019.
- [60] J. P. Freidberg. *Ideal MHD*. Cambridge University Press, 2014.
- [61] R. J. La Haye. Neoclassical tearing modes and their control. In *Physics of Plasmas*, volume 13, pages 1–18, 2006.
- [62] M. E. Foord and E. S. Marmor. Sawtooth oscillations in the visible continuum on alcator c. *Nuclear Fusion*, 25(2):197–202, 1985.
- [63] F. H. Seguin, R. Petrasso, and E. S. Marmor. Effects of Internal Disruptions on Impurity Transport in Tokamaks. *Physical Review Letters*, 51(6):455–458, 8 1983.
- [64] B. B. Kadomtsev and B. B. Disruptive instability in Tokamaks. *Fizika Plazmy*, 1:710–715, 1975.
- [65] J. A. Wesson. Sawtooth reconnection. *Nuclear Fusion*, 30(12):2545–2549, 1990.
- [66] F. Porcelli, D. Boucher, and M. N. Rosenbluth. Model for the sawtooth period and amplitude. *Plasma Physics and Controlled Fusion*, 38(12):2163–2186, 1996.
- [67] P. B. Snyder, H. R. Wilson, J. R. Ferron, L. L. Lao, A. W. Leonard, T. H. Osborne, A. D. Turnbull, D. Mossessian, M. Murakami, and X. Q. Xu. Edge localized modes and the pedestal: A model based on coupled peeling-ballooning modes. In *Physics of Plasmas*, volume 9, pages 2037–2043, 2002.

- [68] P. B. Snyder, J. W. Hughes, T. H. Osborne, C. Paz-Soldan, W. M. Solomon, M. Knolker, D. Eldon, T. Evans, T. Golfopoulos, B. A. Grierson, R. J. Groebner, A. E. Hubbard, E. Kolemen, B. Labombard, F. M. Laggner, O. Meneghini, S. Mordijck, T. Petrie, S. Scott, H. Q. Wang, H. R. Wilson, and Y. B. Zhu. High fusion performance in Super H-mode experiments on Alcator C-Mod and DIII-D. *Nuclear Fusion*, 59(8):86017–86027, 2019.
- [69] P. C. Stangeby. *The Plasma Boundary of Magnetic Fusion Devices*. CRC Press, 1 edition, 2000.
- [70] H. J. Sun, E. Wolfrum, B. Kurzan, T. Eich, K. Lackner, A. Scarabosio, I. P. Pérez, O Kardaun, M Faitsch, S Potzel, and U Stroth. Study of near SOL decay lengths in ASDEX Upgrade under attached and detached divertor conditions. *Plasma Physics and Controlled Fusion*, 59(10), 2017.
- [71] B. LaBombard, E. Marmor, J. Irby, J. L. Terry, R. Vieira, G. Wallace, D. G. Whyte, S. Wolfe, S. Wukitch, S. Baek, W. Beck, P. Bonoli, D. Brunner, J. Doody, R. Ellis, D. Ernst, C. Fiore, J. P. Freidberg, T. Golfopoulos, R. Granetz, M. Greenwald, Z. S. Hartwig, A. Hubbard, J. W. Hughes, I. H. Hutchinson, C. Kessel, M. Kotschenreuther, R. Leccacorvi, Y. Lin, B. Lipschultz, S. Mahajan, J. Minervini, R. Mumgaard, R. Nygren, R. Parker, F. Poli, M. Porkolab, M. L. Reinke, J. Rice, T. Rognlien, W. Rowan, S. Shiraiwa, D. Terry, C. Theiler, P. Titus, M. Umansky, P. Valanju, J. Walk, A. White, J. R. Wilson, G. Wright, and S. J. Zweben. ADX: A high field, high power density, advanced divertor and RF tokamak. *Nuclear Fusion*, 55(5):053020, 2015.
- [72] D. Brunner, B. Labombard, A. Q. Kuang, and J. L. Terry. High-resolution heat flux width measurements at reactor-level magnetic fields and observation of a unified width scaling across confinement regimes in the Alcator C-Mod tokamak. *Nuclear Fusion*, 58(9):094002, 7 2018.
- [73] C. Holland. Validation metrics for turbulent plasma transport. *Physics of Plasmas*, 23(6):060901, 2016.
- [74] M. Greenwald. Verification and validation for magnetic fusion. *Physics of Plasmas*, 17(5):058101, 2010.
- [75] C. Angioni, A. G. Peeters, G. V. Pereverzev, F. Ryter, and G. Tardini. Density Peaking, Anomalous Pinch, and Collisionality in Tokamak Plasmas. *Physical Review Letters*, 90(20):4, 5 2003.
- [76] C. Angioni, E. Fable, M. Greenwald, M. Maslov, A. G. Peeters, H. Takenaga, and H. Weisen. Particle transport in tokamak plasmas, theory and experiment. *Plasma Physics and Controlled Fusion*, 51(12):124017, 2009.
- [77] M. Greenwald, C. Angioni, J. W. Hughes, J. Terry, and H. Weisen. Density profile peaking in low collisionality H-modes: Comparison of Alcator C-Mod data to ASDEX Upgrade/JET scalings. *Nuclear Fusion*, 47(9):L26–L29, 2007.
- [78] C. Angioni, P. Mantica, T. Pütterich, M. Valisa, M. Baruzzo, E. A. Belli, P. Belo, F. J. Casson, C. Challis, P. Drewelow, C. Giroud, N. Hawkes, T. C. Hender, J. Hobirk,

- T. Koskela, L. Lauro Taroni, C. F. Maggi, J. Mlynar, T. Odstrcil, M. L. Reinke, and M. Romanelli. Tungsten transport in JET H-mode plasmas in hybrid scenario, experimental observations and modelling. *Nuclear Fusion*, 54(8):083028, 2014.
- [79] T. Tala, H. Nordman, A. Salmi, C. Bourdelle, J. Citrin, A. Czarnecka, F. Eriksson, E. Fransson, C. Giroud, J. Hillesheim, C. Maggi, P. Mantica, A. Mariani, M. Maslov, L. Meneses, S. Menmuir, S. Mordijck, V. Naulin, M. Oberparleiter, G. Sips, D. Tegnered, M. Tsalas, and H. Weisen. Density peaking in JET—determined by fuelling or transport? *Nuclear Fusion*, 59(12):126030, 2019.
- [80] S. Mordijck, T. L. Rhodes, L. Zeng, A. Salmi, T. Tala, C. C. Petty, G. R. McKee, R. Reksoatmodjo, F. Eriksson, E. Fransson, and H. Nordman. Collisionality driven turbulent particle transport changes in DIII-D H-mode plasmas. *Nuclear Fusion*, 60(6), 2020.
- [81] M. Greenwald, D. Whyte, P. Bonoli, Z. Hartwig, J. Irby, B. Labombard, E. Marmor, J. Minervini, M. Takayasu, J. Terry, R. Vieira, A. White, S. Wukitch, D. Brunner, R. Mungaard, and B. Sorbom. The High-Field Path to Practical Fusion Energy. Technical Report February, 2018.
- [82] H. Sugama and W. Horton. Nonlinear electromagnetic gyrokinetic equation for plasmas with large mean flows. *Physics of Plasmas*, 5(7):2560–2573, 7 1998.
- [83] M. R. Wade, D. L. Hillis, J. T. Hogan, R. Maingi, M. M. Menon, M. A. Mahdavi, W. P. West, K. H. Burrell, R. Gohil, R. J. Groebner, R. M. Hong, D. H. Kellman, J. C. Phillips, R. R. Seraydarian, and D. F. Finkenthal. Helium transport and exhaust studies in enhanced confinement regimes in DIII-D. *Physics of Plasmas*, 2(6):2357–2365, 1995.
- [84] C. Bruhn, R. M. McDermott, C. Angioni, J. Ameres, V. Bobkov, M. Cavedon, R. Dux, A. Kappatou, A. Lebschy, P. Manas, and R. Ochoukov. 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.
- [85] R. Dux. Impurity transport in ASDEX upgrade. *Fusion Science and Technology*, 44(3):708–715, 2003.
- [86] R. Dux, R. Neu, A. G. Peeters, G. Pereverzev, A. Mück, F. Ryter, and J. Stober. Influence of the heating profile on impurity transport in ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 45(9):1815–1825, 2003.
- [87] B. A. Grierson, K. H. Burrell, R. M. Nazikian, W. M. Solomon, A. M. Garofalo, E. A. Belli, G. M. Staebler, M. E. Fenstermacher, G. R. McKee, T. E. Evans, D. M. Orlov, S. P. Smith, C. Chrobak, and C. Chrystal. Impurity confinement and transport in high confinement regimes without edge localized modes on DIII-D. *Physics of Plasmas*, 22(5):055901, 2015.
- [88] T. Odstrčil, N. T. Howard, F. Sciortino, C. Chrystal, C. Holland, E. Hollmann, G. McKee, K. E. Thome, and T. M. Wilks. Dependence of the impurity transport on the dominant turbulent regime in ELM-y H-mode discharges on the DIII-D tokamak. *Physics of Plasmas*, 27(8):082503, 8 2020.

- [89] R. Dux, C. Giroud, and K. D. Zastrow. Impurity transport in internal transport barrier discharges on JET. *Nuclear Fusion*, 44(2):260–264, 2004.
- [90] C. Giroud, R. Barnsley, P. Buratti, I. H. Coffey, M. Von Hellermann, C. Jupén, K. D. Lawson, A. Meigs, M. O’Mullane, A. D. Whiteford, and K. D. Zastrow. Method for experimental determination of Z dependence of impurity transport on JET. *Nuclear Fusion*, 47(4):313–330, 4 2007.
- [91] T. Parisot, R. Guirlet, C. Bourdelle, X. Garbet, N. Dubuit, F. Imbeaux, and P. R. Thomas. Experimental impurity transport and theoretical interpretation in a Tore Supra lower-hybrid heated plasma. *Plasma Physics and Controlled Fusion*, 50(5):1–18, 2008.
- [92] D. Villegas, R. Guirlet, C. Bourdelle, X. Garbet, G.T. Hoang, R. Sabot, F. Imbeaux, and J.L. Segui. Experimental and theoretical study of nickel transport dependence on gradients in Tore Supra. *Nuclear Fusion*, 54:073011, 2014.
- [93] E. Scavino, J. S. Bakos, R. Dux, and H. Weisen. Effects of plasma shape on laser blow-off injected impurity transport in TCV. *Plasma Physics and Controlled Fusion*, 45(11):1961–1974, 2003.
- [94] Z. Y. Cui, K. Zhang, S. Morita, X. Q. Ji, X. T. Ding, Y. Xu, P. Sun, J. M. Gao, C. F. Dong, D. L. Zheng, Y. G. Li, M. Jiang, D. Li, W. L. Zhong, Y. Liu, Y. B. Dong, S. D. Song, L. M. Yu, Z. B. Shi, B. Z. Fu, P. Lu, M. Huang, B. S. Yuan, Q. W. Yang, and X. R. Duan. Study of impurity transport in HL-2A ECRH L-mode plasmas with radially different ECRH power depositions. *Nuclear Fusion*, 58(5), 2018.
- [95] B. Geiger, T. Wegner, C. D. Beidler, R. Burhenn, B. Buttenschön, R. Dux, A. Langenberg, N. A. Pablant, T. Pütterich, Y. Turkin, T. Windisch, V. Winters, M. Beurskens, C. Biedermann, K. J. Brunner, G. Cseh, H. Damm, F. Effenberg, G. Fuchert, O. Grulke, J. H. Harris, C. Killer, J. Knauer, G. Kocsis, A. Krämer-Flecken, T. Kremeyer, M. Krychowiak, O. Marchuk, D. Nicolai, K. Rahbarnia, G. Satheeswaran, J. Schilling, O. Schmitz, T. Schröder, T. Szepesi, H. Thomsen, H. Trimino Mora, P. Traverso, and D. Zhang. Observation of anomalous impurity transport during low-density experiments in W7-X with laser blow-off injections of iron. *Nuclear Fusion*, 59(4):046009, 2019.
- [96] S. S. Henderson, L. Garzotti, F. J. Casson, D. Dickinson, M. O’Mullane, A. Patel, C. M. Roach, H. P. Summers, H. Tanabe, and M. Valovič. Charge dependence of neoclassical and turbulent transport of light impurities on MAST. *Plasma Physics and Controlled Fusion*, 57(9):95001, 2015.
- [97] L. Delgado-Aparicio, D. Stutman, K. Tritz, M. Finkenthal, S. Kaye, R. Bell, R. Kaita, B. Leblanc, F. Levinton, J. Menard, S. Paul, D. Smith, and H. Yuh. Impurity transport studies in NSTX neutral beam heated H-mode plasmas. *Nuclear Fusion*, 49(8), 2009.
- [98] J. E. Rice, J. L. Terry, E. S. Marmor, and F Bombarda. X ray observations of up-down impurity density asymmetries in alcator C-Mod plasmas. *Nuclear Fusion*, 37(2):241–249, 2 1997.

- [99] J. E. Rice, J. A. Goetz, R. S. Granetz, M. J. Greenwald, A. E. Hubbard, I. H. Hutchinson, E. S. Marmor, D. Mossessian, T. Sunm Pedersen, J. A. Snipes, J. L. Terry, and S. M. Wolfe. Impurity toroidal rotation and transport in Alcator C-Mod ohmic high confinement mode plasmas. *Physics of Plasmas*, 7(5):1825–1830, 5 2000.
- [100] T. S. Pedersen, R. S. Granetz, A. E. Hubbard, I. H. Hutchinson, E. S. Marmor, J. E. Rice, and J. Terry. Radial impurity transport in the H mode transport barrier region in Alcator C-Mod. *Nuclear Fusion*, 40(10):1795–1804, 2000.
- [101] J. E. Rice, J. L. Terry, E. S. Marmor, R. S. Granetz, M. J. Greenwald, A. E. Hubbard, J. H. Irby, S. M. Wolfe, and T. Sunm Pedersen. Impurity transport in alcator C-Mod plasmas. *Fusion Science and Technology*, 51(3):357–368, 2007.
- [102] N. T. Howard, M. Greenwald, D. R. Mikkelsen, M. L. Reinke, A. E. White, D. Ernst, Y. Podpaly, and J. Candy. Quantitative comparison of experimental impurity transport with nonlinear gyrokinetic simulation in an Alcator C-Mod L-mode plasma. *Nuclear Fusion*, 52(6), 2012.
- [103] M. A. Chilenski, M. J. Greenwald, Y. M. Marzouk, J. E. Rice, and A. E. White. Efficient design and verification of diagnostics for impurity transport experiments. *Review of Scientific Instruments*, 89(1):013504, 2018.
- [104] M. L. Reinke, I. H. Hutchinson, J. E. Rice, N. T. Howard, A. Bader, S. Wukitch, Y. Lin, D. C. Pace, A. Hubbard, J. W. Hughes, and Y. Podpaly. Poloidal variation of high-Z impurity density due to hydrogen minority ion cyclotron resonance heating on Alcator C-Mod. *Plasma Physics and Controlled Fusion*, 54(4):045004, 2012.
- [105] T. Odstrčil, T. Pütterich, C. Angioni, R. Bilato, A. Gude, and M. Odstrcil. The physics of W transport illuminated by recent progress in W density diagnostics at ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 60(1):014003, 1 2018.
- [106] R. M. Churchill, C. Theiler, B. Lipschultz, I.H. Hutchinson, M. L. Reinke, D. Whyte, J. W. Hughes, P. Catto, M. Landreman, D. Ernst, C. S. Chang, R. Hager, A. Hubbard, P. Ennever, and J. R. Walk. Poloidal asymmetries in edge transport barriers. *Physics of Plasmas*, 22(5):056104, 2015.
- [107] J. E. Rice, M. L. Reinke, N. Cao, J. W. Hughes, J. M.A. Ashbourn, D. R. Ernst, A. E. Hubbard, and J. H. Irby. Up/down impurity density asymmetries in C-Mod plasmas. *Nuclear Fusion*, 58(12):126008, 9 2018.
- [108] F. J. Casson, C. Angioni, E. A. Belli, R. Bilato, P. Mantica, T. Odstrčil, T. Pütterich, M. Valisa, L. Garzotti, C. Giroud, J. Hobirk, C. F. Maggi, J. Mlynar, and M. L. Reinke. Theoretical description of heavy impurity transport and its application to the modelling of tungsten in JET and ASDEX upgrade. *Plasma Physics and Controlled Fusion*, 57(1):14031, 2015.
- [109] C. Angioni. Impurity transport in tokamak plasmas, theory, modelling and comparison with experiments. *Plasma Physics and Controlled Fusion*, 63(7):073001, 2021.
- [110] C. Angioni, M. Sertoli, R. Bilato, V. Bobkov, A. Loarte, R. Ochoukov, T. Odstrčil, T. Pütterich, and J. Stober. A comparison of the impact of central ECRH and central ICRH on the tungsten behaviour in ASDEX Upgrade H-mode plasmas. *Nuclear Fusion*, 57(5):aa6453, 2017.

- [111] F. J. Casson, A. G. Peeters, C. Angioni, Y. Camenen, W. A. Hornsby, A. P. Snodin, and G. Szepesi. Gyrokinetic simulations including the centrifugal force in a rotating tokamak plasma. *Physics of Plasmas*, 17(10), 2010.
- [112] T. Fujimoto. *Plasma Spectroscopy*. Oxford Science Publications, International Series of Monographs on Physics, 1 edition, 2004.
- [113] D. R. Bates, A. E. Kingston, and R. W. P. McWhirter. Recombination between electrons and atomic ions I . Optically thin plasmas. *Proceedings of the Royal Society of London. Series A, Mathematical and physical sciences*, 267(1330), 1962.
- [114] H. P. Summers, W. J. Dickson, M. G. O’Mullane, N. R. Badnell, A. D. Whiteford, D. H. Brooks, J. Lang, S. D. Loch, and D. C. Griffin. Ionization state, excited populations and emission of impurities in dynamic finite density plasmas: I. The generalized collisional-radiative model for light elements. *Plasma Physics and Controlled Fusion*, 48(2):263–293, 2006.
- [115] I. H. Hutchinson. *Principles of Plasma Diagnostics*. Cambridge University Press, 2 edition, 2002.
- [116] M. F. Gu. The flexible atomic code. *Canadian Journal of Physics*, 86(5):675–689, 2008.
- [117] T. Schlummer. *Charge exchange recombination in X-ray spectra of He-like argon measured at the tokamak TEXTOR*. PhD thesis, Heinrich-Heine-Universität Düsseldorf, 2014.
- [118] R. C. Isler. Review Paper: Impurities in tokamaks. *Nuclear Fusion*, 24(12):1599–1678, 1984.
- [119] W. W. Heidbrink, D. Liu, Y. Luo, E. Ruskov, and B. Geiger. A code that simulates fast-ion $D\alpha$ and neutral particle measurements. *Communications in Computational Physics*, 10(3):716–741, 2011.
- [120] C. A. Johnson, S. D. Loch, and D. A. Ennis. ColRadPy: A Python collisional radiative solver. *Nuclear Materials and Energy*, 20:100579, 8 2019.
- [121] W. J. Karzas and R. Latter. Electron Radiative Transitions in a Coulomb Field. *The Astrophysical Journal Supplement Series*, 6:167, 5 1961.
- [122] B. A. Trubnikov. Yield Coefficient of Cyclotron Radiation from a Thermonuclear Plasma. *JETP Letters*, 16:25–27, 1972.
- [123] H. Zohm. On the Use of High Magnetic Field in Reactor Grade Tokamaks. *Journal of Fusion Energy*, 38(1):3–10, 2019.
- [124] The Nobel Prize in Physics 1915. NobelPrize.org. Nobel Media AB 2021.
- [125] S. D. Loch. *Spectral series emission and atomic populations in solar astrophysical and laboratory fusion plasmas*. PhD thesis, 2001.
- [126] A. H. Gabriel and C. Jordan. Interpretation of Solar Helium-Like Ion Line Intensities. *Monthly Notices of the Royal Astronomical Society*, 145(2):241–248, 7 1969.

- [127] L. A. Vainshtein and U. I. Safronova. Wavelengths and transition probabilities of satellites to resonance lines of H- and He-like ions, 1978.
- [128] F. Bely-Dubau, J. Dubau, P. Faucher, A. H. Gabriel, M. Loulergue, L. Steenman-Clark, S. Volonte, E. Antonucci, and C. G. Rapley. Dielectronic satellite spectra for highly-charged helium-like ions - VII: Ca spectra. *Monthly Notices of the Royal Astronomical Society*, 201:1155–1169, 1982.
- [129] A. H. Gabriel. Dielectronic Satellite Spectra for Highly-Charged Helium-Like Ion Lines. *Monthly Notices of the Royal Astronomical Society*, 160(1):99–119, 1972.
- [130] D. Porquet, J. Dubau, and N. Grosso. He-like ions as practical astrophysical plasma diagnostics: From stellar coronae to active galactic nuclei. *Space Science Reviews*, 157(1-4):103–134, 2010.
- [131] A. R. Foster and K. Heuer. PyAtomDB: Extending the AtomDB atomic database to model new plasma processes and uncertainties. *Atoms*, 8(3):49, 9 2020.
- [132] G. Y. Liang and N. R. Badnell. R-matrix electron-impact excitation data for the Li-like iso-electronic sequence including Auger and radiation damping. *Astronomy and Astrophysics*, 528:A69, 4 2011.
- [133] N. R. Badnell. Radiative Recombination Data for Modeling Dynamic Finite-Density Plasmas. *The Astrophysical Journal Supplement Series*, 167(2):334–342, 2006.
- [134] M. A. Bautista and N. R. Badnell. Dielectronic recombination data for dynamic finite-density plasmas XII. The helium isoelectronic sequence. *Astronomy and Astrophysics*, 466(2):755–762, 5 2007.
- [135] P. Palmeri, P. Quinet, C. Mendoza, M. A. Bautista, J. García, and T. R. Kallman. Radiative and Auger Decay of K-Vacancy Levels in the Ne, Mg, Si, S, Ar, and Ca Isonuclear Sequences. *The Astrophysical Journal Supplement Series*, 177(1):408–416, 7 2008.
- [136] A. Kramida and Y. Ralchenko and J. Reader and NIST ASD Team. NIST Atomic Spectra Database (ver. 5.8), [Online]. Available: <https://physics.nist.gov/asd> [2021, May 26]. National Institute of Standards and Technology, Gaithersburg, MD., 2020.
- [137] S. Li, J. Yan, C. Y. Li, R. Si, X. L. Guo, M. Huang, C. Y. Chen, and Y. M. Zou. Radiative rates and electron-impact excitation for the $n \leq 6$ fine-structure levels in H-like ions with $13 \leq Z \leq 42$. *Astronomy and Astrophysics*, 583:A82, 11 2015.
- [138] A. R Foster. private communication.
- [139] R. Mewe. Interpolation Formulae for the Electron Impact Excitation of Ions in the H-, He-, Li- and Ne-Sequences. *Astronomy and Astrophysics*, 20:215–221, 1972.
- [140] R. Mewe and J. Schrijver. Helium-like Ion Line Intensities. *Astronomy and Astrophysics*, 65:99–114, 1978.
- [141] D. Porquet and J. Dubau. X-ray photoionized plasma diagnostics with helium-like ions. Application to warm absorber-emitter in active galactic nuclei. *Astronomy and Astrophysics Supplement Series*, 143(3):495–514, 5 2000.

- [142] D. Porquet, R. Mewe, J. Dubau, A. J.J. Raassen, and J. S. Kaastra. Line ratios for helium-like ions: Applications to collision-dominated plasmas. *Astronomy and Astrophysics*, 376(3):1113–1122, 2001.
- [143] E. Källne, J. Källne, E. S. Marmar, and J. E. Rice. High resolution x-ray spectroscopy diagnostics of high temperature plasmas. *Physica Scripta*, 31(6):551–564, 1985.
- [144] J. E. Rice, M. A. Graf, J. L. Terry, E. S. Marmar, K. Giesing, and F. Bombarda. X-ray observations of helium-like scandium from the alcator c-mod tokamak. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 28(5):893–905, 1995.
- [145] M. Bitter, M. F. Gu, L. A. Vainshtein, P. Beiersdorfer, G. Bertschinger, O. Marchuk, R. Bell, B. LeBlanc, K. W. Hill, D. Johnson, and L. Roquemore. New benchmarks from tokamak experiments for theoretical calculations of the dielectronic satellite spectra of heliumlike ions. *Physical Review Letters*, 91(26), 2003.
- [146] O. Marchuk, M. Z. Tokar, G. Bertschinger, A. Urnov, H. J. Kunze, D. Pilipenko, X. Loozen, D. Kalupin, D. Reiter, A. Pospieszczyk, W. Biel, M. Goto, and F. Goryaev. Comparison of impurity transport model with measurements of He-like spectra of argon at the tokamak TEXTOR. *Plasma Physics and Controlled Fusion*, 48(11):1633–1646, 11 2006.
- [147] J. E. Rice, M. L. Reinke, J. M.A. Ashbourn, C. Gao, M. M. Victora, M. A. Chilenski, L. Delgado-Aparicio, N. T. Howard, A. E. Hubbard, J. W. Hughes, and J. H. Irby. X-ray observations of Ca19 +, Ca18 + and satellites from Alcator C-Mod tokamak plasmas. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 47(7):075701, 2014.
- [148] J. E. Rice, M. L. Reinke, J. M.A. Ashbourn, C. Gao, M. Bitter, L. Delgado-Aparicio, K. Hill, N. T. Howard, J. W. Hughes, and U. I. Safronova. 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*, 48(14):1–8, 2015.
- [149] J. E. Rice, E. S. Marmar, E. Källne, and J. Källne. Radial profiles of ground-state transitions of heliumlike argon from the Alcator-C tokamak. *Physical Review A*, 35(7):3033–3045, 4 1987.
- [150] A. S. Rosen, M. L. Reinke, J. E. Rice, A. E. Hubbard, and J. W. Hughes. Validation of x-ray line ratios for electron temperature determination in tokamak plasmas. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 47(10), 2014.
- [151] T. Kato, K. Masai, T. Fujimoto, F. Koike, E. Källne, E. S. Marmar, and J. E. Rice. Effects of state-selective charge-exchange processes on the He-like spectra from the Alcator C tokamak. *Physical Review A*, 44(10):6776–6784, 1991.
- [152] J. E. Rice, E. S. Marmar, J. L. Terry, E. Kallne, and J. Kallne. Observation of charge-transfer population of high-n levels in Ar+16 from neutral hydrogen in the ground and excited states in a tokamak plasma. *Physical Review Letters*, 56(1):50–53, 1 1986.
- [153] N. T. Howard, M. Greenwald, and J. E. Rice. Characterization of impurity confinement on Alcator C-Mod using a multi-pulse laser blow-off system. *Review of Scientific Instruments*, 82(3):0–5, 2011.

- [154] M. A. Chilenski, M. Greenwald, Y. Marzouk, N. T. Howard, A. E. White, J. E. Rice, and J. R. Walk. Improved profile fitting and quantification of uncertainty in experimental measurements of impurity transport coefficients using Gaussian process regression. *Nuclear Fusion*, 55(2):23012, 2015.
- [155] F. Sciortino, N. T. Howard, E. S. Marmor, T. Odstroil, N. M. Cao, R. Dux, A. E. Hubbard, J. W. Hughes, J. H. Irby, Y. M. Marzouk, L. M. Milanese, M. L. Reinke, J. E. Rice, and P. Rodriguez-Fernandez. Inference of experimental radial impurity transport on Alcator C-Mod: Bayesian parameter estimation and model selection. *Nuclear Fusion*, 60(12), 8 2020.
- [156] F. Sciortino, N. M. Cao, N. T. Howard, E. S. Marmor, and J. E. Rice. Particle transport constraints via Bayesian spectral fitting of multiple atomic lines. *Review of Scientific Instruments*, 92:53508, 2021.
- [157] J. W. Hughes, D. Mossessia, K. Zhurovich, M. DeMaria, K. Jensen, and A. Hubbard. Thomson scattering upgrades on Alcator C-Mod. In *Review of Scientific Instruments*, volume 74, pages 1667–1670, 3 2003.
- [158] N. P. Basse, A. Dominguez, E. M. Edlund, C. L. Fiore, R. S. Granetz, A. E. Hubbard, J. W. Hughes, I. H. Hutchinson, J. H. Irby, B. LaBombard, L. Lin, Y. Lin, B. Lipschultz, J. E. Liptac, E. S. Marmor, D. A. Mossessian, R. R. Parker, M. Porkolab, J. E. Rice, J. A. Snipes, V. Tang, J. L. Terry, S. M. Wolfe, S. J. Wukitch, K. Zhurovich, R. V. Bravenec, P. E. Phillips, W. L. Rowan, G. J. Kramer, G. Schilling, S. D. Scott, and S. J. Zweben. Diagnostic systems on Alcator C-Mod. *Fusion Science and Technology*, 51(3):476–507, 2007.
- [159] J. W. Hughes, B. LaBombard, D. A. Mossessian, A. E. Hubbard, J. Terry, and T. Biewer. Advances in measurement and modeling of the high-confinement-mode pedestal on the Alcator C-Mod tokamak. In *Physics of Plasmas*, volume 13, page 056103. American Institute of PhysicsAIP, 5 2006.
- [160] T. N. Carlstrom, G. L. Campbell, J. C. DeBoo, R. Evanko, J. Evans, C. M. Greenfield, J. Haskovec, C. L. Hsieh, E. McKee, R. T. Snider, R. Stockdale, P. K. Trost, and M. P. Thomas. Design and operation of the multipulse Thomson scattering diagnostic on DIII-D (invited). *Review of Scientific Instruments*, 63(10):4901–4906, 10 1992.
- [161] D. M. Ponce-Marquez, B. D. Bray, T. M. Deterly, C. Liu, and D. Eldon. Thomson scattering diagnostic upgrade on DIII-D. *Review of Scientific Instruments*, 81(19D525), 2010.
- [162] M. E. Austin and J. Lohr. Electron cyclotron emission radiometer upgrade on the DIII-D tokamak. In *Review of Scientific Instruments*, volume 74, pages 1457–1459. American Institute of PhysicsAIP, 3 2003.
- [163] C. M. Bishop. *Pattern Recognition and Machine Learning*. 2006.
- [164] M. L. Reinke, P. Beiersdorfer, N. T. Howard, E. W. Magee, Y. Podpaly, J. E. Rice, and J. L. Terry. Vacuum ultraviolet impurity spectroscopy on the Alcator C-Mod tokamak. In *Review of Scientific Instruments*, volume 81, page 10D736, 2010.
- [165] J. K. Lepson, P. Beiersdorfer, R. Kaita, R. Majeski, and D. Boyle. Responsivity calibration of the LoWEUS spectrometer. In *Review of Scientific Instruments*, volume 87, page 11D614. AIP Publishing LLC, 11 2016.

- [166] L. L. Lao, H. St John, R. D. Stambaugh, A. G. Kellman, and W. Pfeiffer. Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nuclear Fusion*, 25(11):1611–1622, 1985.
- [167] R. J. Fonck, A. T. Ramsey, and R. V. Yelle. Multichannel grazing-incidence spectrometer for plasma impurity diagnosis: SPRED. *Applied Optics*, 21(12):2115, 1982.
- [168] R. D. Wood and S. L. Allen. Absolute calibration of a SPRED spectrograph for use on the DIII-D tokamak. In *Topical conference on high temperature plasmas*, 1988.
- [169] K. P. Dere, G. Del Zanna, P. R. Young, E. Landi, and R. S. Sutherland. CHIANTI—An Atomic Database for Emission Lines. XV. Version 9, Improvements for the X-Ray Satellite Lines. *The Astrophysical Journal Supplement Series*, 241(2):22, 2019.
- [170] E. Träbert, P. Beiersdorfer, J. K. Lepson, M. L. Reinke, and J. E. Rice. EUV Spectra and Line Ratios of Multiply Ionized Calcium and Argon Atoms in a Laboratory Plasma. *The Astrophysical Journal*, 865(2):148, 2018.
- [171] C. Chrystal, K. H. Burrell, B. A. Grierson, S. R. Haskey, R. J. Groebner, D. H. Kaplan, and A. Briesemeister. Improved edge charge exchange recombination spectroscopy in DIII-D. In *Review of Scientific Instruments*, volume 87, page 11E512. American Institute of Physics Inc., 11 2016.
- [172] C. Chrystal, A. A. Sulyman, S. R. Haskey, K. H. Burrell, and B. A. Grierson. Synthetic diagnostic for assessing spatial averaging of charge exchange recombination spectroscopy measurements. *Review of Scientific Instruments*, 89(10):10–101, 2018.
- [173] R. M. McDermott, R. Dux, T. Pütterich, B. Geiger, A. Kappatou, A. Lebschy, C. Bruhn, M. Cavedon, A. Frank, N. D. Harder, and E. Viezzer. Evaluation of impurity densities from charge exchange recombination spectroscopy measurements at ASDEX Upgrade. *Plasma Physics and Controlled Fusion*, 60(9):95007, 2018.
- [174] D. F. Finkenthal. The measurement of absolute helium ion density profiles on the DIII-D tokamak using charge exchange recombination spectroscopy, 1994.
- [175] M. L. Reinke, Y. A. Podpaly, M. Bitter, I. H. Hutchinson, J. E. Rice, L. Delgado-Aparicio, and C. Gao. X-ray imaging crystal spectroscopy for use in plasma transport research. *Review of Scientific Instruments*, 83(113504), 2012.
- [176] B. A. Grierson, K. H. Burrell, C. Chrystal, R. J. Groebner, D. H. Kaplan, W. W. Heidbrink, J. M. Muñoz Burgos, N. A. Pablant, W. M. Solomon, and M. A. Van Zeeland. Active spectroscopic measurements of the bulk deuterium properties in the DIII-D tokamak. In *Review of Scientific Instruments*, volume 83, page 10D529. American Institute of Physics, 10 2012.
- [177] S. R. Haskey, B. A. Grierson, L. Stagner, K. H. Burrell, C. Chrystal, R. J. Groebner, A. Ashourvan, and N. A. Pablant. Deuterium charge exchange recombination spectroscopy from the top of the pedestal to the scrape off layer in H-mode plasmas. *Journal of Instrumentation*, 12(10), 2017.

- [178] S R Haskey, B A Grierson, L Stagner, C Chrystal, A Ashourvan, A Bortolon, M D Boyer, K H Burrell, C Collins, R J Groebner, D H Kaplan, and N A Pablant. Active spectroscopy measurements of the deuterium temperature, rotation, and density from the core to scrape off layer on the DIII-D tokamak, 2018.
- [179] S. R. Haskey, B. A. Grierson, C. Chrystal, A. Ashourvan, K. H. Burrell, R. J. Groebner, E. A. Belli, L. Stagner, D. J. Battaglia, T. Stoltzfus-Dueck, and A. Bortolon. Main ion and impurity edge profile evolution across the L- to H-mode transition on DIII-D. *Plasma Physics and Controlled Fusion*, 60(10):105001, 2018.
- [180] M L Reinke, Y Podpaly, C Gao, and Plasma Science. Operation and Validation of The HIREXSR Analysis COde (THACO). Technical report, 2013.
- [181] Thomas Sunn Pedersen. *Edge Plasma Phenomena in the Alcator C-Mod Tokamak measured by High Resolution X-Ray Imaging Diagnostics*. PhD thesis, 2000.
- [182] M. Odstrčil, J Mlynar, T Odstrcil, B Alper, and A Murari. Modern numerical methods for plasma tomography optimisation. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 686:156–161, 2012.
- [183] Jeff Byers. The physics of data. *Nature Physics*, 13(8):718–719, 2017.
- [184] M. Chilenski, M. J. Greenwald, Y. Marzouk, J. E. Rice, and A. E. White. On the importance of model selection when inferring impurity transport coefficient profiles. *Plasma Physics and Controlled Fusion*, 61(12):125012, 2019.
- [185] R. E. Morrison, T. A. Oliver, and R. D. Moser. Representing model inadequacy: A stochastic operator approach. *SIAM-ASA Journal on Uncertainty Quantification*, 6(2):457–496, 4 2018.
- [186] F. Sciortino, T. Odstrčil, A. Cavallaro, S. Smith, O. Meneghini, R. Reksoatmodjo, O. Linder, J. D. Lore, N. T. Howard, E. S. Marmor, and S. Mordijck. Modeling of Particle Transport, Neutrals and Radiation in Magnetically-Confined Plasmas with Aurora. 2021.
- [187] U. Von Toussaint. Bayesian inference in physics. *Reviews of Modern Physics*, 83(3):943–999, 9 2011.
- [188] R. Trotta. Bayes in the sky: Bayesian inference and model selection in cosmology. *Contemporary Physics*, 49(2):71–104, 2008.
- [189] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee : The MCMC Hammer. *Publications of the Astronomical Society of the Pacific*, 125(925):306–312, 2013.
- [190] R. Andrae, T. Schulze-Hartung, and P. Melchior. Dos and don'ts of reduced chi-squared. *arXiv preprint*, 12 2010.
- [191] J. Skilling. Nested Sampling. *AIP Conference Proceedings*, 395(2004):395–405, 2004.
- [192] F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettitt. Importance Nested Sampling and the MultiNest Algorithm. *arXiv*, pages 1–28, 2013.

- [193] F. Feroz and M. P. Hobson. 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*, 384(2):449–463, 2008.
- [194] W. J. Handley, M. P. Hobson, and A. N. Lasenby. POLYCHORD: Nested sampling for cosmology. *Monthly Notices of the Royal Astronomical Society: Letters*, 450(1):L61–L65, 2015.
- [195] R. Dux. STRAHL User Manual. Technical report, 2006.
- [196] Russell A Hulse. Numerical studies of impurities in fusion plasmas. *Nuclear Technology-Fusion*, 3(2):259–272, 1983.
- [197] L. Lauro-Taroni and and others. Impurity transport of high performance discharges in JET. *Proc. 21st EPS Conf. Controlled Fusion Plasma Physics (Montpellier, France, 27 June–1 July 1994)*, 18B:102, 1994.
- [198] O. Meneghini, S. P. Smith, L. L. Lao, O. Izacard, Q. Ren, J. M. Park, J. Candy, Z. Wang, C. J. Luna, V. A. Izzo, B. A. Grierson, P. B. Snyder, C. Holland, J. Penna, G. Lu, P. Raum, A. McCubbin, D. M. Orlov, E. A. Belli, N. M. Ferraro, R. Prater, T. H. Osborne, A. D. Turnbull, and G. M. Staebler. Integrated modeling applications for tokamak experiments with OMFIT. *Nuclear Fusion*, 55(8):083008, 8 2015.
- [199] J. A. Stillerman, T. W. Fredian, K. A. Klare, and G. Manduchi. MDSplus data acquisition system. *Review of Scientific Instruments*, 68(1):939–942, 6 1997.
- [200] B LaBombard. *KN1D : A 1-D Space , 2-D Velocity , Kinetic Transport Algorithm for Atomic and Molecular Hydrogen in an Ionizing Plasma*. 2001.
- [201] B. Geiger, L. Stagner, W. W. Heidbrink, R. Dux, R. Fischer, Y. Fujiwara, A. Garcia, A. S. Jacobsen, A. J. van Vuuren, A. N. Karpushov, D. Liu, P. A. Schneider, I. Sfiligoi, P. Zsolt Poloskei, and M. Weiland. Progress in modelling fast-ion d-alpha spectra and neutral particle analyzer fluxes using fidasim. *Plasma Physics and Controlled Fusion*, 2020.
- [202] L. Stagner, B. Geiger, and W. W. Heidbrink. FIDASIM: A Neutral Beam and Fast-ion Diagnostic Modeling Suite.
- [203] D. Reiter, M. Baelmans, and P. Börner. The eirene and B2-eirene codes. *Fusion Science and Technology*, 47(2):172–186, 2005.
- [204] S. Wiesen, D. Reiter, V. Kotov, M. Baelmans, W. Dekeyser, A. S. Kukushkin, S. W. Lisgo, R. A. Pitts, V. Rozhansky, G. Saibene, I. Veselova, and S. Voskoboynikov. The new SOLPS-ITER code package. *Journal of Nuclear Materials*, 463:480–484, 7 2015.
- [205] X. Bonnin, W. Dekeyser, R. Pitts, D. Coster, S. Voskoboynikov, and S. Wiesen. Presentation of the New SOLPS-ITER Code Package for Tokamak Plasma Edge Modelling. 11:1403102, 2016.
- [206] O. Linder, E. Fable, F. Jenko, G. Papp, and G. Pautasso. Self-consistent modeling of runaway electron generation in massive gas injection scenarios in ASDEX Upgrade. *Nuclear Fusion*, 60(9), 2020.

- [207] K. Lackner, K. Behringer, W. Engelhardt, and R. Wunderlich. An Algorithm for the Solution of Impurity Diffusion under Finite Reaction Rates. *Zeitschrift fur Naturforschung - Section A Journal of Physical Sciences*, 37(8):931–938, 1982.
- [208] A. Foster. *On the Behaviour and Radiating Properties of Heavy Elements in Fusion Plasmas*. PhD thesis, University of Strathclyde, 2008.
- [209] H. Summers, A. R. Foster, S. D. Loch, M. O’Mullane, and A. Whiteford. Publ3: Heavy species in fusion plasma modelling and spectral analysis. Technical report, May 2010. ADAS-EU R(10)PU03.
- [210] X. Bonnin and D. Coster. Full-tungsten plasma edge simulations with SOLPS. In *Journal of Nuclear Materials*, volume 415, pages S488–S491. Elsevier B.V., 8 2011.
- [211] J. Rice. personal communication.
- [212] H. Jeffreys. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A, Mathematical and physical sciences*, 186(1007):453–461, 9 1946.
- [213] E.T. Jaynes. Prior Probabilities. *IEEE Transactions on Systems Science and Cybernetics*, 4(3):227–241, 1968.
- [214] Roger B. Nelsen. Introduction. In *An Introduction to Copulas*, pages 1–5. Springer New York, 2nd editio edition, 6 2007.
- [215] C. Angioni, R. M. McDermott, E. Fable, R. Fischer, T. Pütterich, F. Ryter, and G. Tardini. Gyrokinetic modelling of electron and boron density profiles of H-mode plasmas in ASDEX Upgrade. *Nuclear Fusion*, 51(2):023006, 2011.
- [216] P. Manas, Y. Camenen, S. Benkadda, H. Weisen, C. Angioni, F. J. Casson, C. Giroud, M. Gelfusa, and M. Maslov. Gyrokinetic modeling of impurity peaking in JET H-mode plasmas. *Physics of Plasmas*, 24(6):062511, 2017.
- [217] W. J. Handley, M. P. Hobson, and A. N. Lasenby. POLYCHORD: Next-generation nested sampling. *Monthly Notices of the Royal Astronomical Society*, 453(4):4384–4398, 2015.
- [218] O. Lahav, S. L. Bridle, M. P. Hobson, A. N. Lasenby, and L. Sodr . Bayesian ‘hyper-parameters’ approach to joint estimation: The Hubble constant from CMB measurements. *Monthly Notices of the Royal Astronomical Society*, 315(4):45–49, 2000.
- [219] M. P. Hobson, S. L. Bridle, and O. Lahav. Combining cosmological data sets: Hyperparameters and Bayesian evidence. *Monthly Notices of the Royal Astronomical Society*, 335(2):377–388, 2002.
- [220] P. Marshall, N. Rajguru, and A. Slosar. Bayesian evidence as a tool for comparing datasets. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 73(6):1–4, 2006.
- [221] J. W. Hughes, D. A. Mossessian, A. E. Hubbard, E. S. Marmar, D. Johnson, and D. Simon. High-resolution edge Thomson scattering measurements on the Alcator C-Mod tokamak. *Review of Scientific Instruments*, 72(1 II):1107–1110, 1 2001.

- [222] J. W. Hughes, P. B. Snyder, J. R. Walk, E. M. Davis, A. Diallo, B. Labombard, S. G. Baek, R. M. Churchill, M. Greenwald, R. J. Groebner, A. E. Hubbard, B. Lipschultz, E. S. Marmor, T. Osborne, M. L. Reinke, J. E. Rice, C. Theiler, J. Terry, A. E. White, D. G. Whyte, S. Wolfe, and X. Q. Xu. Pedestal structure and stability in H-mode and I-mode: A comparative study on Alcator C-Mod. *Nuclear Fusion*, 53(4), 2013.
- [223] R. J. Groebner and T. M. Carlstrom. Critical edge parameters for H-mode transition in DIII-D. *Plasma Physics and Controlled Fusion*, 40(673), 1998.
- [224] J. W. Hughes, A. E. Hubbard, D. A. Mossessian, B. LaBombard, T. M. Biewer, R. S. Granetz, M. Greenwald, I. H. Hutchinson, J. H. Irby, Y. Lin, E. S. Marmor, M. Porkolab, J. E. Rice, J. A. Snipes, J. L. Terry, S. Wolfe, and K. Zhurovich. H-mode pedestal and L-H transition studies on Alcator C-Mod. *Fusion Science and Technology*, 51(3):317–341, 2007.
- [225] P. T. Bonoli, R. Parker, S. J. Wukitch, Y. Lin, M. Porkolab, J. C. Wright, E. Edlund, T. Graves, L. Lin, J. Liptac, A. Parisot, A. E. Schmidt, V. Tang, W. Beck, R. Childs, M. Grimes, D. Gwinn, D. Johnson, J. Irby, A. Kanojia, P. Koert, S. Marazita, E. Marmor, D. Terry, R. Vieira, G. Wallace, J. Zaks, S. Bernabei, C. Brunkhorse, R. Ellis, E. Fredd, N. Greenough, J. Hosea, C. C. Kung, G. D. Loesser, J. Rushinski, G. Schilling, C. K. Phillips, J. R. Wilson, R. W. Harvey, C. L. Fiore, R. Granetz, M. Greenwald, A. E. Hubbard, I. H. Hutchinson, B. LaBombard, B. Lipschultz, J. Rice, J. A. Snipes, J. Terry, and S. M. Wolfe. Wave-particle studies in the ion cyclotron and lower hybrid ranges of frequencies in alcator C-Mod. *Fusion Science and Technology*, 51(3):401–436, 2007.
- [226] R. Reksoatmodjo, S. Mordijck, J. W. Hughes, J. D. Lore, and X. Bonnin. The role of edge fueling in determining the pedestal density in high neutral opacity Alcator C-Mod experiments. *Nuclear Materials and Energy*, 27:100971, 6 2021.
- [227] R. K. Janev and J. J. Smith. Cross Sections for Collisions Processes of Hydrogen Atoms with Electrons, Protons and Multiply Charged Ions. *Nuclear Fusion*, 4:195, 1993.
- [228] R. Dux, M. Cavedon, A. Kallenbach, R. M. McDermott, G. Vogel, and The Asdex Upgrade Team. Influence of CX-reactions on the radiation in the pedestal region at ASDEX Upgrade. *Nuclear Fusion*, 60(12), 2020.
- [229] B. Lipschultz, B. LaBombard, J. L. Terry, C. Boswell, and I. H. Hutchinson. Divertor physics research on Alcator C-Mod. *Fusion Science and Technology*, 51(3):369–389, 2006.
- [230] A. Niemczewski, I. H. Hutchinson, B. LaBombard, B. Lipschultz, and G. M. McCracken. Neutral particle dynamics in the alcator C-MOD tokamak. *Nuclear Fusion*, 37(2):151–163, 1997.
- [231] B. LaBombard, M V Umansky, R L Boivin, J A Goetz, J Hughes, B Lipschultz, D Mossessian, C S Pitcher, and J L Terry. Cross-field plasma transport and main-chamber recycling in diverted plasmas on Alcator C-Mod. *Nuclear Fusion*, 40(12):2041–2060, 2000.

- [232] P. Virtanen, R. Gommers, T. E. Oliphant, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [233] Y. Yang. Can the strengths of AIC and BIC be shared? A conflict between model identification and regression estimation. *Biometrika*, 92(4):937–950, 12 2005.
- [234] M. R.E. Symonds and A. Moussalli. A brief guide to model selection, multimodel inference and model averaging in behavioural ecology using Akaike’s information criterion. *Behavioral Ecology and Sociobiology*, 65(1):13–21, 2011.
- [235] D. Silvagni, T. Eich, T. Happel, G.F. Harrer, M. Griener, M. Dunne, M. Cavedon, M. Faitsch, L. Gil, D. Nille, B. Tal, R. Fischer, U. Stroth, D. Brida, P. David, P. Manz, and E. Viezzer. I-mode pedestal relaxation events at ASDEX Upgrade. *Nuclear Fusion*, 60(12):126028, 12 2020.
- [236] A. Kallenbach, H. J. Sun, T. Eich, D. Carralero, J. Hobirk, A. Scarabosio, and M. Siccino. Parameter dependences of the separatrix density in nitrogen seeded ASDEX Upgrade H-mode discharges. *Plasma Physics and Controlled Fusion*, 60(4), 2018.
- [237] B. Labombard, J. W. Hughes, N. Smick, A. Graf, K. Marr, R. McDermott, M. Reinke, M. Greenwald, B. Lipschultz, J. L. Terry, D. G. Whyte, and S. J. Zweben. Critical gradients and plasma flows in the edge plasma of Alcator C-Mod. *Physics of Plasmas*, 15(5):56106, 2008.
- [238] B. Labombard, J. E. Rice, A. E. Hubbard, J. W. Hughes, M. Greenwald, J. Irby, Y. Lin, B. Lipschultz, E. S. Marmor, C. S. Pitcher, N. Smick, S. M. Wolfe, and S. J. Wukitch. Transport-driven Scrape-off-layer flows and the boundary conditions imposed at the magnetic separatrix in a tokamak plasma. *Nuclear Fusion*, 44(10):1047–1066, 2004.
- [239] B. LaBombard, J. W. Hughes, D. Mossessian, M. Greenwald, B. Lipschultz, and J. L. Terry. Evidence for electromagnetic fluid drift turbulence controlling the edge plasma state in the Alcator C-Mod tokamak. *Nuclear Fusion*, 45(12):1658–1675, 2005.
- [240] A. M. Rosenthal, J. W. Hughes, A. Bortolon, F. M. Laggner, T. M. Wilks, R. Vieira, R. Leccacorvi, E. Marmor, A. Nagy, C. Freeman, and D. Mauzey. A 1D Lyman-alpha profile camera for plasma edge neutral studies on the DIII-D tokamak. *Review of Scientific Instruments*, 92(3):33523, 2021.
- [241] W. W. Heidbrink. Fast-ion $D\alpha$ measurements of the fast-ion distribution. In *Review of Scientific Instruments*, volume 81, page 10D727. American Institute of PhysicsAIP, 10 2010.
- [242] B. Geiger, L. Stagner, W.W. Heidbrink, R. Dux, R. Fischer, Y. Fujiwara, and And Others. Progress in modelling fast-ion D-alpha spectra and neutral particle analyzer fluxes using FIDASIM. *Plasma Physics and Controlled Fusion*, 2019.
- [243] B. Lipschultz, B. LaBombard, C. S. Pitcher, and R. Boivin. Investigation of the origin of neutrals in the main chamber of Alcator C-Mod. In *Plasma Physics and Controlled Fusion*, volume 44, pages 733–748, 2002.

- [244] C. Angioni, H. Weisen, O. J.W.F. Kardaun, M. Maslov, A. Zabolotsky, C. Fuchs, L. Garzotti, C. Giroud, B. Kurzan, P. Mantica, A. G. Peeters, and J. Stober. Scaling of density peaking in H-mode plasmas based on a combined database of AUG and JET observations. *Nuclear Fusion*, 47(9):1326–1335, 2007.
- [245] J. E. Rice, J. L. Terry, J. A. Goetz, Y. Wang, E. S. Marmor, M. Greenwald, I. Hutchinson, Y. Takase, S. Wolfe, H. Ohkawa, and A. Hubbard. Impurity transport in Alcator C-Mod plasmas. In *Physics of Plasmas*, volume 4, pages 1605–1609. American Institute of Physics AIP, 6 1997.
- [246] M. A. Chilenski, M. J. Greenwald, A. E. Hubbard, J. W. Hughes, J. P. Lee, Y. M. Marzouk, J. E. Rice, and A. E. White. Experimentally testing the dependence of momentum transport on second derivatives using Gaussian process regression. *Nuclear Fusion*, 57(12):aa8387, 2017.
- [247] J. E. Rice, C. Gao, M. L. Reinke, P. H. Diamond, N. T. Howard, H. J. Sun, I. Cziegler, A. E. Hubbard, Y. A. Podpaly, W. L. Rowan, J. L. Terry, M. A. Chilenski, L. Delgado-Aparicio, P. C. Ennever, D. Ernst, M. J. Greenwald, J. W. Hughes, Y. Ma, E. S. Marmor, M. Porkolab, A. E. White, and S. M. Wolfe. Non-local heat transport, rotation reversals and up/down impurity density asymmetries in Alcator C-Mod ohmic L-mode plasmas. *Nuclear Fusion*, 53(3):033004, 2013.
- [248] N. T. Howard, A. E. White, M. Greenwald, C. Holland, J. Candy, and J. E. Rice. Impurity transport, turbulence transitions and intrinsic rotation in Alcator C-Mod plasmas. *Plasma Physics and Controlled Fusion*, 56(12):124004, 12 2014.
- [249] O. Sauter, C. Angioni, and Y. R. Lin-Liu. Erratum: Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime (Physics of Plasmas (1999) 6 (2834)). *Physics of Plasmas*, 9(12):5140, 12 2002.
- [250] C. L. Fiore, D. R. Ernst, J. E. Rice, K. Zhurovich, N. Basse, P. T. Bonoli, M. J. Greenwald, E. S. Marmor, and S. J. Wukitch. Internal transport barriers in Alcator C-Mod. *Fusion Science and Technology*, 51(3):303–316, 2006.
- [251] C. K. I. Williams C. E. Rasmussen. *Gaussian processes for machine learning*. Cambridge, Massachusetts, 2006.
- [252] J. E. Rice, M. L. Reinke, C. Gao, N. T. Howard, M. A. Chilenski, L. Delgado-Aparicio, R. S. Granetz, M. J. Greenwald, A. E. Hubbard, J. W. Hughes, J. H. Irby, Y. Lin, E. S. Marmor, R. T. Mumgaard, S. D. Scott, J. L. Terry, J. R. Walk, A. E. White, D. G. Whyte, S. M. Wolfe, and S. J. Wukitch. Core impurity transport in Alcator C-Mod L-, I- and H-mode plasmas. *Nuclear Fusion*, 55(3):033014, 2015.
- [253] F. Sciortino, N. T. Howard, R. Reksoatmodjo, A. R. Foster, J. W. Hughes, E. S. Marmor, S. Mordijck, T. Odstrčil, T. Pütterich, M. L. Reinke, J. E. Rice, and P. Rodriguez-Fernandez. Experimental Inference of Neutral and Impurity Transport in Alcator C-Mod Using High-Resolution X-Ray and Ultra-Violet Spectra. *in preparation*.
- [254] N. M. Cao and F. Sciortino. Bayesian Spectral Moment Estimation and Uncertainty Quantification. *IEEE Transactions on Plasma Science*, pages 1–9, 2019.

- [255] G M Staebler, J E Kinsey, and R E Waltz. A theory-based transport model with comprehensive physics. In *Physics of Plasmas*, volume 14, page 055909, 2007.
- [256] J. Breslau, M. Gorelenkova, F. Poli, J. Sachdev, X. Yuan, and USDOE Office of Science. TRANSP, 2018.
- [257] C. R. Harris, K. J. Millman, S. J. van der Walt, et al. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020.
- [258] G. Verdoolaege, S. M. Kaye, C. Angioni, O. J.W.F. Kardaun, M. Maslov, M. Romanelli, F. Ryter, K. Thomsen, The Asdex Upgrade Team, The Eurofusion Mst1 Team, and J. E.T. Contributors. The updated ITPA global H-mode confinement database: Description and analysis. *Nuclear Fusion*, 61(7), 2021.
- [259] C. Angioni, Y. Camenen, F. J. Casson, E. Fable, R. M. McDermott, A. G. Peeters, and J. E. Rice. Off-diagonal particle and toroidal momentum transport: A survey of experimental, theoretical and modelling aspects, 2012.
- [260] A. E. White, N. T. Howard, A. J. Creely, M. A. Chilenski, M. Greenwald, A. E. Hubbard, J. W. Hughes, E. Marmor, J E Rice, J M Sierchio, C Sung, J R Walk, D G Whyte, D R Mikkelsen, E M Edlund, C Kung, C Holland, J Candy, C C Petty, M L Reinke, and C Theiler. Nonlinear gyrokinetic simulations of the I-mode high confinement regime and comparisons with experiment. *Physics of Plasmas*, 22(5):056109, 2015.
- [261] X. R. Fu, W. Horton, I. O. Bespamyatnov, W. L. Rowan, S. Benkadda, C. L. Fiore, S. Futatani, and K. T. Liao. Turbulent impurity transport modeling for Alcator C-Mod. *Journal of Plasma Physics*, 79(5):837–846, 2013.
- [262] C. Sung, A. E. White, D. R. Mikkelsen, M. Greenwald, C. Holland, N. T. Howard, R. Churchill, and C. Theiler. Quantitative comparison of electron temperature fluctuations to nonlinear gyrokinetic simulations in C-Mod Ohmic L-mode discharges. *Physics of Plasmas*, 23(4), 2016.
- [263] P. Rodriguez-Fernandez, A. E. White, A. J. Creely, M. J. Greenwald, N. T. Howard, F. Sciortino, and J. C. Wright. VITALS: A surrogate-based optimization framework for the accelerated validation of plasma transport codes. *Fusion Science and Technology*, 74(1-2):65–76, 8 2018.
- [264] H. Sugama, T. H. Watanabe, and M. Nunami. Linearized model collision operators for multiple ion species plasmas and gyrokinetic entropy balance equations. *Physics of Plasmas*, 16(11):112503, 2009.
- [265] R. J. Hawryluk. An empirical approach to tokamak transport. In *Physics of Plasmas Close to Thermonuclear Conditions*, pages 19–46. Pergamon, 1 1981.
- [266] D. R. Mikkelsen, N. T. Howard, A. E. White, and A. J. Creely. Verification of GENE and GYRO with L-mode and I-mode plasmas in Alcator C-Mod. *Physics of Plasmas*, 25(4):042505, 4 2018.
- [267] A. J. Creely, N. T. Howard, P. Rodriguez-Fernandez, N. Cao, A. E. Hubbard, J. W. Hughes, J. E. Rice, A. E. White, J. Candy, G. M. Staebler, G. D. Conway, S. J. Freethy,

- and C. Sung. Validation of nonlinear gyrokinetic simulations of L- and I-mode plasmas on Alcator C-Mod. *Physics of Plasmas*, 24(5):056104, 2017.
- [268] G. J. Wilkie, A. Iantchenko, I. G. Abel, E. Highcock, and I. Pusztai. First principles of modelling the stabilization of microturbulence by fast ions. *Nuclear Fusion*, 58(8), 2018.
- [269] F. Eriksson, M. Oberparleiter, A. Skyman, H. Nordman, P. Strand, A. Salmi, and T. Tala. Impact of fast ions on density peaking in JET: Fluid and gyrokinetic modeling. *Plasma Physics and Controlled Fusion*, 61(7):075008, 5 2019.
- [270] B. S. Victor, T. Odstreil, C. A. Paz-Soldan, B. A. Grierson, E. Hinson, A. Jarvinen, E. M. Hollmann, C. Chrystal, C. S. Collins, K. E. Thome, and S. L. Allen. Impurity transport in the pedestal of H-mode plasmas with resonant magnetic perturbations. *Plasma Physics and Controlled Fusion*, 62(9):095021, 9 2020.
- [271] M. Newville, R. Otten, A. Nelson, et al. lmfit/lmfit-py 1.0.2, February 2021.
- [272] J.E. Rice, F. Sciortino, M. Gu, N.M. Cao, J.W. Hughes, J.H. Irby, E.S. Marmor, S. Mordijck, M.L. Reinke, and R. Reksoatmodjo. The very high-n Rydberg series of Ar16+ in Alcator C-Mod tokamak plasmas. *submitted to Journal of Physics B: Atomic, Molecular and Optical Physics*, 2021.
- [273] L. Casali, D. Eldon, J. A. Boedo, T. Leonard, and B. Covele. Neutral leakage, power dissipation and pedestal fueling in open vs closed divertors. *Nuclear Fusion*, 60(7), 2020.
- [274] F. Wagner. A quarter-century of H-mode studies. *Plasma Physics and Controlled Fusion Related content*, 49(B1), 2007.
- [275] T. E. Evans, R. A. Moyer, P. R. Thomas, J. G. Watkins, T. H. Osborne, J. A. Boedo, E. J. Doyle, M. E. Fenstermacher, K. H. Finken, R. J. Groebner, M. Groth, J. H. Harris, R. J. La Haye, C. J. Lasnier, S. Masuzaki, N. Ohyaabu, D. G. Pretty, T. L. Rhodes, H. Reimerdes, D. L. Rudakov, M. J. Schaffer, G. Wang, and L. Zeng. Suppression of large edge-localized modes in high-confinement DIII-D plasmas with a stochastic magnetic boundary. *Physical Review Letters*, 92(23):9–12, 2004.
- [276] P. T. Lang, A. Loarte, G. Saibene, L. R. Baylor, M. Becoulet, M. Cavinato, S. Clement-Lorenzo, E. Daly, T. E. Evans, M. E. Fenstermacher, Y. Gribov, L. D. Horton, C. Lowry, Y. Martin, O. Neubauer, N. Oyama, M. J. Schaffer, D. Stork, W. Suttrop, P. Thomas, M. Tran, H. R. Wilson, A. Kavin, and O. Schmitz. ELM control strategies and tools: Status and potential for ITER. *Nuclear Fusion*, 53(4), 2013.
- [277] B. A. Grierson, K. H. Burrell, A. M. Garofalo, W. M. Solomon, A. Diallo, and M. O’Mullane. Response of impurity particle confinement time to external actuators in QH-mode plasmas on DIII-D. *Nuclear Fusion*, 54(11), 2014.
- [278] S. Mordijck, R. A. Moyer, and G. R. McKee. Changes in density fluctuations as a result of resonant magnetic perturbations correlate with the density inverse scale length. *Physics of Plasmas*, 19(2):24504, 2 2012.

- [279] B. A. Grierson, X. Yuan, M. Gorelenkova, S. Kaye, N. C. Logan, O. Meneghini, S. R. Haskey, J. Buchanan, M. Fitzgerald, S. P. Smith, L. Cui, R. V. Budny, and F. M. Poli. Orchestrating TRANSP simulations for interpretative and predictive tokamak modeling with OMFIT. *Fusion Science and Technology*, 74(1-2):101–115, 8 2018.
- [280] A. Marinoni, S. Brunner, Y. Camenen, S. Coda, J. P. Graves, X. Lapillonne, A. Pochelon, O. Sauter, and L. Villard. The effect of plasma triangularity on turbulent transport: Modeling TCV experiments by linear and non-linear gyrokinetic simulations. *Plasma Physics and Controlled Fusion*, 51(5), 2009.
- [281] A. Marinoni and M.E. Austin and A.W. Hyatt and S. Saarelma and F. Scotti and Z. Yan and C. Chrystal and S. Coda and F. Glass and J.M. Hanson and A.G. McLean and D.C. Pace and C. Paz-Soldan and C.C. Petty and M. Porkolab and L. Schmitz and F. Sciortino and S.P. Smith and K.E. Thome and F. Turco and the DIII-D Team. Diverted Negative Triangularity plasmas on DIII-D: the benefit of high confinement without the liability of an edge pedestal. *submitted*.
- [282] M. Fontana, L. Porte, S. Coda, and O. Sauter. The effect of triangularity on fluctuations in a tokamak plasma. *Nuclear Fusion*, 58(2), 2018.
- [283] M. E. Austin, A. Marinoni, M. L. Walker, M. W. Brookman, J. S. DeGrassie, A. W. Hyatt, G. R. McKee, C. C. Petty, T. L. Rhodes, and C. Sung. Exploration of negative triangularity discharges on DIII-D. In *44th EPS Conference on Plasma Physics, EPS 2017*, pages 5–8, 2017.
- [284] A. Marinoni, M. E. Austin, A W Hyatt, M L Walker, J Candy, C Chrystal, C J Lasnier, G. R. McKee, T. Odstrčil, C. C. Petty, M. Porkolab, J. C. Rost, O Sauter, S. P. Smith, G. M. Staebler, C. Sung, K. E. Thome, A D Turnbull, and L. Zeng. H-mode grade confinement in L-mode edge plasmas at negative triangularity on DIII-D. *Physics of Plasmas*, 26(4), 2019.
- [285] M. E. Austin, A. Marinoni, M. L. Walker, M. W. Brookman, J. S. Degrassie, A. W. Hyatt, G. R. McKee, C. C. Petty, T. L. Rhodes, S. P. Smith, C. Sung, K. E. Thome, and A. D. Turnbull. Achievement of Reactor-Relevant Performance in Negative Triangularity Shape in the DIII-D Tokamak. *Physical Review Letters*, 122(11):115001, 2019.
- [286] S. Yu Medvedev, M. Kikuchi, L. Villard, T. Takizuka, P. Diamond, H. Zushi, K. Nagasaki, X. Duan, Y. Wu, A. A. Ivanov, A. A. Martynov, Yu Yu Poshekhonov, A. Fasoli, and O. Sauter. The negative triangularity tokamak: Stability limits and prospects as a fusion energy system. *Nuclear Fusion*, 55(6), 2015.
- [287] A. Marinoni, J. C. Rost, M. Porkolab, A. E. Hubbard, T. H. Osborne, A. E. White, D. G. Whyte, T. L. Rhodes, E. M. Davis, D. R. Ernst, and K. H. Burrell. Characterization of density fluctuations during the search for an I-mode regime on the DIII-D tokamak. *Nuclear Fusion*, 55(9):93019, 2015.
- [288] S. Saarelma, M.E. Austin, A. Marinoni, M. Knolker, C. Paz-Soldan, L. Schmitz, P.B. Snyder, and DIII-D Team Team. LH-transition Prevented by Ballooning Instability in Negative Triangularity DIII-D Discharges. In *APS Division of Plasma Physics Meeting Abstracts*, volume 2020 of *APS Meeting Abstracts*, page GP19.003, January 2020.

- [289] W. Han, N. Offeddu, T. Golfopoulos, C. Theiler, C. K. Tsui, J. A. Boedo, E. S. Marmor, and TCV Team. Suppression of first-wall interaction in negative triangularity plasmas on TCV. *Nuclear Fusion*, 61(034003), 2021.
- [290] M. Cavedon, T. Pütterich, E. Viezzer, R. Dux, B. Geiger, R. M. McDermott, H. Meyer, and U. Stroth. A fast edge charge exchange recombination spectroscopy system at the ASDEX Upgrade tokamak. *Review of Scientific Instruments*, 88(4):43103, 2017.
- [291] R. K. Janev, D. Reiter, and U. Samm. Collision Processes in Low-Temperature Hydrogen Plasmas. Technical report, 2003.
- [292] R. A. Pitts, X. Bonnin, F. Escourbiac, H. Frerichs, J. P. Gunn, T. Hirai, A. S. Kukushkin, E. Kaveeva, M. A. Miller, D. Moulton, V. Rozhansky, I. Senichenkov, E. Sytova, O. Schmitz, P. C. Stangeby, G. De Temmerman, I. Veselova, and S. Wiesen. Physics basis for the first ITER tungsten divertor. *Nuclear Materials and Energy*, 20(July):100696, 2019.
- [293] E. Higson, W. Handley, M. Hobson, and A. Lasenby. Dynamic nested sampling: an improved algorithm for parameter estimation and evidence calculation. *Statistics and Computing*, 29(5):891–913, 9 2019.
- [294] L. A. Vainshtein and U. I. Safronova. Dielectronic satellite spectra for highly charged H-like ions ($2l'3l''-1s2l$, $2l'3l''-1s3l$) and He-like ions ($1s2l'3l''-1s22l$, $1s2l'3l''-1s23l$) with $Z = 6-33$. *Atomic Data and Nuclear Data Tables*, 25(4):311–385, 1980.
- [295] D. Stotler and C. Karney. Neutral Gas Transport Modeling with DEGAS 2. *Contributions to Plasma Physics*, 34(2-3):392–397, 1994.