

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Tokamak Pedestal Region Profile Prediction in Transport-Limited Regimes

Permalink

<https://escholarship.org/uc/item/74w9m78d>

ISBN

9798193291312

Author

Molesworth, Steven

Publication Date

2026-10-13

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Tokamak Pedestal Region Profile Prediction in Transport-Limited Regimes

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Engineering Sciences (Engineering Physics)

by

Steven Molesworth

Committee in charge:

Christopher Holland, Chair
Alessandro Marinoni
Matthias Morzfeld
George Tynan

2026

Copyright

Steven Molesworth, 2026

All rights reserved, except where otherwise noted.

Chapter 3, with the exception of the leading paragraph, is a reprint of *Filamentary velocity scaling validation and spin dynamics in the DIII-D tokamak*, © 2024 by S. C. Molesworth, J. A. Boedo, C. K. Tsui, R. Perillo and D. L. Rudakov, published in *Physics of Plasmas* **31**, 042515 (2024), and is licensed under CC BY 4.0. To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

The dissertation of Steven Molesworth is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2026

TABLE OF CONTENTS

Dissertation Approval Page	iii
Table of Contents	iv
List of Abbreviations	vi
List of Symbols	viii
List of Figures	x
List of Tables	xiii
Acknowledgements	xiv
Vita	xv
Abstract of the Dissertation	xvi
Chapter 1 Introduction	1
1.1 Motivation: from energy demand to transport-limited fusion	1
1.2 The core–edge gap	2
1.3 State of the art and the gap this thesis addresses	4
1.4 Aim, scope, and approach	7
Chapter 2 Background	10
2.1 Core transport	10
2.2 Scrape-off layer	13
2.3 Near-edge and pedestal physics	15
Chapter 3 Filamentary Velocity Scaling and Spin Dynamics in DIII-D	19
Chapter 4 The Near-Edge Profile Prediction Framework	38
4.1 Framework overview	38
4.2 Profile parameterizations	39
4.2.1 Expressiveness against the alternatives	42
4.2.2 Tail inference is the binding lever	43
4.3 Sources and losses	48
4.4 Rotation and the radial electric field	52
4.5 Transport	53
4.5.1 Neoclassical fluxes	53
4.5.2 Quasilinear theory and the gyrofluid closure	55
4.5.3 Saturation rules	56
4.5.4 Near-edge configuration	57
4.5.5 MHD ceiling	57

4.5.6	A minimal analytic alternative	58
4.6	Flux matching	59
4.7	Uncertainty quantification	60
Chapter 5	Validation	64
5.1	Validation approach	64
5.2	Results across the set	66
5.2.1	The flux match improves every case and closes none	66
5.2.2	Pedestal top pressure	67
5.2.3	Turbulent fingerprints across the set	69
5.2.4	Pedestal error does not track any dimensionless parameter	71
5.3	L-mode	72
5.4	High-collisionality H-mode	73
5.5	Low-collisionality H-mode	74
5.6	Negative triangularity	75
5.7	Alcator C-Mod I-mode	76
Chapter 6	Error Analysis	88
6.1	Approach	88
6.2	Limitations from the inputs alone	88
6.3	The objective	90
6.3.1	Equal residuals, different pedestals	90
6.3.2	Most of the pedestal is never constrained	91
6.3.3	Interior gradients depend on the separatrix conditions	93
6.4	The transport model	94
Chapter 7	Conclusions and Future Work	100
7.1	Conclusions	100
7.2	Future directions	101
Appendix A	Gaussian-Process Flux Surrogates	103
A.1	The surrogate problem	103
A.2	Gaussian-process regression	103
A.3	Training	105
A.4	Uncertainty: two separate sources	106
A.5	Validation of the predictive variance	106
Appendix B	A Minimal Near-Edge Transport Model	107
B.1	Equations and geometry	108
B.2	Fluxes	108
B.3	Boundary conditions and solution	111
B.4	Standing limitations	111
Bibliography		113

LIST OF ABBREVIATIONS

ADAS	Atomic Data and Analysis Structure (atomic rate database)
ARD	Automatic relevance determination
AURORA	Impurity transport and collisional–radiative package
BOUT++	Boundary plasma turbulence code
CGYRO	Nonlinear gyrokinetic turbulence code
C-Mod	Alcator C-Mod tokamak
DIII-D	DIII-D tokamak
EIRENE	Kinetic neutral transport code
ELM	Edge-localized mode
EPED	Peeling–ballooning/kinetic-ballooning pedestal model
EPED1NN	Neural-network surrogate of EPED
ETG	Electron temperature gradient (mode)
GACODE	General Atomics code suite and its plasma state format
GKEYLL	Continuum gyrokinetic edge code
GP	Gaussian process
H-mode	High-confinement mode
I-mode	Improved energy confinement mode
ITER	International burning-plasma tokamak experiment
ITG	Ion temperature gradient (mode)
KBM	Kinetic ballooning mode
LCFS	Last closed flux surface (separatrix)
L-mode	Low-confinement mode
MHD	Magnetohydrodynamic
NEO	Drift-kinetic neoclassical transport code
NT	Negative triangularity

PORTALS	Surrogate-based core transport optimization framework
QH	Quiescent H-mode
RMS	Root mean square
SAT	Saturation rule (SAT2, SAT3) of TGLF
SMARTS	Solver ecosystem this framework is built into
SOL	Scrape-off layer
SSF	Sheared spectral filament (separatrix gradient model)
TEM	Trapped electron mode
TGLF	Trapped gyro-Landau fluid quasilinear transport model
TGYRO	Steady-state core transport flux matching solver
TLP	Transport-limited pedestal
WPED	Wide-pedestal (H-mode state)
XGC1	Particle-in-cell gyrokinetic edge code

LIST OF SYMBOLS

Latin

a	Plasma minor radius
a/L_y	Normalized inverse scale length of a profile y , $-(a/y) dy/dr$
B_T	Toroidal magnetic field
c_s	Ion sound speed
E_r	Radial electric field
f_Z	Impurity concentration, n_Z/n_e
k_y	Binormal wavenumber; $k_y \rho_s$ when normalized
$k_\perp$	Perpendicular wavenumber
L_y	Scale length of a profile y
l_{mix}	Turbulent mixing length
n_e, n_i	Electron and main-ion density
n_Z	Impurity density
p_{ped}	Pedestal-top pressure
P_{fus}	Fusion power
P_{ei}	Electron–ion collisional energy exchange
q	Safety factor
Q_e, Q_i	Electron and ion energy flux
r	Flux-surface minor radius; r/a when normalized
R_0	Plasma major radius
S_{wall}	Wall neutral particle source
T_e, T_i	Electron and ion temperature
V'	Flux-surface volume derivative, dV/dr
Z_{eff}	Effective ion charge

Greek

α_{MHD}	Normalized pressure-gradient (ballooning) parameter
α_t	Turbulence (resistive ballooning) parameter
β	Ratio of kinetic to magnetic pressure
Γ_e	Electron particle flux
γ_E	$E \times B$ shearing rate
γ_{lin}	Linear growth rate
δ	Triangularity
ℓ_d	Gaussian-process kernel length scale in dimension d
κ	Elongation
λ_q	Scrape-off-layer heat-flux width
ν_*	Normalized collisionality
ρ_s	Ion sound gyroradius
ρ_*	Normalized gyroradius, ρ_s/a
σ	Standard deviation
ϕ	Electrostatic potential
χ	Thermal diffusivity
ψ_N	Normalized toroidal magnetic flux

Subscripts and superscripts

e, i, Z	Electron, main-ion and impurity species
sep	Evaluated at the separatrix
ped	Evaluated at the pedestal top
nc	Neoclassical contribution
gB	Normalized in gyroBohm units

LIST OF FIGURES

Figure 1.1.	Tokamak magnetic geometry	2
Figure 1.2.	Poloidal cross-section of a diverted tokamak	4
Figure 1.3.	MHD-limited versus transport-limited pedestals	5
Figure 2.1.	Near-edge profiles and their inverse scale lengths	18
Figure 3.1.	Normalized velocity scaling against collisionality and size for the RB RX Ci and Cs regimes	21
Figure 3.2.	Toroidal blob spin and the weakening of the filament charge dipole	22
Figure 3.3.	Reciprocating probe head pin layout and example magnetic equilibrium ..	23
Figure 3.4.	Poloidal and radial velocities against $R - R_{\text{sep}}$ for plunge 1 of shot 174208	24
Figure 3.5.	Radial velocity distribution from the $E_p \times B_T$ drift	24
Figure 3.6.	Filament diameter and normalized radius distributions	25
Figure 3.7.	Instability time scales for small and large filaments	25
Figure 3.8.	Normalized collisionality estimates against $R - R_{\text{sep}}$	26
Figure 3.9.	Normalized collisionality against normalized size coloured by measured velocity	26
Figure 3.10.	Distribution of measured-to-theoretical normalized velocity ratios	27
Figure 3.11.	Normalized velocity against size at low collisionality and in the RB regime	27
Figure 3.12.	Normalized velocity against collisionality for intermediate normalized sizes	28
Figure 3.13.	Filament electron temperature traces at varying normalized temperature ratio	28
Figure 3.14.	Radial velocity against normalized temperature ratio	28
Figure 3.15.	Spin tangential velocity against temperature ratio and radial velocity against spin	29
Figure 3.16.	Sheath current and spin timescales against normalized temperature ratio ..	29
Figure 3.17.	Flux surface against collisionality and radial velocity adjusted for $T_i > T_e$.	31

Figure 3.18.	Difference between measured and theoretical normalized velocity	32
Figure 3.19.	Distribution of velocity ratios across the sensitivity analysis	32
Figure 4.1.	The near-edge profile-prediction workflow	39
Figure 4.2.	Profile parameterization expressiveness (H-mode)	44
Figure 4.3.	Profile parameterization expressiveness (high- v^* H-mode)	45
Figure 4.4.	Profile parameterization expressiveness (L-mode)	46
Figure 4.5.	Sensitivity of the inferred gradient tail to the flux match knots	47
Figure 4.6.	Sensitivity to the separatrix (LCFS) gradient	48
Figure 4.7.	Neutral module behavior	50
Figure 4.8.	Rotation terms at the initial state	54
Figure 4.9.	Propagation of the input covariance to the objective	63
Figure 5.1.	Flux match scorecard across the validation set	68
Figure 5.2.	Pedestal top scorecard across the validation set	70
Figure 5.3.	L-mode (DIII-D 128913): converged near edge	78
Figure 5.4.	L-mode: turbulent-flux mode breakdown	79
Figure 5.5.	High- v^* H-mode: converged near edge	80
Figure 5.6.	High- v^* H-mode: turbulent-flux mode breakdown	81
Figure 5.7.	Low- v^* H-mode: converged near edge	82
Figure 5.8.	Low- v^* H-mode: turbulent-flux mode breakdown	83
Figure 5.9.	Negative triangularity: converged near edge	84
Figure 5.10.	Negative triangularity: turbulent-flux mode breakdown	85
Figure 5.11.	Alcator C-Mod I-mode: converged near edge	86
Figure 5.12.	Alcator C-Mod I-mode: turbulent-flux mode breakdown	87
Figure 6.1.	Input covariance band against tolerance and against error	90

Figure 6.2.	Equal residuals, different pedestals	92
Figure 6.3.	The unconstrained pedestal and the cost of knot placement	93
Figure 6.4.	Converged profiles under a perturbed separatrix gradient	98
Figure 6.5.	A missing collisionality dependence, and what it does not fix	99
Figure A.1.	What the GP surrogate does	104

LIST OF TABLES

Table 2.1.	Physics and modeling maturity across the three regions	11
Table 3.1.	Database statistics of the DIII-D discharge set	22
Table 3.2.	Summary of measurement uncertainties	30
Table 4.1.	$E \times B$ shear verification against Haskey Table I	53
Table 5.1.	Global parameters of the validation discharges	66
Table 5.2.	Per-case validation inputs	66
Table 5.3.	Predicted versus measured pedestal top pressure	69

ACKNOWLEDGEMENTS

I am grateful to my advisor, Chris Holland, who supported my change of research topic partway through this degree, well aware that my background and life circumstances would make it challenging. His commitment to my success undoubtedly helped me persevere.

I would also like to thank the members of my doctoral committee for their time and input on this thesis; the UCSD CER boundary group for mentoring me through the completion of our filament velocity scaling study; the MIT MFE integrated modeling group for fielding code-related questions; and Sergei Krashennnikov for originally welcoming me into the program.

To Michelle, your patience, encouragement, and support powered me through these years. You agreed to a major life change, then proceeded to make countless sacrifices along the way. Thank you for keeping things in perspective for me.

To Sloane, you were not with us for the entirety of this adventure, but you certainly added motivation for me to complete it. More sleep in the final months of this program would have been helpful, but that feels quite meaningless relative to seeing you grow.

To Kingsley, you never judged me by my research, only by the quantity of belly rubs and walks I could provide. Man's best friend is an apt phrase for you.

This material is based upon work supported by the US Department of Energy, Office of Science, Office of Fusion Energy Sciences, under Award Number DE-SC0024399. This report was prepared as an account of work sponsored by an agency of the United States Government. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Chapter 3, with the exception of the leading paragraph, is a reprint of the material as it appears in Molesworth, Steven; Boedo, Jose; Tsui, Cedric; Perillo, Renato; Rudakov, Dmitry, "Filamentary velocity scaling validation and spin dynamics in the DIII-D tokamak," *Physics of Plasmas* 31 (4), 042515 (1 April 2024), <https://doi.org/10.1063/5.0195025>; published by AIP Publishing under a Creative Commons Attribution (CC BY) license, which grants permission to reproduce it here. The dissertation author was the primary investigator and author of this paper.

VITA

2014	Bachelor of Science in Mechanical Engineering, Western Michigan University
2015	Master of Science in Mechanical Engineering, Western Michigan University
2026	Doctor of Philosophy in Engineering Sciences (Engineering Physics), University of California San Diego

PUBLICATIONS

S. C. Molesworth, J. A. Boedo, C. K. Tsui, R. Perillo, and D. L. Rudakov, “Filamentary velocity scaling validation and spin dynamics in the DIII-D tokamak,” *Physics of Plasmas* **31**, 042515 (2024).

S. C. Molesworth and C. Holland, “Modeling of transport-limited, near-edge tokamak profiles using control volumes,” *Bulletin of the American Physical Society*, 65th Annual Meeting of the APS Division of Plasma Physics, JP11.096 (2023).

“Hydraulic Mount,” U.S. Patent 11,287,011 B2 (2022).

“Hydraulic Mount,” U.S. Patent 10,927,922 B2 (2021).

“Hydraulic Body Mount,” U.S. Patent 10,215,252 B2 (2019).

ABSTRACT OF THE DISSERTATION

Tokamak Pedestal Region Profile Prediction in Transport-Limited Regimes

by

Steven Molesworth

Doctor of Philosophy in Engineering Sciences (Engineering Physics)

University of California San Diego, 2026

Christopher Holland, Chair

Accurately predicting temperature and density profiles in the near-edge region—roughly the outer 10% of closed flux surfaces—is critical for modeling tokamak performance because fusion power scales with the square of pedestal-top pressure. Building upon an existing framework for accelerated core transport solutions, this dissertation presents a workflow for self-consistently solving transport-limited pedestals using 1D, steady-state flux matching. Inverse scale lengths are represented using an asymmetric hyperbolic tangent parameterization, anchored at the separatrix. Modular near-edge physics including neutral and impurity penetration are added to account for their effects on source terms in the flux matching solution. Diamagnetic and neoclassical terms in the rotation profile are self-consistently updated, while the toroidal velocity is held

constant at this stage. Neoclassical and turbulent transport models developed for core plasmas were chosen to determine their range of validity. Five discharges spanning L-mode, H-mode at contrasting collisionalities, negative triangularity, and I-mode are used to test the workflow. Only the NT result falls within the desired sub-10% error in pedestal-top pressure for reactor predictions, while the remaining cases miss the experimental pedestal-top pressure by 17–84%. That error is attributed to three factors: uncertainties in the separatrix boundary conditions, degeneracy between solver objective and pedestal values, and transport model limitations most severe in the region that most strongly determines pedestal shape. Together, these errors do not exclude flux matching as a solution to the transport-limited pedestal prediction problem, but do suggest limitations of the approach, particularly until an accurate near-edge turbulence model can be made tractable within an iterative framework. This work is thus a step toward a unified, self-consistent treatment of core-edge coupling suitable for integrated tokamak modeling.

Chapter 1

Introduction

1.1 Motivation: from energy demand to transport-limited fusion

Rising global energy demand strains the sustainability of current power systems [1]. Nuclear fusion is an attractive response, offering abundant energy without carbon emissions or the long-lived waste of fission, yet it demands conditions that are extraordinarily hard to create and sustain. Nuclei must collide frequently and with enough energy to overcome their Coulomb repulsion, a tall order because the scattering cross section exceeds the fusion cross section by orders of magnitude [2]. Electrons dissociate from their nuclei at these temperatures, leaving the charged-particle gas called plasma. The central challenge is therefore confinement, holding this “sun in a bottle” at high temperature and pressure while meeting the material constraints of exhaust and energy extraction.

Magnetic confinement is the leading approach to maintaining these conditions. The Lorentz force ties charged particles to magnetic field lines, forcing them to gyrate about the field and suppressing motion across it, while leaving rapid transport along the field unimpeded. Among magnetic configurations the tokamak has proven the most successful architecture for approaching net-energy conditions [3]. It is a toroidal device combining external toroidal field coils with the poloidal field of an induced plasma current, shown in Figure 1.1. The superposed fields trace out nested, closed toroidal flux surfaces on which plasma parameters are nearly




Figure 1.1. Tokamak magnetic geometry. External toroidal field coils and the poloidal field of an induced plasma current combine into a helical field tracing out nested, closed toroidal flux surfaces, with a wedge cut away to expose them. The field line is wound at one poloidal turn over the arc shown for legibility.

constant. This geometry must eventually be opened to exhaust fusion ash and impurities thereby introducing a different set of dominant physics.

Reactor-relevant tokamaks are designed for high-confinement regimes whose performance is set by local *transport*, the turbulent and collisional leakage of heat and particles across flux surfaces, rather than by global magnetohydrodynamic (MHD) stability limits. Predicting the plasma profiles that emerge in these transport-limited regimes is the problem this work aims to address, particularly across the steep near-edge region that couples the hot core to the cool exhaust.

1.2 The core–edge gap

The near-edge is the hinge of integrated tokamak modeling. It sets the boundary condition for core transport solvers, supplies the upstream condition for scrape-off-layer (SOL) and divertor

models, and largely fixes the global confinement on which fusion performance depends, because stiff core transport carries the pedestal top up through the profile beneath it. Figure 1.2 locates the region, a thin band spanning the outer $\sim 10\text{--}15\%$ of the closed flux surfaces. It is bounded on the inside by the stiff core and on the outside by the last closed flux surface (LCFS), or separatrix, beyond which field lines are diverted to the target plates.

Prediction, however, effectively stops at an assumed boundary. Core flux matching solvers terminate at a fixed core–edge interface at which density and temperatures are supplied rather than solved, and the pedestal tools that take over beyond it return a pressure rather than the three kinetic channels separately (Section 1.3). The ion temperature is as critical as the electron temperature, both for its contribution to the pedestal pressure and for setting ion-scale turbulent transport. The gap is therefore sharp, because no method self-consistently predicts the near-edge density and both temperatures together across the full region. Compounding this, the pedestal takes qualitatively different forms depending on what limits it (Figure 1.3). An MHD-limited pedestal is clamped at the peeling–ballooning boundary in ELMy H-mode, whereas the softer, transport-limited pedestals (TLPs) of regimes like quiescent H-mode (QH), the wide-pedestal state (WPED), and I-mode are the target of this thesis.

Reactor performance is steeply sensitive to the pedestal top. Stiff ion-temperature-gradient turbulence holds the core gradient near its critical value, so the core profile rises and falls with the pedestal top and the fusion power follows it. In the ARC physics basis a 20% increase in the assumed pedestal pressure raises predicted fusion power from 878 MW to ~ 1.24 GW (+41%), while varying the pedestal T_i/T_e over 0.9–1.1 spans 698–1134 MW [4]. Reactor scoping thus rests on pedestal values that are, today, assumed [5].

This thesis narrows that chasm by extending self-consistent, flux matched prediction of all three kinetic channels from $\rho \approx 0.9$ out to the LCFS. Note that throughout this text ρ may be used loosely to mean either the normalized effective radius $\sqrt{\psi_N}$, where ψ_N is the normalized toroidal magnetic flux, or the normalized minor radius, r/a , because of near equivalence approaching the separatrix. The approach remains, by construction, naive to the SOL, requiring semi-informed,




Figure 1.2. Poloidal cross-section of a diverted tokamak, identifying the core, the near-edge region modeled in this work (shaded, $0.85 \leq r/a \leq 1$), the separatrix, and the scrape-off layer. The equilibrium is the analytic single-null solution of [6] at DIII-D-like shaping, and is not fitted to any discharge.

user-supplied $(n_e, T_e, T_i)_{\text{LCFS}}$ and $(a/L_{n_e}, a/L_{T_e}, a/L_{T_i})_{\text{LCFS}}$. In the future, these could be evolved with the profiles through scrape-off layer modeling, but that is beyond the scope of this work. Why the near-edge resists the direct application of core methods, through the loss of scale separation and the breakdown of quasilinear transport theory, is taken up in Sections 1.3 and 2.3.

1.3 State of the art and the gap this thesis addresses

Predicting kinetic profiles by flux matching is well established in the core. TGYRO [7] solves the flux surface averaged, steady-state transport balance by adjusting the inverse gradient scale lengths a/L_γ at a handful of radii until the modeled turbulent-plus-neoclassical fluxes match




Figure 1.3. Illustrative (synthetic, normalized) edge pressure profiles contrasting an MHD-limited pedestal, clamped at the peeling–ballooning limit as in EPED-predicted ELMy H-mode, with a softer transport-limited pedestal characteristic of QH-, WPED-, and I-mode regimes. The shaded band marks the near-edge region modeled in this work. Profiles are synthetic and normalized to the MHD pedestal top pressure, and are not fitted to any discharge.

the volume-integrated sources, with the profiles recovered by integrating a/L_y inward from a fixed boundary. PORTALS accelerates this by replacing the expensive turbulence model with Gaussian process surrogates refined during Bayesian optimization, and by training on physically motivated dimensionless parameters rather than raw profile values, which sharply reduces the number of high-fidelity evaluations [5, 8]. These tools predict core temperature and density reliably, but treat everything beyond $\rho \approx 0.9$ as fixed.

Extending them to the separatrix is not simply a matter of moving that boundary, because the near-edge is qualitatively harder and not merely quantitatively so. The quantitative part is real enough, turbulent flux amplitudes there being roughly an order of magnitude larger than in the core while several instability branches compete across a narrow region, so the resulting optimization landscape carries many more local minima than the gentle, near-convex core problem. What changes in kind is the standing of the models themselves, because the separation between fluctuation and equilibrium scales that underpins local, quasilinear transport weakens, so

the reduced models are being asked to work outside the ordering they were derived under rather than merely at harder values of their arguments. Attempts to apply core machinery regardless have had qualified success. The TGLF SAT3 saturation rule [9, 10] lowered transport errors in TGYRO calculations of DIII-D L-mode edges, and full-radius TGLF gave acceptable confinement predictions for ASDEX Upgrade L-modes [11], even as the underlying quasilinear approximation has since been shown to break down in the edge outright [12]. Direct core gyrokinetic codes are seldom converged beyond $\rho \approx 0.9$ owing to cost and convergence difficulty.

For the pedestal proper, the state of the art is EPED [13]. It predicts the height and width of the pressure pedestal by locating the state marginally stable to both ideal peeling–ballooning and kinetic-ballooning modes. The profiles are represented by a tanh anchored at a prescribed pedestal top density and separatrix temperature. EPED is well validated for ELMy H-modes at or near the Type-I ELM threshold, and its neural-network surrogate EPED1NN [14] delivers the same prediction at negligible cost. Its scope is narrower than is sometimes appreciated, because it returns a pressure rather than independent density and temperature, and its marginal peeling–ballooning premise pertains to MHD-limited pedestals, not to the transport-limited regimes that increasingly define reactor-relevant scenarios. Recent work has begun to relax these limits. For example, the Parisi model [15] extends EPED with an analytic electron-temperature-gradient transport model and gyrokinetic kinetic-ballooning constraints. The two additions decouple the density pedestal from the temperature pedestal, and they also identify $E \times B$ flow shear as a necessary ingredient of a TLP. The WPED model [16] instead targets the weak pedestals of negative triangularity. These are genuine advances, but they remain within the stability-and-analytic paradigm rather than solving the coupled transport problem across all channels.

True flux matching across the near-edge has only recently become tractable, and so far only at the high-fidelity extreme. TCV profiles have been evolved to a flux matched steady state with the edge gyrokinetic code GKEYLL and a self-consistent particle-accounting scheme, reproducing measured profiles [17]. That is a landmark result, but its per-case cost precludes

routine or multi-regime use. The gap this thesis addresses lies between these poles, because there is no *fast*, self-consistent, transport-based solver that predicts the near-edge density, electron temperature, and ion temperature together across the full L-mode-to-H-mode range, connects naturally to existing core and SOL models, and is free of per-device calibration.

1.4 Aim, scope, and approach

The framework built to fill that gap is neither a comprehensive theory of near-edge transport nor a finished production tool. It is a starting framework, intended for everyday scenario and device design rather than for a bespoke research campaign mounted per discharge, and several constraints bound it deliberately from the outset. Profiles are obtained by matching transport fluxes to sources rather than by imposing an MHD stability ceiling as in EPED, so that density and both temperatures are predicted self-consistently rather than a pressure alone. Turbulent and neoclassical fluxes are evaluated with the quasilinear TGLF and drift-kinetic NEO models. Nonlinear, full- f gyrokinetic and particle-in-cell codes are more faithful but far too costly to iterate to a flux matched steady state in routine use, an approach only now being attempted [17, 18]. The intent is to carry the established core transport models as far toward the separatrix as they will credibly go, not to develop a new edge-specific transport theory.

The remaining restrictions are simplifications adopted to keep the method tractable and honest about what it does *not* attempt. Only stationary, transport-limited profiles are sought, so ELM cycles, L–H transition dynamics, and the zonal- and mean-flow interplay that governs them are out of scope. The magnetic equilibrium is held fixed during the solve, which defers self-consistent equilibrium update and full core–edge coupling to future work. Edge-specific sources and sinks that cannot be omitted, namely neutral penetration and ionization, impurity radiation, and a self-consistent diamagnetic flow, are included through reduced models like a prescribed impurity profile shape rather than first-principles treatments. The framework is built to plug into the existing SMARTS solver ecosystem, reusing its core machinery rather than standing

alone.

One could object that a surrogate trained on high-fidelity gyrokinetics by active learning over a dimensionless parameter space would sidestep much of this. That path is real and, for raw accuracy, likely superior, but it is complementary rather than precluding. The training campaign is enormous and must be re-extended per device, where extrapolation is least trusted. A flux surrogate also predicts *fluxes*, not profiles, so it still requires the profile representation, edge source and sink modules, and boundary treatment developed here. The modular workflow is built so that such a surrogate is exactly the drop-in the transport component anticipates (Section 4.1), while the reduced, physics-based route yields understanding a black box does not, namely which instability carries which channel and where the quasilinear picture fails.

A recurring theme is that the near-edge rewards exploiting the known structure of its profiles. Core flux matching solvers construct profiles from piecewise-linear a/L_y segments, which suffice where a/L_y varies by $\mathcal{O}(1)$ across the domain, but a/L_y varies by $\mathcal{O}(10)$ over the narrow near-edge. The framework therefore pairs a compact profile parameterization with modular edge-physics modules and an edge-tailored transport model, coupled by flux matching and developed in Chapter 4.

The contributions that follow are as much a *characterization of the problem* as a tool. The framework does not predict near-edge profiles to the reactor-handoff tolerance, and the account of why is the more durable result. The central deliverable is an integrated, modular near-edge workflow that carries the core flux matching loop out to the last closed flux surface and predicts n_e , T_e , and T_i together. Its edge sources, boundary treatment, stability ceiling, and transport model are individually validatable, replaceable components, and every prediction carries a propagated input-covariance error bar (Chapter 4). Within it, a profile parameterization built for the edge, a modified hyperbolic tangent with a superposed corrector spline defined on a/L_y , reproduces the profiles and gradients of real discharges with a handful of design variables (Section 4.2). The neutral, impurity, rotation, and stability modules are each checked against expected physical behavior rather than calibrated to the cases (Chapter 4).

The results that reframe the rest concern the problem statement itself. States that the objective scores alike carry materially different pedestals, so the near-edge flux match as posed here does not determine the pedestal, and the share of the profile lying outside the constrained region accounts for that quantitatively (Sections 6.3.1 and 6.3.2). The transport model is separately and identifiably wrong, because reverse-engineering the saturation layer against plasma state and power-balance targets identifies a collisionality dependence that no existing saturation rule contains, and two distinct quasilinear failures are documented at the radii that set the pedestal height (Section 6.4). Propagating the declared input covariance gives a 1σ band on pedestal top pressure that already exceeds the tolerance a device study would need, before any transport-model error is counted. Together these order the remaining work rather than merely listing it (Chapters 5 and 7).

The remainder of the thesis proceeds as follows. Chapter 2 reviews the physics of the core, scrape-off layer, and near-edge. Chapter 3 reprints an experimental filament study that motivated the turn to computation. Chapter 4 develops the framework, Chapter 5 tests it across a span of operating regimes, and Chapter 6 investigates the shortcomings the validation exposes. Chapter 7 concludes.

Chapter 2

Background

This chapter sketches the physics of the three regions this thesis straddles, the core, the scrape-off layer, and the near-edge between them, with Table 2.1 setting the three side by side. The emphasis throughout is on physics, and its algorithmic treatment is deferred to Chapter 4.

2.1 Core transport

The viability of a magnetically confined plasma as a net energy source is commonly gauged by the fusion triple product $nT\tau_E$, where n and T are characteristic density and temperature. The energy confinement time $\tau_E = W/P_{\text{loss}}$ is the ratio of the stored plasma kinetic energy to the rate of power loss [2, 19]. Operational ceilings exist on the ingredients, the Troyon limit on the normalized pressure β [20] and the empirical Greenwald limit on density [3, 21]. Within those bounds it is *transport*, the cross-field leakage of heat and particles, that sets τ_E and hence the achievable n and T , which makes core transport the central task of performance modeling.

Transport is conventionally decomposed into neoclassical and turbulent contributions. Neoclassical transport is collisional diffusion corrected for toroidal magnetic geometry. Trapping of particles in banana orbits raises the effective diffusivity roughly two orders of magnitude above the classical estimate, and it is computed accurately by drift-kinetic codes like NEO [22]. In most regimes, though, turbulent (“anomalous”) transport driven by micro-instabilities dominates, and its resistance to prediction is a large part of why fusion energy has proven so difficult [19].

Table 2.1. Physics and modeling maturity differ sharply across the core, the near-edge, and the scrape-off layer. The middle column is the region this thesis addresses.

	Core	Near-edge	Scrape-off layer
Characteristics	Weak gradients, closed surfaces, small fluctuation amplitude	Steep gradients over centimeters, sheared flow, neutral penetration	Open field lines, parallel losses dominant, fluctuations $\mathcal{O}(1)$
Orderings	Gyrokinetic ordering holds, gyroBohm diffusivity scaling	Gyrokinetic orderings break down, transition toward Bohm scaling	Scale separation absent, decay lengths from reduced models
Solved today	Profiles from flux matching, routinely and with validated accuracy	ELMy pedestal height from stability, not a transport solution	Decay lengths from reduced models, turbulence codes costly
Role here	Supplies source fluxes entering the near-edge	Self-consistent n_e, T_e, T_i with no imposed stability closure	Supplies the outer boundary condition, coupling left to future work

The instabilities responsible, the ion- and electron-temperature-gradient (ITG, ETG) modes and the kinetic ballooning mode (KBM) among others, share a critical gradient character. Each is destabilized only once its driving gradient exceeds a threshold, above which the associated flux rises [23]. How steeply it rises is the *stiffness*, and it varies by mode and by regime. Core ion transport is stiff enough that gradients are pinned near their thresholds, so the temperature and density at the top of the edge pedestal largely dictate the profiles beneath, and thus the fusion power, which scales approximately as $P_{\text{fus}} \sim p_{\text{ped}}^2$ [4, 5]. Stiffness is therefore the reason core prediction is both tractable and acutely sensitive to the boundary supplied at the pedestal top.

Where the fluctuation scale, the ion gyroradius ρ_i , is far smaller than the profile scale length, the ordering $\rho_* = \rho_i/a \ll 1$ permits the local, scale-separated gyrokinetic reduction on which most transport models rest [24]. Fluxes can then be evaluated across a hierarchy of fidelity and cost: (semi-)analytic formulae, the quasilinear TGLF model, trained on gyrokinetic data and equipped with saturation rules like SAT3 [9, 10], and nonlinear gyrokinetic codes like CGYRO [25], which are the most faithful but far the most expensive.

Transport modeling begins from the fluid conservation laws obtained by taking velocity moments of the gyrokinetic equation [3, 24]. For each species s the particle density, thermal

energy, and toroidal momentum obey balances of the generic conservation form $\partial_t(\cdot) + \nabla \cdot (\text{flux}) = (\text{source})$,

$$\frac{\partial n_s}{\partial t} + \nabla \cdot \mathbf{\Gamma}_s = S_{n,s}, \quad (2.1)$$

$$\frac{\partial}{\partial t} \left(\frac{3}{2} n_s T_s \right) + \nabla \cdot \mathbf{Q}_s = S_{E,s}, \quad (2.2)$$

$$\frac{\partial}{\partial t} (m_i n_i R^2 \Omega) + \nabla \cdot \mathbf{\Pi} = S_L, \quad (2.3)$$

with $\mathbf{\Gamma}_s$ the particle flux, $\mathbf{Q}_s$ the energy flux (conductive plus convective), and $\mathbf{\Pi}$ the flux of toroidal angular momentum $m_i n_i R^2 \Omega$. Of the sources and sinks on the right-hand sides, fueling and ionization enter $S_{n,s}$ and the injected torque enters S_L . The energy source $S_{E,s}$ carries auxiliary and ohmic heating, line and bremsstrahlung radiation, and the collisional electron–ion energy exchange $P_{ei} \propto n_e (T_e - T_i)$. The equations are coupled through the exchange term, through quasineutrality, and through the electromagnetic fields on which the fluxes depend, so they form a system rather than three independent balances. The present framework evolves the electron density and the electron and ion temperatures, with the radial electric field regenerated from the predicted pressure each iteration while the toroidal rotation is held at its measured value (Section 4.4), and the impurities entering through quasineutrality and radiation at a prescribed profile shape.

To make prediction tractable, the flux surface averaged (FSA) transport equations are solved in one radial dimension. Plasma parameters vary negligibly over a flux surface in the core, so the total volume is discretized into nested control volumes. In steady state the balance for a conserved scalar y reduces to

$$\frac{1}{V'} \frac{\partial}{\partial r} (V' \Gamma_y) = S_y, \quad \implies \quad \Gamma_y(r) = \frac{1}{V'} \int_0^r S_y V' dr', \quad (2.4)$$

where $V' = \partial V / \partial r$, Γ_y is the transport flux, and S_y is the volumetric source. Transport solvers like TGYRO [7] pose this as a root-finding problem. The inverse gradient scale lengths $z_y =$

$a/L_y = -(a/y) \partial y / \partial r$ at a handful of radii are adjusted until the modeled fluxes match the integrated sources, with the profiles recovered by integrating z_y inward from a fixed boundary. Fluxes are gyroBohm-normalized, for example $Q_{\text{gB}} = n_e T_e c_s (\rho_s/a)^2$ with sound speed c_s and sound-radius ρ_s , to fold in the known physical scaling and aid convergence.

The problem is hard even so. A representative case evolves T_e , T_i , n_e , and rotation at several radii, giving tens of coupled unknowns. The critical gradient stiffness renders the flux response nearly singular near threshold, so the landscape carries many local minima and demands many iterations of an expensive flux model [7]. PORTALS addresses this by replacing the flux model with Gaussian process surrogates refined under Bayesian optimization. Those surrogates are trained on a compact set of physically motivated dimensionless parameters rather than raw profile values, which sharply reduces the number of high-fidelity evaluations needed to converge [5, 8]. The Gaussian process surrogate machinery is summarized in Appendix A, and the present work utilizes it where applicable.

2.2 Scrape-off layer

Outside the separatrix, field lines are open and terminate on the divertor targets. Fast parallel streaming along them drains heat and particles before they can spread radially, producing a narrow boundary layer whose width is characterized by the heat-flux decay length λ_q , effectively the fraction of the divertor area doing the exhausting. A small λ_q concentrates the power, and peak divertor heat fluxes of $10\text{--}20 \text{ MW m}^{-2}$ are anticipated for ITER [26], near the tolerance of plasma-facing components. Setting the upstream conditions that fix λ_q and the target loads is thus a first-order design concern, and it is these upstream, separatrix conditions that a near-edge profile solver must ultimately be handed.

Turbulence broadens the layer. Alongside the usual eddies, the edge is populated by coherent field-aligned structures, blobs or filaments, that carry a disproportionate share of the cross-field transport [27]. Born near the LCFS, a blob is polarized by the curvature and ∇B drifts,

and the resulting internal electric field drives a rapid, radially outward $E \times B$ motion that ejects it into the SOL [28]. Filament scaling and spin dynamics are the subject of Chapter 3 [29].

Edge and SOL transport are considerably harder to model than the core because the scale separation underpinning local, quasilinear models breaks down in both space and time. The steep gradients that link the hot core to the cool boundary drive fluctuations of order unity relative to the background, with pronounced intermittency and nonlocality. Open field lines with a directional bias in the parallel flows produce in–out (high- versus low-field-side) asymmetries, and the larger poloidal variation of magnetic curvature erodes the validity of flux surface averaging.

The interchange and ballooning drive, set by the product of the field-line curvature with the pressure gradient, is destabilizing (“bad”) on the low-field side, where curvature and ∇p align to drive Rayleigh–Taylor-like motion, and stabilizing (“good”) on the high-field side. The unstable modes consequently balloon at the outboard midplane and are evanescent on the inboard. A flux surface stability measure like the normalized pressure gradient α_{MHD} averages over this poloidal localization [30]. That approximation itself weakens toward the separatrix, because the flux surfaces there are increasingly shaped and the local magnetic shear diverges as the field line approaches the X-point. Both the drive and the mode structure therefore vary strongly from the outboard midplane to the divertor legs. The framework inherits this limitation, since it evaluates transport on flux surface averaged profiles. Any flux surface averaged stability measure of this kind should thus be read with the poloidal localization of the true drive in mind.

These difficulties, together with the computational cost of resolving SOL turbulence over useful volumes and times, have so far precluded a TGYRO-like transport solver for the open-field-line region, so practical prediction instead relies on reduced models. The Goldston heuristic-drift model [31] estimates λ_q in the neoclassical limit, scaling it with the ion poloidal gyroradius, while the empirical Eich regression [32] reproduces measured λ_q across many devices. High-fidelity SOL codes, the fluid BOUT++ [33] and the full- f gyrokinetic XGC1 [34], predict departures from these scalings for ITER that are attributed to turbulence. For the upstream state specifically, the Stangeby two-point model [35] relates separatrix and target density and temperature through

particle, momentum, and power balance along the field, and has been extended to two-fluid forms [36]. Models of this class may supply the semi-informed separatrix values on which the framework of Chapter 4 depends.

2.3 Near-edge and pedestal physics

The near-edge is where the competing demands of confinement and exhaust are reconciled, and the operating regimes form a spectrum set by that tension [37]. Low-confinement L-mode has a broad λ_q , a weak pressure pedestal, and no edge-localized modes (ELMs). High-confinement H-mode has a narrow λ_q , a tall pedestal, and, in its standard form, Type-I ELMs, repetitive magnetohydrodynamic (MHD) relaxations that expel heat in short, filamentary bursts and threaten plasma-facing components [13]. Between and around these lie a zoo of alternatives, I-mode, quiescent H-mode (QH), the wide-pedestal state, and negative triangularity, that sustain a pedestal while remaining below the Type-I ELM threshold [16, 38]. The distinction that organizes this thesis is between the ELMy H-mode, whose pedestal is capped by MHD peeling–ballooning stability [13, 30] (*MHD-limited*), and every other regime, whose pedestal gradient is instead capped by turbulent transport (*transport-limited*).

Physically the near-edge overlaps core and SOL. The flux surfaces are closed, yet the steep gradients required to bridge the hot core and cool boundary compress the profile scale length L to only a few ρ_i . The ion temperature falls across the same region, so ρ_i shrinks along with L and the ratio does not run away. What erodes is the margin rather than the inequality, because ρ_i/L passes from the $\ll 1$ of the core to a value merely below unity, and the local reduction of Section 2.1 rests on an expansion in precisely that parameter. Fluctuation amplitudes rise to a comparable fraction of the background and their correlation times approach those of parallel transit and profile evolution [2, 19], so the separation fails in time as well as in space. Blobs are born just inside the separatrix as blob–hole pairs, the inward-propagating holes contributing to turbulence spreading and hence nonlocality. Recycled and impurity neutrals

penetrate the pedestal, cooling and diluting it and radiating a substantial fraction of the exhaust power to the main-chamber wall. Their reach may be curtailed by opacity in future devices, but the pedestal is affected regardless [39, 40].

The defining feature of the near-edge is the pedestal itself, which forms when an edge transport barrier (ETB) develops at the L–H transition [41]. The barrier is a narrow layer of strongly sheared $E \times B$ flow that suppresses long wavelength (e.g. ITG) turbulence by shearing eddies apart faster than they grow [42]. It is seeded by driving power through the ion channel. The growing ion pressure gradient sets, through radial force balance, a radial electric field E_r and an associated poloidal rotation, whose shearing rate γ_E opposes the instability growth rates. The transition occurs above a critical input power P_{L-H} that scales with density, field and machine size [43], and $E \times B$ shear is now recognized as a necessary ingredient of a TLP [15]. A second, temporal and spatially smaller scale shear source, zonal flows, mediates the transition dynamics by storing turbulent energy long enough for the mean flow to build [44]. The present scope is restricted to stationary profiles, so zonal flows are not treated further.

Across machines a broadly consistent picture of pedestal transport has emerged. Ion heat and particle transport can drop to near-neoclassical levels in high-confinement regimes, so neoclassical processes are not always a small correction there [37]. How close they come depends strongly on collisionality and on the local gradients, so the comparison is regime-specific rather than a fixed ratio. Electron heat transport, by contrast, exceeds its neoclassical level by one to two orders of magnitude across the pedestal and must be turbulent. Ratios of measured diffusivities to each other and to their neoclassical values are the “fingerprints” by which candidate instabilities are identified [23]. Each mode carries a characteristic signature across the channels, so comparing measured and simulated ratios discriminates between mechanisms that a single flux cannot. The leading candidates for the electron channel are microtearing and electron temperature gradient (ETG) modes, while ion-scale electrostatic modes are usually held down by $E \times B$ shear, although evidence is accumulating that they matter at low collisionality. ETG turbulence survives that shear, so a faithful treatment of the electron channel is unavoidably multiscale [45]. This

decomposition into neoclassical, ion-scale, and electron-scale contributions is the one the case studies of Chapter 5 resolve.

That stiffness weakens near the edge, where $E \times B$ shear quenching and other stabilizing effects soften the flux response. Near-edge gradients are therefore not pinned to their critical values as in the core and can steepen well beyond them. A gradient is held at a threshold only where the mode limiting it is very stiff, as the KBM is in the peeling–ballooning picture of the ELMy H-mode pedestal. Figure 2.1 shows the consequence for any solver that works in gradient space. Across $0.85 \leq r/a \leq 1$ of a representative H-mode the profiles fall smoothly, by factors of 6 in n_e and 19 in T_e , while the inverse scale lengths swing from order unity at the pedestal top to peaks of 39 and 71 within a few centimeters. Piecewise-linear gradient profiles that suffice in the core therefore become inaccurate across the pedestal without an impractically fine grid. The common alternative is a modified hyperbolic tangent (“mtanh”) anchored at the pedestal top and separatrix [13]. It captures pedestal height and width well, but its relationship to the scale lengths that transport solvers evolve is nonlinear and non-trivial rather than direct, and it offers little freedom at the knees of the profile. Reconciling a flexible, edge-aware profile representation with a gradient-based flux matching solver is one of the problems taken up in Chapter 4.




Figure 2.1. Measured near-edge profiles and their gradients for the low-collisionality H-mode of Chapter 5 (DIII-D 179444). (a) n_e , T_e , and T_i , each normalized to its own value at $r/a = 0.85$. (b) The inverse scale lengths a/L_y . The shaded band is the solver domain.

Chapter 3

Filamentary Velocity Scaling and Spin Dynamics in DIII-D

Before turning to profile prediction, this chapter reprints an experimental study from earlier in the doctoral program that motivated the shift in focus toward predictive modeling. Invasive diagnostics, like the reciprocating probe heavily used in the work, will be decreasingly relevant in next-generation devices with hotter, denser plasma. Arcing across the probe head was already common near the LCFS in the DIII-D L-modes of this study and was the reason H-modes were not included in the database. Another crossover between this study and the primary topic of this thesis is uncertainty quantification and parameter sensitivity, which help contextualize the results in both cases.

Filamentary velocity scaling validation and spin dynamics in the DIII-D tokamak

Cite as: Phys. Plasmas **31**, 042515 (2024); doi: 10.1063/5.0195025

Submitted: 29 December 2023 · Accepted: 3 April 2024 ·

Published Online: 23 April 2024



S. C. Molesworth,^{1,a)} J. A. Boedo,¹ C. K. Tsui,² R. Perillo,¹ and D. L. Rudakov¹

AFFILIATIONS

¹University of California, La Jolla, San Diego, California 92093-0417, USA

²Sandia National Laboratories California, P.O. Box 969, Livermore, California 94551-0969, USA

^{a)}Author to whom correspondence should be addressed: smolesworth@ucsd.edu

ABSTRACT

Measured filament velocities in the DIII-D tokamak are compared against theoretical scalings, finding that the latter often represents an upper limit on experimental velocity distributions with most filaments possessing lower velocity. Filament spin from internal $\mathbf{E} \times \mathbf{B}$ drift is experimentally demonstrated to alter filament radial velocity. A critical spin velocity, where filament radial velocity peaks, is observed and corresponds to approximately 5 km/s. This transition is corroborated using a less direct measure of filament spin in the form of a temperature ratio. These techniques are combined to find that the critical spin velocity closely aligns with transport times along and across filaments becoming comparable. The normalized filament size distribution is consistent with the most stable size as dictated by Kelvin–Helmholtz and curvature-driven instabilities. Overall, the findings suggest filament stability and spin alter filamentary transport that may threaten the integrity of first walls in fusion devices.

© 2024 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0195025>

I. INTRODUCTION

Convective transport in the tokamak scrape-off layer (SOL),¹ mediated by coherent, intermittent structures interchangeably called filaments or blobs, can dominate heat and particle fluxes in the far-SOL^{2,3} that can damage the first wall^{4,5} in future fusion devices like SPARC,⁶ ITER,^{7,8} and DEMO.⁹ These filaments originate at or inside the last closed flux surface (LCFS),^{10,11} inherit the nearby plasma characteristics, and advect hot, dense plasma across the SOL.¹² In the SOL, they can be detected as high intermittent peaks in density and temperature.^{2,13} Understanding the physics of filamentary transport is therefore needed to predict plasma temperatures and fluxes at the wall of future fusion devices and possibly controlling or countering them. Although particle and heat transport in H-mode¹⁴ can be drift-dominated, turbulent transport is still involved, especially if mixed regimes with a wider SOL are sought.

Early observations of blob characteristics came from optical diagnostics^{15,16} and probes.¹⁷ Later on, disagreement between fluxes and diffusive particle balances found in simulations¹⁸ and Alcator C-Mod¹⁹ initiated deeper exploration into the phenomenon. Blobs were proposed as a convective mechanism composed of a density peak with an internal, poloidal electric field, E_p , caused by curvature and ∇B_i drifts (B_T is the toroidal magnetic field), and the resulting $E_p \times B_T$

drift carries the filament radially outward.¹² The poloidal axis is oriented perpendicular to B_T and approximately vertical at the midplane, hence the radial force. Measurements soon thereafter verified the theoretical expectations.² Filaments were observed in some capacity over the following years at Alcator C-Mod,^{20,21} NSTX,²² Tore-Supra,²³ and many other toroidal devices.^{24–28} This period also included expansions of blob theory^{29–34} and many computational studies of filaments.^{10,28,35–45}

Given the need for predicting SOL particle and heat fluxes, a filamentary velocity scaling theory was developed,^{16–49} which we summarize in Sec. I and proceed to compare to a large DIII-D database that is described in Sec. II. Analysis of measured filament size, velocity, and collisionality is performed in Sec. III in addition to sought-after validation of the velocity scaling theory. Finally, a summary and conclusions are presented in Sec. IV.

A. Velocity scaling theory

Analytical velocity scaling is derived by considering plasma filaments as a computational circuit model that closes at the divertor floor and defines a filament split into two regions, midplane, and X-point.³² Filaments gain charge through ion polarization drift as described above, and currents flow along and across the magnetic field lines.¹²

Each balance of current closures, dependent on the effective resistivities within a filament, corresponds to a different velocity scaling regime.⁴⁶ In all scenarios, resisting parallel current flow allows a larger internal electric field magnitude and resulting radial $E_p \times B_T$ drift velocity, v_r .

The proxy for total parallel resistivity is an effective collisionality, Λ , that relates parallel connection length to the ion–electron collision mean free path. For a calculation from the divertor target to the outer midplane, Λ is defined as⁴⁶

$$\Lambda = \int_{\text{target}}^{\text{OMP}} \frac{\nu_{ei}}{\Omega_e \rho_s} ds_{\parallel}, \quad (1)$$

where ν_{ei} is the electron–ion collision frequency, Ω_e is electron gyrofrequency, ρ_s is the ion gyroradius at ion sound speed c_s , and $s_{\parallel}$ is the parallel distance to a point along the field line as measured from the target. Filament cross sections near the X-point in diverted plasmas are elongated elliptically because in that region the magnetic poloidal flux expansion, f_x , increases. This shortens the perpendicular current return path for ions and alters the filament velocity.

Filament size also affects cross-field dynamics. Normalizing the filament radius by the scale size that balances inertial, curvature, and sheath terms in the charge conservation equation²⁹ leads to a characteristic filament radius expressed as

$$\hat{a} = a_b \frac{R^{1/5}}{(L_{\parallel} \rho_s^2)^{2/5}}, \quad (2)$$

where R is the major radius, $L_{\parallel}$ is the parallel connection length from the target to outer midplane, and a_b is the filament radius, where $\hat{a}$ is often parameterized^{38,46} to $\Theta = \hat{a}^{5/2}$. The local density gradient drives filament instabilities according to the intrinsic condition $\partial \ln(n)/\partial r' < 0$, where r' is the local filament radial coordinate, and instability growth rates depend on $\hat{a}$. Blobs of $\hat{a} \ll 1$ are more susceptible to the Kelvin–Helmholtz instabilities developing in a timescale $\tau_{KH} \approx (4a_b^3 R) / (\hat{a}^{5/2} c_s \rho_s^2 L_{\parallel})$, where c_s represents the ion sound speed. Curvature-driven instabilities are most prevalent when $\hat{a} \gg 1$ and evolve in the timescale $\tau_{CD} \approx (4a_b^3 R) / (c_s \rho_s^2 L_{\parallel})$.^{31,38,50} These time scales therefore only vary by the dimensionless term $\hat{a}^{5/2} = \Theta$ and are equivalent when $\hat{a} = 1$, which is a transition where blobs are suspected to be most stable.³⁸

Each factor discussed thus far locates a filament within the regions defined by collisionality and size, shown in Fig. 1. The last parameter is normalized velocity

$$\hat{v} = \frac{v_r}{v^*} = v_r \frac{R^{3/5}}{(2L_{\parallel} \rho_s^2)^{1/5} c_s}, \quad (3)$$

where v^* is the ideal filament velocity when $\hat{a} = 1$.⁴⁶ Theory predicts $\hat{v}$ to scale differently depending on a filament's normalized radius, Θ , and effective collisionality, Λ . The four regions with their associated $\hat{v}$ scaling expressions²⁹ are as follows:

- Resistive Ballooning (RB)— Λ exceeding Θ terminates filaments above the X-point leaving ion polarization drift as the primary return current. This regime may be referred to as “inertial-limited.”




FIG. 1. Normalized velocity $\hat{v}$ scaling approximation against normalized collisionality Λ and size $\Theta = \hat{a}^{5/2}$ for expected RB, RX, C_i , and C_s regimes. Reprinted from Tsui et al., Phys. Plasmas **25**, 072506 (2018) with the permission of AIP Publishing.

$$\hat{v} \sim \hat{a}^{1/2}. \quad (4)$$

- Resistive Interchange (RX)—The magnetic field near the X-point squeezes filament cross sections, effectively shortening filaments thereby amplifying cross-field resistivity. Unlike in other regimes, parallel current modulated by Λ regulates filament velocity on a similar order as the perpendicular current with its $\hat{a}$ dependence,

$$\hat{v} \sim \Lambda / \hat{a}^2. \quad (5)$$

- Connected Ideal-Interchange (C_i)—Low-collisional, small filaments extend to the limiter sheath with the primary current closure from ion polarization drift. This regime depends on maximum divertor region magnetic flux expansion, $f_{x,\text{max}}$, expressed inversely in terms of the magnetic fanning parameter, $\varepsilon_x \sim 1/f_{x,\text{max}}$,

$$\hat{v} \sim \varepsilon_x \hat{a}^{1/2}. \quad (6)$$

- Sheath-Connected (C_s)—Low-collisional, large filaments extending to the outer target with current primarily flowing through the sheath,

$$\hat{v} \sim 1 / \hat{a}^2. \quad (7)$$

Mathematical simplifications limit the accuracy of normalized velocity scaling theory, yet there have been many attempts to reproduce fundamental aspects. Experimental comparisons to velocity scaling are available for devices such as NSTX,^{51,52} TCV,^{53,54} ASDEX Upgrade,^{55,56} and the Wendelstein 7-X stellarator.⁵⁷ An intermachine comparison revealed normalized velocities generally fall between the RB and C_s scalings, albeit with significant uncertainties.⁵⁸ 2D^{59–62} and 3D^{57,63} simulations have also reproduced theoretical trends to varying degrees and with the caveat that some inputs are not experimentally determined.

A large spread in experimental velocities at constant blob size and collisionality has been noted.⁵⁴ The spread of velocities found in

experiment may stem from physics excluded by theory. Simulations have shown filament velocity dependencies include warm ions,⁶⁰ shape irregularity,^{57,64} charge smearing,^{34,37} blob-blob interactions,^{62,65,66} Prandtl number,⁶¹ neutral interactions,^{67–70} parallel resistivity,⁷¹ and triangularity,⁷² among other things.⁷³ Furthermore, testing velocity scaling theory is thus justified because it may inform future analytical and computational work.

B. Spinning blobs

Filament spin about the parallel axis was proposed³⁸ as a mechanism to reduce radial velocities by smearing out the ∇B polarization. The original theory applied to sheath-connected blobs³⁴ and showed that, neglecting vorticity injection,²⁹ spin occurs when blob temperature differs from the background plasma.³⁴ It has since been demonstrated that all blobs may spin due to parallel variation in blob characteristics.^{74–76} Blob spin has been observed by probe,^{2,77,78} GPI,⁷⁹ and in simulations.^{75,76,80}

Rotational velocity, v_{rot} , for nonthermalized blobs may be determined using plasma potential, Φ_p , to calculate a local, radial electric field, $E_{r'} = -\nabla\Phi_p \approx \Delta\Phi_p/a_b$ according to $v_{rot} = E_{r'}/B_t$, where r' represents the filament radial coordinate as depicted in Fig. 2.²

The electron temperature, T_e , can be used as a proxy for plasma potential, $\Phi_p = \Phi_f + 2.5k_bT_e$, where the floating potential, Φ_f , is assumed constant and k_b is the Boltzmann constant.³⁴ Assuming a Gaussian profile for filament T_e , the relative difference [Eq. (8a)] between peak filament electron temperature, $T_{e,fil}$, and background electron temperature, $T_{e,bg}$, is predicted to reach a threshold [Eq. (8b)] where polarization dissipation from spin and through parallel transport at the sheath become comparable.^{32,34} Quantitatively, this is when blob angular velocity $\Omega = v_{rot}/a_b$ reaches the sheath current rate⁸¹ $\nu_{sheath} = c_s a_b^2 / (L_{||} \rho_s^2)$, and beyond this level “smearing” gradually

decreases the charge dipole, which consequently reduces the radial velocity,

$$T_{ratio} = \frac{T_{e,fil} - T_{e,bg}}{T_{e,bg}}, \quad (8a)$$

$$T_{ratio,crit} = \frac{a_b^4}{\rho_s^3 L_{||}}. \quad (8b)$$

In cases of too little and too much filament spin, the coherent structure may be altogether lost to instabilities.^{34,58,82} The Kelvin–Helmholtz and curvature-driven instabilities can develop without spin whereas rotational instabilities occur proportional to angular velocity in the timescale $\tau_R \sim \Omega^{-1}$.⁸² These instabilities inhibit filamentary transport when their respective timescales are much smaller than the convective timescale,³⁴ $\tau_C \approx 2a_b/v_r$, that describes the time to move one blob diameter.^{31,82} Between these unstable extremes of the filament spin spectrum are a range where spin stabilizes against curvature-driven instabilities by moving perturbations from “bad” to “good” curvature.^{34,82} Blob coherence becomes more resilient to background shear flows^{34,83} in this range too. Based on this knowledge, filament detection and velocity should peak at some T_{ratio} between zero and $T_{ratio,crit}$.

II. EXPERIMENTAL METHODS

A database was built from existing DIII-D tokamak⁸⁴ deuterium discharges to capture a larger filament scaling parameter space than previous works. The DIII-D tokamak features a 1.7 m major radius, moderate plasma shaping, electron cyclotron heating (ECH), neutral beam injection (NBI), and a substantial diagnostic array.⁸⁵ 74 low-confinement (L-mode), lower single-null diverted discharges spanning the years 2015–2019 are included in this study and summarized in Table I. B_T polarity varies throughout the discharges with a typical $|B_T|$ of 1.5 T. Divertor region gas puffing provided varying divertor plasma conditions, which supports a wide range of collisionality that included measurements by the divertor region Thomson scattering (DTS) system.⁸⁶ An example magnetic geometry generated by the tokamak equilibrium code EFIT⁸⁷ within OMFIT⁸⁸ is shown in Fig. 3 (right) with relevant diagnostic information.

The principal diagnostic for this study was a high-bandwidth, nine-pin reciprocating probe (RCP) located slightly below the mid-plane on the low-field side (LFS) (Fig. 3).⁸⁹ This Langmuir probe array includes a cylindrical, boron nitride probe head; poloidal and toroidal angling to align with typical DIII-D magnetic geometries;




FIG. 2. Toroidal blob spin initiates from a lack of thermalization with the background plasma and can weaken the filament charge dipole.

TABLE I. Database statistics of current I_p , line-averaged density $n_{e,avg}$, instantaneous total power input P_{total} , elongation κ , mean triangularity δ_{mean} , and max divertor region magnetic flux expansion $f_{x,max}$. $|B_T| \approx 1.5$ T for all discharges.

Parameter	Median	Minimum	Maximum
I_p (MA)	0.98	0.86	1.27
$n_{e,avg}$ ($10^{19}/m^3$)	2.07	0.79	7.28
P_{total} (MW)	1.59	0.66	7.10
κ	1.69	1.56	1.82
δ_{mean}	0.39	0.32	0.54
$f_{x,max}$	5.35	3.10	9.18




FIG. 3. Left: Side view of the reciprocating probe (RCP) head. Center: RCP pin layout with relative scale. Right: Example magnetic equilibrium from EFIT for lower single-null discharge with midplane reciprocating probe and DTS locations approximated by the blue arrows and red box.

and 1.25-mm-diameter graphite pins exposed only 2 mm. An actuator plunges the RCP through the SOL, often reaching the LCFS, during a typical 250 ms cycle.

RCP measurements of temperature and density profiles have been benchmarked against other diagnostics at DIII-D including Thomson scattering,^{89–92} reflectometry,⁹¹ and charge-exchange recombination spectroscopy⁹⁰ (CER). Likewise, RCP systems were utilized for SOL characterization at NSTX,^{51,93} C-MOD,⁹⁴ ASDEX,⁹⁵ and JET,⁹⁶ among many others.^{23,77,97–101} Only GPI—a method unavailable at DIII-D in the database time range—can currently measure blob-scale fluctuations in the tokamak SOL comparably to RCP.^{56,102–104} For these reasons, RCP was deemed an appropriate and reliable diagnostic for this work.

Ion saturation I_{sat} and floating potential Φ_f signals are collected at 4 MHz over multiple, spatially separated pins (Fig. 3, center). Floating potential signals are digitally bandpass filtered between 20 and 380 kHz to isolate oscillating filament signals from harmonic electron temperature hardware^{92,105} and bulk plasma oscillations. Electron temperature T_e was collected via the harmonic method^{92,105} cycled at 400 kHz and thermal equilibrium, $T_i = T_e$, is assumed. Previous results^{106–110} demonstrate $T_i \neq T_e$ in the edge, and this is addressed within Appendix B. Electron density was calculated using $n_e = 2I_{sat}[k_b(T_i + T_e)/m_i]^{-1/2}/A_{pr}e$ where A_{pr} is the projected probe tip area, e is one unit charge, n_e is electron density, T_i is the ion temperature, and m_i is the deuterium mass. Ion density n_i relates to n_e by $n_i = n_e/Z_{eff}$, where Z_{eff} is the effective ion charge and taken as a constant $Z_{eff} = 2$.

The fluctuating poloidal electric field, a characteristic of filaments, is determined by $E_p = (\Phi_{f2} - \Phi_{f1})/\delta_p$. This calculation assumes equivalent sheath drops and T_e at each probe pin, and that passing structures are wider than the poloidal inter-pin distance, δ_p , of 5.4 mm. Verification of the results using Φ_p in place of Φ_f is documented in Appendix B out of concern for temperature influences.¹¹¹ Potential errors¹¹² in E_p due to variations in a_b relative to δ_p are also considered in Appendix B. The background radial electric field, E_R , was calculated using the plasma potential gradient, $E_R = -d\Phi_p/dR$, at the probe with 2 ms average values of Φ_p .

Each sample was defined by a time window in which signals were conditionally averaged to calculate typical filament properties.^{113,114} The conditional averaging procedure applied uses I_{sat} as the primary

signal from which events are detected, accumulated, and averaged. Local maxima beyond $+2\sigma$ in I_{sat} were used for identifying 100 μ s event slices centered about each respective maximum. This threshold is below other studies but provides additional data for analysis with negligible consequences.⁵⁴ A 25 μ s minimum between subsequent $+2\sigma$ peaks was determined by inspection to prevent averaging on the same filament repeatedly.

Every sample aligns the I_{sat} peaks of 10 consecutive events and results in a coherent signal nearly isolated from stochastic turbulence. This quantity of events per sample equated to 1.5–2.5 ms, 0.5 cm of stroke, or about two exposed probe tip lengths. More events per sample may reduce conditionally averaged signal noise but also reduces spatial resolution because the time span of each sample increases, and the probe moves throughout that time. Result sensitivity to 50 events per sample is considered in Appendix B. Slices of the remaining fluctuating signals, defined by the 100 μ s surrounding each I_{sat} peak in a sample, were overlaid and averaged, while non-fluctuating parameters like $L_{||}$ were averaged over the sample time window.

The critical filament-specific measurements required for normalization are the radial velocity, poloidal velocity, and duration. Event duration was found by fitting a Gaussian curve to the conditionally averaged I_{sat} signal with rise in signals weighted higher than the return to mean levels to prevent biasing from density wake.⁵⁴ Furthermore, emphasis was placed on filament core sizing, rather than the core plus tail, by extracting the full width half maximum (FWHM) time, τ_{FWHM} . Filament radial velocity was calculated using $v_r = (E_p \times B_t)/B^2$, where E_p here is the conditionally averaged peak poloidal electric field. Radial probe velocities were at most 3 m/s and therefore unnecessary to account for compared to mean v_r as shown in Sec. III.

Although velocity scaling depends explicitly on radial velocity, poloidal velocities are also necessary for estimating filament cross-sectional widths. This was accomplished using time-delay cross correlation of Φ_f signals in both radial and poloidal directions. The peak cross correlation between floating potential tips over the sample's time window provided time delays, τ_p and τ_r , which result in a poloidal velocity $v_{p,cc} = \delta_p/\tau_p$ and radial velocity $v_{r,cc} = \delta_r/\tau_r$ where the radial probe tip distance δ_r is 3 mm. These velocities were corrected on the premise of filaments forming 2D wavefronts at angle $\theta = \tan^{-1}(v_{r,cc}/v_{p,cc})$.¹¹⁵ Measured velocities are therefore $v_{p,2D} = v_{p,cc}(\sin \theta)^2$ and $v_{r,2D} = v_{p,2D}/\tan(\theta)$. Without this correction, a very steep or flat filament trajectory causes the apparent velocity to be nearly infinite. Filament radii were then found using⁵⁴

$$a_b = \frac{\tau_{FWHM}}{2} \sqrt{v_r^2 + v_{p,2D}^2}. \quad (9)$$

Measurement uncertainties were derived to provide context for observations throughout Sec. III. Previously proven^{89,116} quantities, such as B_T , $L_{||}$, and $R - R_{sep}$, have relatively negligible errors and were therefore omitted. Systemic error in v_r was found to be 5% relative to $v_{r,2D}$ meaning a slight overestimation in $E_p \times B_T$ velocity. Random errors on direct measurements were determined case-by-case, whereas indirect quantities without multiple nonlinear terms, e.g., Θ , were evaluated using standard deviations of truncated Taylor expansions according to GUM.¹¹⁷ Covariances were included in these estimates when the correlation coefficient between parameters was larger in magnitude than 0.1. Finally, Monte Carlo simulations were used on highly nonlinear terms like $T_{ratio,crit}$. A table of uncertainties and

derivation details is provided in Appendix A. Error bars corresponding to these uncertainties on the median value of the applicable parameter are shown throughout the remaining text. The size of error bars decreases or increases proportionally with a decrease or increase in said parameter.

Aside from uncertainties, a primary concern about RCP measurement validity stems from the perturbative nature of probes. An oversized probe or one with potential biasing can alter the background plasma locally¹¹⁸ and filter fluctuations.¹¹⁹ Consideration must therefore be taken to minimize and account for probe-related plasma disturbances.^{111,120} The RCP design mitigates perturbations through pin tip sizes of a few ρ_s , which is small relative to blobs. By terminating proud of the probe body, the pin tips also allow data collection before a blob intersects the much larger probe body.

In terms of timescales, the parallel particle transit time, $\tau_{\parallel} \sim L_{\parallel}/c_s \sim 25 \mu\text{s}$, may act as the key metric for probe perturbations. A blob typical of the description in Sec. III A crosses the probe in $\sim 10 \mu\text{s}$, so perturbative effects on filament measurements should be minimal. For background plasma replenishment, we may use the probe diameter relative to mean plasma poloidal flow speed, or $\tau_{p,bg} \sim 20 \mu\text{s}$. This demonstrates that the background plasma near the probe is refreshed slightly faster than the time a particle may spend in the SOL.

III. RESULTS

Nearly 1000 samples from plunges during L-mode operation are included in the ensuing plots. Filament velocity, size, and collisionality are the required inputs for scaling, so these are established first. The scaling regime distribution and adherence to formulas (4)–(7) are then analyzed to determine velocity scaling theory validity in DIII-D. A comparison between filament temperature and velocity is also provided for understanding the influence of filament spin.

Well-documented indicators of filament presence are skewness and excess kurtosis in T_e , E_p , and I_{sat} fluctuation distributions.^{44,99,121–126} The intermittent bursts in signal that accompany filaments are unidirectionally at least two standard deviations away from the mean fluctuation level. Positive values of skewness and excess kurtosis are therefore anticipated in I_{sat} for regions of frequent filament detection. This was the case for 97% of samples in the database.

Density “holes”^{127–129} are another source of positive excess kurtosis, but with negative skewness in I_{sat} .¹²⁹ These are low pressure regions left behind by filament transport that tend to drift radially inward due to reversed charge polarization relative to blobs.^{127,128} Density holes have been omitted from this study by only detecting filaments with $v_r \geq 0$.

A. Filament characteristics

Poloidal and radial velocities determined by cross correlation and $E \times B$ are plotted against $R - R_{sep}$ for an example plunge in Fig. 4. In this example, little discrepancy between the $E_p \times B_T$ and cross correlation methods when $R - R_{sep} > 1 \text{ cm}$ validates the 2D cross correlation technique used. Uncertainties in $v_{r,2D}$ and $v_{p,2D}$ can account for the differences against $E \times B$ at nearly all datapoints. Furthermore, filament poloidal velocity may deviate from background $E \times B$ flows near the LCFS,^{56,80,130,131} so disagreements in poloidal velocity between measurement methods are less concerning here.




FIG. 4. Poloidal (top) and radial (bottom) velocities vs $R - R_{sep}$ for plunge 1 of shot 174 208 using $E \times B$ (blue) and plane wave-corrected 2D cross correlation (green). Agreement between the methods improves with increasing $R - R_{sep}$. Error bars are representative of medians and scale proportionally with v_p or v_r .

Radial velocities were mostly distributed between 0.6 and 1.6 km/s, and the distribution peaked near 1.2 km/s as Fig. 5 indicates. These results agree with previous measurements at DIII-D.^{2,91} Also, worth noting is that an equivalent distribution from TCV nearly matches Fig. 5 despite machine and operating parameter differences.⁵⁴

Velocities and durations were used to calculate filament diameters $2a_b$, the distribution of which is shown in Fig. 6 (left) and spans 0.1–10.2 cm. A simple validation, which can be done at this time, is to check that the filament diameters are reasonable, i.e., larger than the




FIG. 5. Radial velocity distribution from $v_r = E_p \times B_T / B^2$. All samples were in the range of 0.1–3.3 km/s. The error bar represents median velocities and scales proportionally with v_r .




FIG. 6. Filament diameter $2a_b$ (left) and normalized radius $\hat{a}$ (right) distributions are centered near one with long tails toward larger sizes. Red indicates diameters smaller than the maximum probe tip distance, while 14 samples distributed evenly between $5 < 2a_b < 10$ cm have been omitted for readability. Error bars are representative of medians and scale proportionally with a_b or $\hat{a}$.

electrode spacing (0.5 cm) and smaller than the SOL width (~ 5 cm or $\sim 30 \rho_s$). Indeed, 94% of the filament diameters fall within this bracket. Blobs wider than the SOL, of which 11 were measured, need not be erroneous. These can be explained by filaments with elliptical cross sections, as are commonly observed by gas-puff imaging.^{53,56,127,130,132,133} Evidence of ellipticity is also apparent in the probe results, e.g., filaments with diameters greater than 5 cm had an average $v_{p,2D}/v_r$ ratio of 3.5.

The 45 samples satisfying $2a_b < \delta_p$ are included throughout the remaining figures with unique markings because of potentially larger uncertainties in measured quantities. It was decided to include these after thoroughly reviewing the conditionally averaged time traces of I_{sat} , E_p , T_e , and Φ_j cross correlation for many samples in this subset. Distinct peaks in each signal were observed at or near the center of the time windows, as with other samples in the database. The only common difference between each of these samples and the remaining database was a smaller τ_{FWHM} . Once again, elliptical cross sections may offer an explanation. If the major axis of the ellipse does not align with the blob velocity vector, then the I_{sat} probe can pass through the blob along a shortened path. The angle formed between cross-sectional major axis and the blob velocity vector was included in the τ_{FWHM} uncertainty, and these 45 samples may be extreme cases.

Normalized filament sizes, $\hat{a}$, were also compared against the theorized most stable size of unity^{38,50} (Fig. 6, right). Measurements support this with a median $\hat{a}$ of 1.15. Observing blobs smaller or larger than the ideal size is reasonable because instabilities grow on finite time scales. This was confirmed by comparing instability time scales τ_y against convection time scales τ_c ,³⁸ shown in Fig. 7. Here, τ_y may refer to τ_{KH} for $\hat{a} < 1$ or τ_{CD} for $\hat{a} > 1$. For nearly all samples, $\tau_y > \tau_c$, assuming a constant v_r , which suggests the blobs propagate radially more than their respective spatial scales, $2a_b$. Contrary to predictions by stability analysis,³¹ Fig. 7 indicates blobs with $\hat{a} > 1$ generally traveled further relative to their size than blobs with $\hat{a} < 1$. This deviation for larger blobs relates to the a_b^3 dependency in τ_{CD} , whereas

$\tau_{KH} \sim a_b^{1/2}$ for smaller blobs. Using $\tau_y = \tau_{KH}$ for all blobs brings measurements into agreement with the trends from stability analysis and indicates the Kelvin-Helmholtz instability is dominant for blobs throughout this database.

A key parameter in the scaling space is collisionality, Λ , defined by (1). Λ has traditionally been locally approximated^{53–55,134} by




FIG. 7. Instability time scales τ_y for small filaments ($\hat{a} < 1$, $\tau_y = \tau_{KH}$) and large filaments ($\hat{a} > 1$). Large filaments are shown for $\tau_y = \tau_{CD}$ (green $\times$) and $\tau_y = \tau_{KH}$ (orange $+$) and are compared against convective time scales τ_c . Samples further above the dashed line remain coherent longer. Error bars are representative of the median and scale proportionally with τ_y or τ_c .

$\Lambda \sim 1.7 \times 10^{-14} (n_e/T_e^2) L_{\parallel}$ with constant plasma conditions assumed along the magnetic field line, n_e in cm^{-3} , T_e in eV, and $L_{\parallel}$ in cm. Applying this technique with RCP data, which is collected at the mid-plane, will be herein referred to as Λ_{probe} . Updating T_e and n_e along the parallel connection length using divertor Thomson scattering (DTS) measurements on the average flux surface of an individual sample allows approximation of (1) by trapezoidal rule, yielding Λ_{hybrid} . The DTS time and spatial resolution prevent measurements of the same filaments being identified by RCP, so filament collisionality is assumed equal to the background plasma in the divertor region.

Λ_{hybrid} was compared against Λ_{probe} using three existing discharges with varying amounts of lower, outer divertor region gas puffing 120° deg toroidally from the DTS collection locations. A hot, attached divertor leg should yield $\Lambda \ll 1$ and negligible differences in Λ between estimation techniques. Adding cold neutrals to the divertor leg locally reduces T_e and increases n_e temporarily, which results in a dominant divertor region in the Λ_{hybrid} estimate. As shown in Fig. 8, each step-up in gas puffing was accompanied by a larger increase in Λ_{hybrid} relative to Λ_{probe} for the near-SOL region. Λ_{hybrid} estimations during low and high amounts of gas puffing also appear below or above $\Lambda = 1$, respectively, which is where adequately sized filaments transition between connected and resistive regimes (see Fig. 1). Notably, Λ_{probe} rarely exceeds unity independent of divertor gas puffing.

B. Velocity scaling

The range of attainable plasma temperatures and densities in DIII-D enabled analysis of filaments covering a larger parameter space than previous studies. Accordingly, samples were found in all four scaling regions over four orders of magnitude in Λ and Θ as shown in Fig. 9. Measured normalized velocities $\hat{v}_m$ were also in the range expected by velocity scaling theory. A majority of conditionally averaged samples classified as resistive (RB or RX), which may be in part due to the relative sizes of the resistive and connected (C_s or C_i) scaling regimes for the measured ranges of Λ and Θ .




FIG. 8. Normalized collisionality estimates by Λ_{probe} (×) and Λ_{hybrid} (•) vs $R-R_{\text{sep}}$ increasingly vary from low (blue), to intermediate (green), to high (red) amounts of divertor region gas puffing. $\Lambda = 1$ (black, dashed line) is at the transition between scaling regions C_s and RX. The error bar represents the median and scales proportionally with Λ .




FIG. 9. Normalized collisionality Λ vs normalized size $\Theta = \bar{a}^{5/2}$ with color denoting measured normalized velocity $\hat{v}_m$, scaling regions divided by solid lines, and dashed lines relevant to Figs. 11 and 12. The median fanning parameter of $v_x = 0.19$ dictates the division of C_s from C_i . Samples with diameters smaller than the probe tip distance are denoted by ×. Error bars are representative of medians and scale proportionally with Λ or Θ .

Most $\hat{v}_m$ fall below the theoretical normalized velocity $\hat{v}_t$ (notice the subscript “t” for theoretical) per velocity scalings (4)–(7), as seen in Fig. 10 (left). This supports the interpretation that measured filament velocities are often slower than their theoretical counterparts due to mechanisms not treated in the scaling model.⁵⁴ Figure 10 (right) focuses on samples where $\hat{v}_m > \hat{v}_t$ and shows this most often occurred in the RX region where $\hat{v}_t$ is most sensitive to Λ and Θ . The largest values of $\hat{v}_m/\hat{v}_t$ appeared in or near the C_s region where $\hat{v}_t$ is less than anywhere else in the measured parameter space.

Scanning Θ revealed further evidence that scaling theory is predictive of the highest $\hat{v}_m$. First, the scatter in $\hat{v}_m$ vs Θ (Fig. 11, left) nearly matches the area capped by $\hat{v}_t$. Most samples with $\hat{v}_m > \hat{v}_t$ remain within uncertainties of $\hat{v}_t$. The largest deviations occur for $\Theta > 1$ where $\hat{v}_t$ rapidly decreases and depends strongly on Λ . Samples in the RB regime (Fig. 11, right), with the largest $\hat{v}_m$ at a given Θ likewise follow the predicted trend toward higher $\hat{v}_t$, while RB samples were predominately slower than $\hat{v}_t$ overall. In both parameter scans, the minimum $\hat{v}_m$ at each Θ was nearly constant at $\hat{v}_m \approx 0.25$.

Velocity scaling dependence on Λ was similarly studied using a subset of Θ . In Fig. 12, $\hat{v}_m$ rises rapidly in the RX region like $\hat{v}_t$. Varying Θ alters the Λ range where the RX region applies, so measurement uncertainties may be the cause of so many points breaching the $\hat{v}_t$ line. The RX region is also where the highest concentration of samples appeared in Fig. 12. Finally, there are indications that the maximum $\hat{v}_m$ is lower at $\Lambda \ll 1$ than at $\Lambda \gg 1$ despite $\hat{v}_m$ exceeding $\hat{v}_t$ for the few samples in the C_i region, where $\Lambda < 10^{-1}$. This runs counter to the previously discussed measurements in the RB and RX regions. More data are needed in the C_i and C_s regions to better evaluate these details of velocity scaling theory.

C. Spinning blobs

One mechanism known to reduce v_r is spin about the filament longitudinal axis.³⁴ Filaments hotter than the background plasma




FIG. 10. Left: Distribution of measured-to-theoretical normalized velocity ratios, $\hat{v}_m/\hat{v}_t$, with a bias toward $\hat{v}_t$ overprediction. Right: Ratio of $\hat{v}_m/\hat{v}_t$ for samples with $\hat{v}_m > \hat{v}_t$. Most samples in this subset occur in the highly nonlinear RX region of $\hat{v}_t$. Samples with diameters smaller than the probe tip distance are denoted by $\times$. Error bars are representative of the medians and scale proportionally with measured $\hat{v}_m$, Λ , or Θ .

develop internal electric fields that cause them to spin, as shown in Fig. 2. High-frequency T_e measurements by the DIII-D RCP provide filament electron temperature $T_{e,fil}$ against time in addition to that of the background plasma temperature $T_{e,bg}$. For the ensuing analysis, we assume ∇T_e is the primary cause of blob spin without regard for sheath connection. This definition provides no discernment for the direction of spin.

To test the hypothesis outlined in Sec. 1B, the normalization $\hat{T}_{ratio} = T_{ratio}/T_{ratio,crit}$ using Eqs. (8a) and (8b) is first introduced. Examples of conditionally averaged T_e signals in Fig. 13 show $T_{e,fil}$ (central peaks) relative to $T_{e,bg}$ (mean values) for various $\hat{T}_{ratio}$. Wider primary peaks near $0 \mu s$ strongly correlate ($R=0.67$) with a larger $R - R_{sep}$. Time shifts of the primary peaks compared with $0 \mu s$ stem from the magnitude of v_p , whereas the direction of a time shift is caused by the direction of v_p , which results in a blob intersecting the

T_e pin before or after the I_{sat} pin. By the $\hat{T}_{ratio}$ metric, v_r should decrease when $\hat{T}_{ratio}$ either exceeds unity or approaches 0 due to charge smearing and decoherence, respectively (see Sec. 1B).

In Fig. 14, v_r peaks in the range $0.1 < \hat{T}_{ratio} < 1$ and decreases when $\hat{T}_{ratio} > 1$ for the bulk distribution. Like in Sec. III B, the spread in v_r near a given $\hat{T}_{ratio}$ may be attributable to influences from mechanisms like those mentioned in Sec. 1A. More samples were measured with $\hat{T}_{ratio} < 1$, which supports the stabilizing effect of moderate spin. The fit curve in Fig. 14 was robust to changes in the fitting parameters, e.g., polynomial degree, such that v_r repeatedly peaked at $0.50 < \hat{T}_{ratio} < 1$. Finally, $\hat{T}_{ratio}$ typically declines with increasing $R - R_{sep}$, indicating filament rotation is highest near the LCFS before $T_{e,fil}$ rapidly decays.^{29,129,135}

A second measurement of blob spinning effects was made using the plasma potential gradient to estimate $v_{rot} = E_r/B_T$ (see Fig. 2)




FIG. 11. Left: Normalized velocity $\hat{v}_m$ vs normalized size Θ for samples with normalized collisionality $\Lambda \leq 1$ where color indicates Λ magnitude and lines demonstrate theoretical $\hat{v}_t$ at $\Lambda = 1$ or $\Lambda = 2$. Right: $\hat{v}_m$ vs Θ for RB-regime samples where $\hat{v}_t$ scales independent of Λ . Samples with diameters smaller than the probe tip distance are denoted by $\times$. Error bars are representative of medians and scale proportionally with Θ or $\hat{v}_m$.




FIG. 12. Scan of normalized velocity $\hat{v}_m$ vs normalized collisionality Λ for samples with normalized size $0.25 < \Theta < 1.75$. The RX region is the positive slope in $\hat{v}_t$, the RB region is where $\hat{v}_t \approx 1$, and the C_r region is where $\hat{v}_t \approx 0.2$. Error bars are representative of medians and scale proportionally with Λ or $\hat{v}_m$.

where $E_r = (\Phi_{p,fil} - \Phi_{p,bg})_{max}/a_b$. It is evident from this expression for v_{rot} that blobs can spin in opposite directions when B_T polarity is reversed although the following results utilize $|B_T|$. Directly comparing the spin estimation techniques, Fig. 15 (left), exposes a power-law correspondence between v_{rot} and $\hat{T}_{ratio}$. A similar peak in v_r to that of




FIG. 13. Filament T_e time traces with varying normalized temperature ratio $\hat{T}_{ratio}$ indicated by color. The error bar is representative of the median T_e and scales proportionally with T_e .




FIG. 14. Filament radial velocity v_r vs normalized temperature ratio $\hat{T}_{ratio}$ with $R - R_{sep}$ given by symbol color. The dashed line is a fourth-degree polynomial fit to the medians of 62 bins along $\hat{T}_{ratio}$. The fit curve peaks in v_r at $\hat{T}_{ratio} = 0.85$. Error bars represent median values and scale proportionally with $\hat{T}_{ratio}$ or v_r .

Fig. 14 occurs near $v_{rot} \sim 5$ km/s in Fig. 15 (right). Previously reported values^{2,7,79} of v_{rot} also fall in the measured range reported here. At both ends of the v_{rot} range of Fig. 15 (right) were decreases in v_r , and $R - R_{sep}$ was inversely correlated with more spin.

The v_r dependence on blob spin described so far can be framed as a competition between the time scales τ_γ , τ_{sheath} , and τ_{rot} . To test this, an ideal filament was constructed from typical database measurements of $T_{e,fil}$, Φ_f , a_b , c_s , $L_{||}$, and B_T . Like the results from Sec. III A and Fig. 7, the Kelvin–Helmholtz is the dominate instability ($\tau_{KH} < \tau_{CD}$) for these conditions, so $\tau_\gamma = \tau_{KH}$. The rotation timescale $\tau_{rot} = \pi/\Omega = \pi a_b/v_{rot}$ represents the time for a particle to be spun to the opposite side of a blob cross section. This “test” blob has a radial velocity estimated as $v_{r,test} = \frac{E_{p,test}}{B_T} = 3590$ m/s where $E_{p,test} \sim 2\Phi_{f,test}/2a_b$. Next, two scalars that act on $v_{r,test}$ are defined as

$$\alpha_\gamma = \frac{\tau_\gamma^{-1}}{\tau_\gamma^{-1} + \tau_{rot}^{-1}}, \quad (10a)$$

$$\alpha_{\nabla B} = \frac{\tau_{sheath}^{-1}}{\tau_{sheath}^{-1} + \tau_{rot}^{-1}}. \quad (10b)$$

τ_{sheath} is used as a proxy for the ∇B charge polarization timescale although this is oversimplified. A linear combination of (10a) and (10b) applied to $v_{r,test}$ forms a v_r that matches the trends observed in Fig. 15 (right)

$$v_r = \alpha_{eff} v_{r,test} = (\alpha_{\nabla B} - \alpha_\gamma) v_{r,test}. \quad (11)$$

Recalling Sec. I B, charge smearing from spin begins to dominate ∇B polarization near $\hat{T}_{ratio} = 1$ because the sheath current rate ν_{sheath} becomes slower than the blob spin angular frequency. Relating time scales for these processes can contextualize the observations from Figs. 14 and 15 (right). The time scales for rotation τ_{rot} and sheath dissipation $\tau_{sheath} = 1/\nu_{sheath}$, plotted in Fig. 16, show that $\hat{T}_{ratio}$ is slightly larger than unity when $\tau_{rot} = \tau_{sheath}$. At $\hat{T}_{ratio} = 1$, $\tau_{rot} \approx 4 \mu s$ and corresponds to $v_{rot} \approx 2$ –8 km/s, depending on a_b of a given sample. This range of v_{rot} includes where v_r peaks in Fig. 15 (right).




FIG. 15. Left: Spin tangential velocity v_{rot} is plotted against $\hat{T}_{ratio}$ and a power-law relation (black line) is evident between filament spin estimation techniques. Right: Filament radial velocity v_r vs v_{rot} with $R - R_{sep}$ given by symbol color and Eq. (11) shown by the dashed black line that caps most samples. Error bars represent median values and scale proportionally with $\hat{T}_{ratio}$, v_{rot} , or v_r .

IV. CONCLUSIONS

Evidence comparing blob radial velocity scaling theory⁴⁶ against experiment was sought and achieved by this work through a combination of techniques applied to DIII-D tokamak probe data. This analysis included a large plasma parameter space achieved by extracting data from L-mode, lower single-null discharges with operating conditions. Samples were identified in all four scaling areas, but most were in the resistive RB and RX regions.

Internal, poloidal electric fields and two-dimensional cross correlation were calculated on conditionally averaged filaments using outer

midplane reciprocating probe data and later became inputs for velocity scaling validation. Filaments traveled radially outward from 0.1 to 3.3 km/s with diameters spanning 0.2–10.2 cm. Only 6% of samples were larger in diameter than the SOL or smaller in diameter than the probe tip gaps, and these samples carry additional uncertainty. Conditionally averaged signals of these outlier filaments appeared comparable to more typical ones, and the sizes may be attributable to noncircular cross sections.

Normalized filament radial velocities were mostly distributed between 40% and 120% of the theoretical values, i.e., the maximum measured velocities tended to follow the simplistic velocity scaling laws. This trend supports the interpretation that the theory predicts the maximum velocity for an ideal filament, while the measured filaments have their velocities reduced by other mechanisms, which are not treated in the model. Overall, 56% of samples measured slower than the velocity scaling predicts, and the remaining 44% of samples were most often in the RX region that has the most nonlinear velocity scaling.

One mechanism anticipated to reduce filament radial velocities is spin about the filament longitudinal axis. This was tested on nonthermalized blobs relative to the background plasma using internal T_e and Φ_p peaks. Spin was first represented in the form of a temperature ratio. Radial velocity was shown to increase up to a predicted critical temperature ratio where charge smearing inside a blob outpaces sheath dissipation. Filament velocities decreased monotonically past the critical temperature ratio and suggest a slowing effect in agreement with theory.³⁴ The v_r dependence on spin was also observed with a more direct measurement using plasma potential gradients. Time scale estimates were utilized to explain the observed velocity trends.^{74,83}

Blobs were mostly observed within measurement uncertainties of the theorized most stable size.³⁸ Estimated instability time scales were longer than convection time scales for nearly all samples, which supports non-local transport intrinsic to convection. Contrary to theory,³⁸ the Kelvin–Helmholtz instability appeared dominate for filaments of all sizes.³⁸ The cause of this was not investigated beyond a scaling argument.

Demonstrating blob velocity scaling validity, stability conditions, and spin dynamics improves our understanding of intermittent




FIG. 16. Filaments with normalized temperature ratio $\hat{T}_{ratio} \gg 1$ have a larger sheath current timescale τ_{sheath} than spin timescale τ_{rot} and appear below the dashed line that represents $\tau_{sheath} = \tau_{rot}$. The color scale indicates $\log_{10}(\hat{T}_{ratio})$ while blue has been used to denote $0.5 < \hat{T}_{ratio} < 1.5$. Error bars represent median values and scale proportionally with τ_{rot} or τ_{sheath} .

filaments that dominate particle and heat fluxes across the tokamak SOL. Making further progress requires continuing the effort to quantify experimental and theoretical uncertainties, while reducing them where feasible. Likewise, assumptions such as $T_i = T_e$ must be limited for an accurate assessment of blobs. Velocity scaling should also be tested in confinement regimes relevant to future power plants, e.g., H-mode. Preliminary analysis of DIII-D H-mode discharges suggests similar results overall, but a larger sample size is needed. Finally, capturing the reported blob velocity spread analytically by adding physics, like blob spin, appears necessary to bring predictive capabilities within reach. Doing so may present a path to an “enhanced convective” SOL¹¹⁰ that distributes divertor heat fluxes over a larger area than witnessed in existing devices, while preventing interactions between filaments and the main chamber first wall.

ACKNOWLEDGMENTS

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award Nos. DE-FC02-04ER54698, DE-FG02-07ER54917, and DE-SC0010529.

We gratefully acknowledge the experimental operations of the DIII-D team.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Steven Molesworth: Conceptualization (lead); Data curation (lead); Formal analysis (lead); Methodology (equal); Software (lead); Visualization (lead); Writing – original draft (lead). **Jose A. Boedo:** Conceptualization (lead); Data curation (supporting); Formal analysis (supporting); Funding acquisition (lead); Investigation (lead); Methodology (equal); Project administration (lead); Resources (lead); Software (supporting); Supervision (lead); Writing – review & editing (lead). **Cedric K. Tsui:** Conceptualization (supporting); Formal analysis (supporting); Investigation (supporting); Methodology (equal); Software (supporting); Supervision (supporting); Writing – review & editing (supporting). **Renato Perillo:** Investigation (supporting); Writing – review & editing (supporting). **Dmitry Rudakov:** Investigation (supporting); Writing – review – editing (supporting).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

APPENDIX A: UNCERTAINTY QUANTIFICATION

Table II outlines all parameters, in no specific order, with assigned uncertainties from this study. Each uncertainty is presented as a standard deviation σ_f normalized to the expected value of parameter f .

For generic parameters “x” and “y,” the parameter-specific standard deviations are σ_x and σ_y , their correlation is ρ_{xy} , and their

TABLE II. Summary of measurement uncertainties. All two-digit percentages are rounded to the nearest 5%.

Parameter	Description	σ_f (%)
Z_{eff}	Effective atomic number inside filament	50
m_i	Ion mass	50
$T_{e,probe}$	Probe T_e (e.g., $T_{e,fil}$, $T_{e,bg}$)	30
$n_{e,probe}$	Probe n_e	30
τ_{FWHM}	Filament signal duration	5
v_r	Filament $E_p \times B_T$ radial velocity	30
$v_{p,2D}$	Filament cross-correlated poloidal velocity	30
a_b	Filament radius	25
Φ_f	Probe floating potential	20
Φ_p	Probe plasma potential	25
E_p	Poloidal electric field at probe	30
E_r	Radial electric field at probe	35
c_s	Ion sound speed	30
ρ_s	Ion Larmor radius at c_s	30
$T_{e,TS}$	DTS T_e	25
$n_{e,TS}$	DTS n_e	15
$\hat{a}$	Normalized filament radius	30
Θ	$\hat{a}^{5/2}$	75
$\hat{v}$	Normalized filament velocity	35
Λ_{probe}	Filament collisionality estimate by probe	55
Λ_{hybrid}	Filament collisionality approximation by trapezoidal rule	25
τ_{KH}	Kelvin–Helmholtz instability time scale	20
τ_{CD}	Curvature-driven instability time scale	15
τ_C	Filament convection time scale	35
τ_{sheath}	Sheath dissipation time scale	50
τ_{rot}	Filament spin time scale	60
T_{ratio}	Filament T_e ratio, (8a)	25
$T_{ratio,crit}$	Critical filament T_e ratio, (8b)	30
$\tilde{T}_{ratio}$	Normalized filament T_e ratio, (8a)/(8b)	40
v_{rot}	Filament rotational velocity	45

covariance is $\sigma_{xy} = \rho_{xy}\sigma_x\sigma_y$. Correlations are omitted when $|\rho_{xy}| < 0.1$. Finally, the “bar over,” like $\bar{x}$, denotes an average across the database.

- Z_{eff} : Randomly selected manual checks of Z_{eff} near the LCFS using charge–exchange recombination spectroscopy revealed typical values of 1–3.
- m_i : Assumed equal in uncertainty to Z_{eff} .
- $T_{e,probe}$ and $n_{e,probe}$: Standard deviations of T_e and n_e across each sample’s time span were normalized by the sample mean T_e and n_e , respectively, and the mean from this dataset was used. $T_{e,probe}$ and $n_{e,probe}$ were found to nearly overlay with upstream, SOL Thomson scattering measurements for $R - R_{sep} > 1$ cm. For $R - R_{sep} \leq 1$ cm, $T_{e,probe}$ and $n_{e,probe}$ were still within the determined uncertainty of Thomson scattering values.

- τ_{FWHM} : Drawing from synthetic probe error estimate techniques,¹¹² a Monte Carlo simulation of I_{sat} pin path through an elliptic blob cross section was completed. A constant effective elliptic radius r_{eff} was set to 0.5 cm and the ellipse area equaled πr_{eff}^2 . Angle θ was assigned a normal distribution centered at 0° and with a 15° standard deviation, whereas the velocity ratio v_r/v_p distribution was representative of the database. One standard deviation of chord length error (relative to the ellipse major radius) was taken as the time duration uncertainty, $\sigma_{\tau_{FWHM}}$. Although less than the 10% error previously reported for circular cross sections,¹¹² this may be expected in elongated structures.
- v_r : Median 2D cross-correlated radial velocity $v_{r,2D}$ error relative to radial $E_p \times B_T$ velocity v_r for filaments. The value found is comparable to previous experimental reports.^{54,130,135}
- $v_{p,2D}$: Uncertainty assumed equal to that of v_r due to interdependence of $v_{p,2D}$ and $v_{r,2D}$. Comparing this against v_p from $E_r \times B_T$ is informative but incomplete, especially near the LCFS, due to effects by diamagnetic drifts¹³¹ and blob path irregularities shortly after birth.¹³⁰
- a_b : $\sigma_{a_b}^2 = \sigma_{\tau_{FWHM}}^2 + \sigma_{|v|}^2$ where $\sigma_{|v|}^2 = \sigma_{v_r}^2 + \sigma_{v_{p,2D}}^2 + 2\sigma_{v_r v_{p,2D}}$ and $\rho_{v_r v_{p,2D}} = 0.13$.
- Φ_f : The 20% value was assumed and later found to nearly reproduce σ_{v_r} according to $\sigma_{v_r}^2 = 2\sigma_{\Phi_f}^2$.
- Φ_p : $\sigma_{\Phi_p}^2 = \sigma_{\Phi_f}^2 + \sigma_{T_e}^2 + 2\sigma_{\Phi_f T_e}$ where $\rho_{\Phi_f T_e} = -0.5$.
- E_p : $\sigma_{E_p}^2 = 2\sigma_{\Phi_f}^2$.
- E_r : $\sigma_{E_r}^2 = 2\sigma_{\Phi_p}^2$.
- c_s : $\sigma_{c_s}^2 = 0.25(\sigma_{T_e}^2 + \sigma_{Z_{eff}}^2)$.
- ρ_s : $\sigma_{\rho_s}^2 = 0.25(\sigma_{T_e}^2 + \sigma_{Z_{eff}}^2)$.
- $T_{e,TS}$ and $n_{e,TS}$: Errors were extracted from the DTS system for the shots and channels used in the construction of Fig. 8, which covers low- to high-collisionality conditions. A description of how DIII-D DTS errors are approximated is beyond the scope of this paper. It was observed that $\sigma_{T_{e,TS}}$ and $\sigma_{n_{e,TS}}$ respond inversely to changes in Λ , so a fit was made for each of these uncertainties. The mean values are shown in Table II.
- $\hat{a}$: $\sigma_{\hat{a}}^2 = \sigma_{a_b}^2 + \sigma_{\rho_s^{0.8}}^2 - 2\sigma_{a_b \rho_s^{0.8}}$ where $\rho_{a_b \rho_s^{0.8}} = 0.2$.
- θ : $\sigma_{\theta} = 2.5\sigma_{\hat{a}}$.

- $\hat{v}$: $\sigma_{\hat{v}}^2 = \sigma_{v_r}^2 + \sigma_{1/(\rho_s^{0.4} c_s)}^2 - 2\sigma_{v_r/(\rho_s^{0.4} c_s)}$, where $\rho_{v_r/(\rho_s^{0.4} c_s)} = 0.56$, $\sigma_{1/(\rho_s^{0.4} c_s)}^2 = \sigma_{c_s}^2 + \sigma_{\rho_s^{0.4}}^2 - 2\sigma_{\rho_s^{0.4} c_s}$, and $\rho_{\rho_s^{0.4} c_s} = 0.95$.
- Λ_{probe} : $\sigma_{\Lambda_{probe}}^2 = \sigma_{n_{e,probe}}^2 + \sigma_{T_{e,probe}^2}^2 - 2\sigma_{n_{e,probe} T_{e,probe}^2}$, where $\rho_{n_{e,probe} T_{e,probe}^2} = 0.7$.
- Λ_{hybrid} : The refined collisionality Λ_{hybrid} relies on the trapezoidal rule for estimating the integral in (1). The uncertainty in Λ_{hybrid} is therefore a summation with each component weighted by its fraction of the parallel connection length. Four Thomson scattering lines dispersed in the 30% of connection leg nearest the divertor in addition to the single probe location yields $n = 5$. Each datapoint had an uncertainty determined by n_e/T_e for the applicable diagnostic (assuming $\sigma_{L_{||}} = 0$). The distance between datapoints is $\Delta L_{||,i+1} = L_{||,i+1} - L_{||,i}$ and the total uncertainty is then

$$\sigma_{\Lambda_{hybrid}}^2 = 0.25 \sum_{i=1}^{n-1} \left[\left(\left(\frac{n_e}{T_e^2} L_{||} \right)_i^2 + \left(\frac{n_e}{T_e^2} L_{||} \right)_{i+1}^2 \right) + 2L_{||,i} L_{||,i+1} \sigma_{probe-TS} \right] \Delta L_{||,i}^2,$$

where $\sigma_{probe-TS} = 0.13$.

- τ_{KH} : $\sigma_{\tau_{KH}}$ determined by Monte Carlo method with $\hat{a}$, c_s , ρ_s , R , and $L_{||}$ given unique distributions representative of the database.
- τ_{CD} : $\sigma_{\tau_{CD}}$ determined by Monte Carlo method with a_b , c_s , ρ_s , R , and $L_{||}$ given unique distributions representative of the database.
- τ_c : $\sigma_{\tau_c}^2 = \sigma_{a_b}^2 + \sigma_{v_r}^2$.
- τ_{sheath} : $\sigma_{\tau_{sheath}}$ determined by Monte Carlo method with a_b , c_s , ρ_s , and $L_{||}$ given unique distributions representative of the database.
- τ_{rot} : $\sigma_{\tau_{rot}}^2 = \sigma_{v_{rot}}^2 + \sigma_{a_b}^2 - 2\sigma_{v_{rot} a_b}$ where $\rho_{v_{rot} a_b} = -0.37$.
- T_{ratio} : $\sigma_{T_{ratio}}^2 = 2\sigma_{T_{e,probe}}^2 - 2\sigma_{T_{e,fil} T_{e,bg}}$ where $\rho_{T_{e,fil} T_{e,bg}} = 0.67$ and $\sigma_{T_{e,probe}} = \sigma_{T_{e,fil}} = \sigma_{T_{e,bg}}$ but $T_{e,fil} \neq T_{e,bg}$.
- $T_{ratio,crit}$: $\sigma_{T_{ratio,crit}}$ determined by Monte Carlo method with a_b , ρ_s , and $L_{||}$ given unique distributions representative of the database.
- $\hat{T}_{ratio}$: $\sigma_{\hat{T}_{ratio}}^2 = \sigma_{T_{ratio}}^2 + \sigma_{T_{crit}}^2$.
- v_{rot} : $\sigma_{v_{rot}}^2 = \sigma_{a_b}^2 + \sigma_{T_{e,probe}}^2 - 2\sigma_{a_b T_{e,probe}}$ where $\rho_{a_b T_{e,probe}} = -0.12$.




FIG. 17. Left: Grid of flux surface ψ against normalized collisionality Λ with color indicating T_i/T_e . Right: Filament radial velocity v_r vs normalized temperature ratio $\hat{T}_{ratio}$ with $R - R_{sep}$ given by symbol color and all values adjusted for $T_i > T_e$. Error bars represent median values and scale proportionally with $\hat{T}_{ratio}$ or v_r .




FIG. 18. Difference between measured and theoretical normalized velocity $\hat{v}$ at each sample's measured normalized collisionality Λ and normalized size Θ when $v_r = E_p \times B_T/B^2$ with $E_p = -\nabla\phi_p$. More samples exceed the velocity scalings than in Fig. 12. Samples with diameters smaller than the probe tip distance are denoted by $\times$. Error bars are representative of median values and scale proportionally with $\hat{v}$ or Θ .

APPENDIX B: SENSITIVITY ANALYSIS

Velocity scaling and blob spin analysis was repeated five additional times with one deviation in assumptions or conditions (see Sec. II) per iteration. Described below are the methods by which this was achieved and a discussion of the affected results. Anything not discussed may be assumed to agree with the observations of Sec. III.

The first assumption checked was $T_i = T_e$. Previous studies of the SOL have reported T_i/T_e to vary directly with ψ and inversely to Λ .^{106–110} The grid shown in Fig. 17 (left) was constructed with these results projected onto ψ and Λ ranges representative of this work. Being a sensitivity check, the exact values of T_i/T_e are not critical. T_i was calculated for each sample by interpolating T_i/T_e with the sample-specific ψ and Λ , then multiplying the value by that sample's T_e . As visible in Fig. 17 (right), T_{ratio} is shifted to higher values from Fig. 14 and v_r peaks at $T_{ratio} \sim 5$.

Next, the number of events per sample was set to 50. Doing so greatly reduces spatial resolution by increasing the time span per sample but also reduces signal noise after conditional averaging. A major concern with a reduction of spatial resolution is the drastic change in blob characteristics with respect to SOL location.⁹¹ Another detriment to this change is an 80% reduction in sample count. Despite this, trends and observations were still evident albeit with less confidence due to data sparsity.

Other methods to determine v_r were also checked. Substituting the 2D cross-correlated $v_{r,2D}$ in place of v_r from $E_p \times B_T$ resulted in hardly any differences from the primary analysis. Alternatively, replacing floating potential Φ_f with plasma potential Φ_p in E_p may change v_r . If $T_i = T_e$ and blobs pass centered over the probe, then this should have minimal effect because the temperature components of Φ_p are canceled. We find that blob velocities and, subsequently sizes, have larger tails of high values using this technique. This is expected because the T_e pin spacing, δ_{T_e} , is larger than δ_p , so blobs with $2a_b < \delta_{T_e}$ will retain a T_e term in $E_p = -\nabla\phi_p$. Figure 18

can be contrasted against Fig. 11 to find a shift in v_r and a weaker adherence to scaling.

The final sensitivity considered relates to the fixed probe pin gap, δ_p . Approximations of the error caused by blob diameter relative to the pin gap can be made by extrapolating synthetic probe simulation results.¹¹² The relation $e_{v_r} = -0.34\log(2a_b/\delta_p)$, where e_{v_r} is the assumed radial velocity error, determines the degree of overestimation ($e_{v_r} < 0$) or underestimation ($e_{v_r} > 0$) in v_r by




FIG. 19. Distribution of measured-to-theoretical normalized velocity ratios, $\hat{v}_m/\hat{v}_t$, for all trials in the sensitivity analysis. A bias toward $\hat{v}_t$ overprediction remains evident despite varying assumptions and conditions. The error bar is representative of the median and scales proportionally with measured $\hat{v}_m$.

measured $E_p \times B_T$. Setting the radial velocity to $v_r(1 + e_v)$ shifted the v_r distribution median by -30% , moved nearly all samples below the theoretical $\hat{v}$ cap, and improved the visibility of scaling trends in parameter scans like Fig. 11. Combining these results with the previously mentioned sensitivity checks leads to the distribution of $\hat{v}_m/\hat{v}_t$ ratios in Fig. 19. As with Fig. 10, there remains a skew toward $\hat{v}_m < \hat{v}_t$.

REFERENCES

- ¹P. C. Stangeby, *The Plasma Boundary of Magnetic Fusion Devices* (CRC Press, 2000).
- ²J. A. Boedo, D. Rudakov, R. Moyer, S. Krashenninnikov, D. Whyte, G. McKee, G. Tynan, M. Schaffer, P. Stangeby, P. West, S. Allen, T. Evans, R. Fonck, E. Hollmann, A. Leonard, A. Mahdavi, G. Porter, M. Tillack, and G. Antar, "Transport by intermittent convection in the boundary of the DIII-D tokamak," *Phys. Plasmas* **8**(11), 4826–4833 (2001).
- ³D. L. Rudakov, J. A. Boedo, R. A. Moyer, S. Krashenninnikov, A. W. Leonard, M. A. Mahdavi, G. R. McKee, G. D. Porter, P. C. Stangeby, J. G. Watkins, W. P. West, D. G. Whyte, and G. Antar, "Fluctuation-driven transport in the DIII-D boundary," *Plasma Phys. Control. Fusion* **44**, 717 (2002).
- ⁴G. Federici, C. H. Skinner, J. N. Brooks, J. P. Coad, C. Grisolia, A. A. Haasz, A. Hassanein, V. Philipps, C. S. Pitcher, J. Roth, W. R. Wampler, and D. G. Whyte, "Plasma-material interactions in current tokamaks and their implications for next step fusion reactors," *Nucl. Fusion* **41**(12), 1967–2137 (2001).
- ⁵J. N. Brooks, "Modeling of sputtering erosion/redeposition—status and implications for fusion design," *Fusion Eng. Des.* **60**, 515–526 (2002).
- ⁶S. B. Ballinger, A. Q. Kuang, M. V. Umansky, D. Brunner, J. M. Canik, M. Greenwald, J. D. Lore, B. LaBombard, J. L. Terry, M. Wigram, and SPARC Team, "Simulation of the SPARC plasma boundary with the UEDGE code," *Nucl. Fusion* **61**(8), 086014 (2021).
- ⁷A. Loarte, B. Lipschultz, A. S. Kukushkin, G. F. Matthews, P. C. Stangeby, N. Asakura, G. F. Counsell, G. Federici, A. Kallenbach, K. Krieger, A. Mahdavi, V. Philipps, D. Reiter, J. Roth, J. Strachan, D. Whyte, R. Doerner, T. Eich, W. Fundamenski, A. Herrmann, M. Fenstermacher, P. Ghendrih, M. Groth, A. Kirschner, S. Konoshima, B. Labombard, P. Lang, A. W. Leonard, P. Monier-Garbet, R. Neu, H. Pacher, B. Pegourie, R. A. Pitts, S. Takamura, J. Terry, and E. Tsitrone, "Chapter 4: Power and particle control," *Nucl. Fusion* **47**(6), S203 (2007).
- ⁸J. N. Brooks, J. P. Allain, and T. D. Rognien, "Erosion/redeposition analysis of the ITER first wall with convective and non-convective plasma transport," *Phys. Plasmas* **13**(12), 122502 (2006).
- ⁹G. Federici, W. Biel, M. R. Gilbert, R. Kemp, N. Taylor, and R. Wenninger, "European DEMO design strategy and consequences for materials," *Nucl. Fusion* **57**(9), 092002 (2017).
- ¹⁰N. Bisai, A. Das, S. Deshpande, R. Jha, P. Kaw, A. Sen, and R. Singh, "Edge and scrape-off layer tokamak plasma turbulence simulation using two-field fluid model," *Phys. Plasmas* **12**(7), 027520 (2005).
- ¹¹N. Bisai, S. Banerjee, S. J. Zweben, and A. Sen, "Experimental validation of universal plasma blob formation mechanism," *Nucl. Fusion* **62**(2), 026027 (2022).
- ¹²S. I. Krashenninnikov, "On scrape off layer plasma transport," *Phys. Lett. A* **283**, 368–370 (2001).
- ¹³G. Decristoforo, A. Theodorsen, J. Omatani, T. Nicholas, and O. E. Garcia, "Numerical turbulence simulations of intermittent fluctuations in the scrape-off layer of magnetized plasmas," *Phys. Plasmas* **28**(7), 072301 (2021).
- ¹⁴ASDEX Team, "The H-Mode of ASDEX," *Nucl. Fusion* **29**(11), 1959–2040 (1989).
- ¹⁵D. H. J. Goodall, "High speed cine film studies of plasma behaviour and plasma surface interactions in tokamaks," *J. Nucl. Mater.* **111–112**, 11–22 (1982).
- ¹⁶S. J. Zweben, J. McChesney, and R. W. Gould, "Optical imaging of edge turbulence in the caltech tokamak," *Nucl. Fusion* **23**(6), 825–830 (1983).
- ¹⁷S. J. Zweben, "Search for coherent structure within tokamak plasma turbulence," *Phys. Fluids* **28**(3), 974–982 (1985).
- ¹⁸Y. Sarazin and P. Ghendrih, "Intermittent particle transport in two-dimensional edge turbulence," *Phys. Plasmas* **5**(12), 4214–4228 (1998).
- ¹⁹B. LaBombard, "An interpretation of fluctuation induced transport derived from electrostatic probe measurements," *Phys. Plasmas* **9**(4), 1300 (2002).
- ²⁰J. L. Terry, S. J. Zweben, K. Hallatschek, B. LaBombard, R. J. Maqueda, B. Bai, C. J. Boswell, M. Greenwald, D. Kopon, W. M. Nevins, C. S. Pitcher, B. N. Rogers, D. P. Stotler, and X. Q. Xu, "Observations of the turbulence in the scrape-off-layer of Alcator C-Mod and comparisons with simulation," *Phys. Plasmas* **10**(5), 1739–1747 (2003).
- ²¹J. L. Terry, N. P. Basse, I. Cziegler, M. Greenwald, O. Grulke, B. LaBombard, S. J. Zweben, E. M. Edlund, J. W. Hughes, L. Lin, Y. Lin, M. Porkolab, M. Sampsell, B. Veto, and S. J. Wukitch, "Transport phenomena in the edge of Alcator C-Mod plasmas," *Nucl. Fusion* **45**(11), 1321–1327 (2005).
- ²²S. J. Zweben, R. J. Maqueda, D. P. Stotler, A. Keesee, J. Boedo, C. E. Bush, S. M. Kaye, B. LeBlanc, J. L. Lowrance, V. J. Mastrocola, R. Maingi, N. Nishino, G. Renda, D. W. Swain, J. B. Wilgen, and NSTX Team, "High-speed imaging of edge turbulence in NSTX," *Nucl. Fusion* **44**(1), 134–153 (2004).
- ²³P. Devynck, P. Ghendrih, and Y. Sarazin, "The origin of the long time correlations of the density fluctuations in the scrape-off layer of the Tore Supra Tokamak," *Phys. Plasmas* **12**(5), 050702 (2005).
- ²⁴G. S. Kirnev, V. P. Budaev, S. A. Grashin, E. V. Gerasimov, and L. N. Khimchenko, "Intermittent transport in the plasma periphery of the T-10 tokamak," *Plasma Phys. Control Fusion* **46**(4), 621–637 (2004).
- ²⁵Y. H. Xu, S. Jachmich, R. R. Weynants, and TEXTOR Team, "On the properties of turbulence intermittency in the boundary of the TEXTOR tokamak," *Plasma Phys. Control Fusion* **47**(10), 1841–1855 (2005).
- ²⁶G. Y. Antar, G. Counsell, J.-W. Ahn, Y. Yang, M. Price, A. Tabasso, and A. Kirk, "The poloidal distribution of turbulent fluctuations in the mega-Ampere spherical Tokamak," *Phys. Plasmas* **12**(3), 032506 (2005).
- ²⁷G. Y. Antar, M. Tsalas, E. Wolfrum, and V. Rohde, "Turbulence during H- and I-mode plasmas in the scrape-off layer of the ASDEX Upgrade tokamak," *Plasma Phys. Control Fusion* **50**(9), 095012 (2008).
- ²⁸O. E. Garcia, R. A. Pitts, J. Horacek, A. H. Nielsen, W. Fundamenski, J. P. Graves, V. Naulin, and J. J. Rasmussen, "Turbulent transport in the TCV SOL," *J. Nucl. Mater.* **363–365**(1–3), 575–580 (2007).
- ²⁹D. A. D'Ippolito, J. R. Myra, and S. I. Krashenninnikov, "Cross-field blob transport in tokamak scrape-off-layer plasmas," *Phys. Plasmas* **9**(1), 222 (2002).
- ³⁰D. A. D'Ippolito, J. R. Myra, S. I. Krashenninnikov, G. Q. Yu, and A. Y. Pigarov, "Blob transport in the Tokamak scrape-off-layer," *Contrib. Plasma Phys.* **44**(1–3), 205–216 (2004).
- ³¹D. A. D'Ippolito and J. R. Myra, "Blob stability and transport in the scrape-off-layer," *Phys. Plasmas* **10**(10), 4029–4039 (2003).
- ³²S. I. Krashenninnikov, D. A. D'Ippolito, and J. R. Myra, "Recent theoretical progress in understanding coherent structures in edge and SOL turbulence," *J. Plasma Phys.* **74**(5), 679–717 (2008).
- ³³A. Y. Aydemir, "Convective transport in the scrape-off layer of tokamaks," *Phys. Plasmas* **12**(6), 062503 (2005).
- ³⁴J. R. Myra, D. A. D'Ippolito, S. I. Krashenninnikov, and G. Q. Yu, "Convective transport in the scrape-off-layer by nonthermalized spinning blobs," *Phys. Plasmas* **11**(9), 4267–4274 (2004).
- ³⁵A. Y. Pigarov, S. I. Krashenninnikov, T. D. Rognien, M. J. Schaffer, and W. P. West, "Tokamak edge plasma simulation including anomalous cross-field convective transport," *Phys. Plasmas* **9**(4), 1287 (2002).
- ³⁶S. J. Zweben, D. P. Stotler, J. L. Terry, B. Labombard, M. Greenwald, M. Mutterspaugh, C. S. Pitcher, K. Hallatschek, R. J. Maqueda, B. Rogers, J. L. Lowrance, V. J. Mastrocola, and G. F. Renda, *Phys. Plasmas* **9**, 1981–1989 (2002).
- ³⁷D. A. Russell, D. A. D'Ippolito, J. R. Myra, W. M. Nevins, and X. Q. Xu, "Blob dynamics in 3D BOUT simulations of tokamak edge turbulence," *Phys. Rev. Lett.* **93**(26), 265001 (2004).
- ³⁸G. Q. Yu and S. I. Krashenninnikov, "Dynamics of blobs in scrape-off-layer/shadow regions of tokamaks and linear devices," *Phys. Plasmas* **10**(11), 4413–4418 (2003).
- ³⁹O. E. Garcia, V. Naulin, A. H. Nielsen, and J. J. Rasmussen, "Turbulence simulations of blob formation and radial propagation in toroidally magnetized plasmas," *Phys. Scr.* **2006**(T122), 89–103.
- ⁴⁰N. Bian, S. Benkadda, J.-V. Paulsen, and O. E. Garcia, "Blobs and front propagation in the scrape-off layer of magnetic confinement devices," *Phys. Plasmas* **10**(3), 671–676 (2003).

- ⁴¹W. Fundamenski, O. E. Garcia, V. Naulin, R. A. Pitts, A. H. Nielsen, J. J. Rasmussen, J. Horacek, J. P. Graves, and JET EFDA Contributors, "Dissipative processes in interchange driven scrape-off layer turbulence," *Nucl. Fusion* **47**(5), 417–433 (2007).
- ⁴²P. Ghendrih, Y. Sarazin, G. Attuel, S. Benkadda, P. Beyer, G. Falchetto, C. Figarella, X. Garbet, V. Grandgirard, and M. Ottaviani, "Theoretical analysis of the influence of external biasing on long range turbulent transport in the scrape-off layer," *Nucl. Fusion* **43**(10), 1013–1022 (2003).
- ⁴³Y. Sarazin, P. Ghendrih, G. Attuel, C. Clément, X. Garbet, V. Grandgirard, M. Ottaviani, S. Benkadda, P. Beyer, N. Bian, and C. Figarella, "Theoretical understanding of turbulent transport in the SOL," *J. Nucl. Mater.* **313–316**, 796–803 (2003).
- ⁴⁴P. Ghendrih, Y. Sarazin, G. Attuel, S. Benkadda, P. Beyer, G. Darmet, G. Falchetto, C. Figarella, X. Garbet, V. Grandgirard, and M. Ottaviani, "Statistical analysis of turbulent front propagation in the scrape-off-layer," *J. Nucl. Mater.* **337–339**, 347–351 (2005).
- ⁴⁵X. Q. Xu, W. M. Nevins, R. H. Cohen, J. R. Myra, and P. B. Snyder, "Dynamical simulations of boundary plasma turbulence in divertor geometry," *New J. Phys.* **4**, 53 (2002).
- ⁴⁶J. R. Myra, D. A. Russell, and D. A. D'Ippolito, "Collisionality and magnetic geometry effects on tokamak edge turbulent transport. I. A two-region model with application to blobs," *Phys. Plasmas* **13**(11), 112502 (2006).
- ⁴⁷D. A. Russell, J. R. Myra, and D. A. D'Ippolito, "Collisionality and magnetic geometry effects on tokamak edge turbulent transport. II. Many-blob turbulence in the two-region model," *Phys. Plasmas* **14**(10), 102307 (2007).
- ⁴⁸D. A. D'Ippolito and J. R. Myra, "Thermal transport catastrophe and the tokamak edge density limit," *Phys. Plasmas* **13**(6), 062503 (2006).
- ⁴⁹J. R. Myra and D. A. D'Ippolito, "Edge instability regimes with applications to blob transport and the quasicohherent mode," *Phys. Plasmas* **12**(9), 092511 (2005).
- ⁵⁰O. E. Garcia, N. H. Bian, V. Naulin, A. H. Nielsen, and J. J. Rasmussen, "Mechanism and scaling for convection of isolated structures in nonuniformly magnetized plasmas," *Phys. Plasmas* **12**(9), 090701 (2005).
- ⁵¹J. A. Boedo, J. R. Myra, S. Zweben, R. Maingi, R. J. Maqueda, V. A. Soukhanovskii, J. W. Ahn, J. Canik, N. Crocker, D. A. D'Ippolito, R. Bell, H. Kugel, B. Leblanc, L. A. Roquemore, and D. L. Rudakov, "Edge transport studies in the edge and scrape-off layer of the national spherical torus experiment with Langmuir probes," *Phys. Plasmas* **21**(4), 042309 (2014).
- ⁵²J. R. Myra, D. A. D'Ippolito, D. P. Stotler, S. J. Zweben, B. P. LeBlanc, J. E. Menard, R. J. Maqueda, and J. Boedo, "Blob birth and transport in the tokamak edge plasma: Analysis of imaging data," *Phys. Plasmas* **13**(9), 092509 (2006).
- ⁵³N. Offeddu, W. Han, C. Theiler, T. Golfopoulos, J. L. Terry, E. Marmar, C. Wüthrich, C. K. Tsui, H. De Oliveira, B. P. Duval, D. Galassi, D. S. Oliveira, and D. Mancini, "Cross-field and parallel dynamics of SOL filaments in TCV," *Nucl. Fusion* **62**(9), 096014 (2022).
- ⁵⁴C. K. Tsui, J. A. Boedo, J. R. Myra, B. Duval, B. Labit, C. Theiler, N. Vianello, W. A. J. Vijvers, H. Reimerdes, S. Coda, O. Février, J. R. Harrison, J. Horacek, B. Lipschultz, R. Maurizio, F. Nespoli, U. Sheikh, K. Verhaegh, and N. Walkden, "Filamentary velocity scaling validation in the TCV tokamak," *Phys. Plasmas* **25**(7), 072506 (2018).
- ⁵⁵D. Carralero, P. Manz, L. Aho-Mantila, G. Birkenmeier, M. Brix, M. Groth, H. W. Müller, U. Stroth, N. Vianello, and E. Wolftrum, "Experimental validation of a filament transport model in turbulent magnetized plasmas," *Phys. Rev. Lett.* **115**(21), 215002 (2015).
- ⁵⁶G. Fuchert, G. Birkenmeier, D. Carralero, T. Lunt, P. Manz, H. W. Müller, B. Nold, M. Ramisch, V. Rohde, and U. Stroth, "Blob properties in L- And H-mode from gas-puff imaging in ASDEX upgrade," *Plasma Phys. Control Fusion* **56**(12), 125001 (2014).
- ⁵⁷C. Killer, B. Shanahan, O. Grulke, O. Grulke, M. Endler, K. Hammond, and L. Rudischhauser, "Plasma filaments in the scrape-off layer of Wendelstein 7-X," *Plasma Phys. Control Fusion* **62**(8), 085003 (2020).
- ⁵⁸D. A. D'Ippolito, J. R. Myra, and S. J. Zweben, "Convective transport by intermittent blob-filaments: Comparison of theory and experiment," *Phys. Plasmas* **18**(6), 060501 (2011).
- ⁵⁹R. Kube and O. E. Garcia, "Velocity scaling for filament motion in scrape-off layer plasmas," *Phys. Plasmas* **18**(10), 102314 (2011).
- ⁶⁰J. Olsen, J. Madsen, A. H. Nielsen, J. J. Rasmussen, and V. Naulin, "Temperature dynamics and velocity scaling laws for interchange driven, warm ion plasma filaments," *Plasma Phys. Control Fusion* **58**(4), 044011 (2016).
- ⁶¹O. E. Garcia, N. H. Bian, and W. Fundamenski, "Radial interchange motions of plasma filaments," *Phys. Plasmas* **13**(8), 082309 (2006).
- ⁶²G. Decristoforo, F. Militello, T. Nicholas, J. Omotani, C. Marsden, N. Walkden, and O. E. Garcia, "Blob interactions in 2D scrape-off layer simulations," *Phys. Plasmas* **27**(12), 122301 (2020).
- ⁶³P. Paruta, C. Beadle, P. Ricci, and C. Theiler, "Blob velocity scaling in diverted tokamaks: A comparison between theory and simulation," *Phys. Plasmas* **26**(3), 032302 (2019).
- ⁶⁴J. T. Omotani, F. Militello, L. Easy, and N. R. Walkden, "The effects of shape and amplitude on the velocity of scrape-off layer filaments," *Plasma Phys. Control Fusion* **58**(1), 014030 (2015).
- ⁶⁵F. Militello, B. Dudson, L. Easy, A. Kirk, and P. Naylor, "On the interaction of scrape off layer filaments," *Plasma Phys. Control Fusion* **59**(12), 125013 (2017).
- ⁶⁶B. Nold, P. Manz, T. T. Ribeiro, G. Fuchert, G. Birkenmeier, H. W. Müller, M. Ramisch, B. D. Scott, and U. Stroth, "Turbulent transport across shear layers in magnetically confined plasmas," *Phys. Plasmas* **21**(10), 102304 (2014).
- ⁶⁷T. N. Bernard, F. D. Halpern, M. Francisquez, J. Juno, N. R. Mandell, G. W. Hammett, A. Hakim, E. Humble, and R. Mukherjee, "Effect of neutral interactions on parallel transport and blob dynamics in gyrokinetic scrape-off layer simulations," *Phys. Plasmas* **30**(11), 112501 (2023).
- ⁶⁸D. Schwörer, N. R. Walkden, H. Leggate, B. D. Dudson, F. Militello, T. Downes, and M. M. Turner, "Influence of plasma background on 3D scrape-off layer filaments," *Plasma Phys. Control Fusion* **61**(2), 025008 (2019).
- ⁶⁹A. S. Thyroe, V. Naulin, A. H. Nielsen, and J. Juul Rasmussen, "Dynamics of seeded blobs under the influence of inelastic neutral interactions," *Phys. Plasmas* **27**(5), 052302 (2020).
- ⁷⁰A. S. Thyroe, L. E. H. Tophøj, V. Naulin, J. J. Rasmussen, J. Madsen, and A. H. Nielsen, "The influence of blobs on neutral particles in the scrape-off layer," *Plasma Phys. Control Fusion* **58**(4), 044010 (2016).
- ⁷¹L. Easy, F. Militello, J. Omotani, N. R. Walkden, and B. Dudson, "Investigation of the effect of resistivity on scrape off layer filaments using three-dimensional simulations," *Phys. Plasmas* **23**(1), 012512 (2016).
- ⁷²T. Bernard, see https://mediacentral.princeton.edu/Media/TS07November2022_TBernard/1_sjwax4qa/13468191 for "Effects of blob dynamics and triangularity on blob dynamics in the SOL" (2022).
- ⁷³N. R. Walkden, L. Easy, F. Militello, and J. T. Omotani, "Dynamics of 3D isolated thermal filaments," *Plasma Phys. Control Fusion* **58**(11), 115010 (2016).
- ⁷⁴J. R. Myra, J. Cheng, and S. E. Parker, "Dynamics of rapidly spinning blob-filaments: Fluid theory with a parallel kinetic extension," *Phys. Plasmas* **30**(7), 072302 (2023).
- ⁷⁵J. R. Angus, S. I. Krasheninnikov, and M. V. Umansky, "Effects of parallel electron dynamics on plasma blob transport," *Phys. Plasmas* **19**(8), 082312 (2012).
- ⁷⁶N. R. Walkden, B. D. Dudson, and G. Fishpool, "Characterization of 3D filament dynamics in a MAST SOL flux tube geometry," *Plasma Phys. Control Fusion* **55**(10), 105005 (2013).
- ⁷⁷B. K. Joseph, R. Jha, P. K. Kaw, S. K. Mattoo, C. V. S. Rao, Y. C. Saxena, P. K. Atrey, S. B. Bhatt, B. N. Bach, C. N. Gupta, K. K. Jain, A. Kumar, H. A. Patchak, H. R. Prabhakara, H. D. Pujara, and K. Sathyanarayana, "Observation of vortex-like coherent structures in the edge plasma of the ADITYA tokamak," *Phys. Plasmas* **4**(12), 4292–4300 (1997).
- ⁷⁸M. Spolaore, V. Antoni, R. Cavazzana, G. Regnoli, G. Serianni, E. Spada, N. Vianello, H. Bergsäker, and J. R. Drake, "Effects of E×B velocity shear on electrostatic structures," *Phys. Plasmas* **9**(10), 4110 (2002).
- ⁷⁹M. Lampert, A. Diallo, and S. J. Zweben, "Novel angular velocity estimation technique for plasma filaments," *Rev. Sci. Instrum.* **94**(1), 013505 (2023).
- ⁸⁰J. R. Myra, W. M. Davis, D. A. D'Ippolito, B. Labombard, D. A. Russell, J. L. Terry, and S. J. Zweben, "Edge sheared flows and the dynamics of blob-filaments," *Nucl. Fusion* **53**(7), 073013 (2013).
- ⁸¹M. V. Umansky, S. I. Krasheninnikov, B. LaBombard, and J. L. Terry, "Comments on particle and energy balance in the edge plasma of Alcator C-Mod," *Phys. Plasmas* **5**(9), 3373–3376 (1998).

- ⁸²D. A. D'Ippolito, J. R. Myra, D. A. Russell, and G. Q. Yu, "Rotational stability of plasma blobs," *Phys. Plasmas* **11**(10), 4603–4609 (2004).
- ⁸³J. Cheng, J. Myra, S. H. Ku, R. Hager, C. S. Chang, and S. Parker, "Transport barrier and spinning blob dynamics in the tokamak edge," *Nucl. Fusion* **63**(8), 086015 (2023).
- ⁸⁴J. L. Luxon, "A design retrospective of the DIII-D tokamak," *Nucl. Fusion* **42**(5), 614 (2002).
- ⁸⁵J. Wesson and D. Campbell, in *Tokamaks*, 4th ed. (Oxford U.P., Oxford, 2011), pp. 665–686.
- ⁸⁶T. N. Carlstrom, C. L. Hsieh, R. Stockdale, D. G. Nilson, and D. N. Hill, "Initial operation of the divertor Thomson scattering diagnostic on DIII-D," *Rev. Sci. Instrum.* **68**(2), 1195–1200 (1997).
- ⁸⁷L. L. Lao, H. St. John, R. D. Stambaugh, A. G. Kellman, and W. Pfeiffer, "Reconstruction of current profile parameters and plasma shapes in tokamaks," *Nucl. Fusion* **25**(11), 1611–1622 (1985).
- ⁸⁸O. Meneghini and L. Lao, "Integrated modeling of tokamak experiments with OMFIT," *Plasma Fusion Res.* **8**, 2403009 (2013).
- ⁸⁹J. G. Watkins, J. Salmonson, R. Moyer, R. Doerner, R. Lehmer, L. Schmitz, and D. N. Hill, "A fast scanning probe for DIII-D," *Rev. Sci. Instrum.* **63**(10), 4728–4730 (1992).
- ⁹⁰J. A. Boedo, J. S. DeGrassie, B. Grierson, T. Stoltzfus-Dueck, D. J. Battaglia, D. L. Rudakov, E. A. Belli, R. J. Groebner, E. Hollmann, C. Lasnier, W. M. Solomon, E. A. Unterberg, and J. Watkins, "Experimental evidence of edge intrinsic momentum source driven by kinetic ion loss and edge radial electric fields in tokamaks," *Phys. Plasmas* **23**(9), 092506 (2016).
- ⁹¹D. L. Rudakov, J. A. Boedo, R. A. Moyer, P. C. Stangeby, J. G. Watkins, D. G. Whyte, L. Zeng, N. H. Brooks, R. P. Doerner, T. E. Evans, M. E. Fenstermacher, M. Groth, E. M. Hollmann, S. I. Krashennikov, C. J. Lasnier, A. W. Leonard, M. A. Mahdavi, G. R. McKee, A. G. McLean, A. Y. Pigarov, W. R. Wampler, G. Wang, W. P. West, and C. P. Wong, "Far SOL transport and main wall plasma interaction in DIII-D," *Nucl. Fusion* **45**(12), 1589–1599 (2005).
- ⁹²D. L. Rudakov, J. A. Boedo, R. A. Moyer, R. D. Lehmer, G. Gunner, and J. G. Watkins, "Fast electron temperature diagnostic based on Langmuir probe current harmonic detection on DIII-D," *Rev. Sci. Instrum.* **72**(1 II), 453–456 (2001).
- ⁹³J. A. Boedo, N. Crocker, L. Chousal, R. Hernandez, J. Chalfant, H. Kugel, P. Roney, and J. Wertenbaker, "Fast scanning probe for the NSTX spherical tokamak," *Rev. Sci. Instrum.* **80**(12), 123506 (2009).
- ⁹⁴N. Smick, B. Labombard, and I. H. Hutchinson, "Transport and drift-driven plasma flow components in the Alcator C-Mod boundary plasma," *Nucl. Fusion* **53**(2), 023001 (2013).
- ⁹⁵G. Grenfell, J. Adamek, M. Komm, D. Brida, G. D. Conway, P. Manz, P. Tolias, T. Eich, D. Sestak, A. Herrmann, T. Nishizawa, and U. Stroth, "High-heat flux ball-pen probe head in ASDEX-Upgrade," *Rev. Sci. Instrum.* **93**(2), 023507 (2022).
- ⁹⁶C. Silva, B. Gonçalves, C. Hidalgo, M. A. Pedrosa, W. Fundamenski, M. Stamp, and R. A. Pitts, "Intermittent transport in the JET far-SOL," *J. Nucl. Mater.* **390–391**(1), 355–358 (2009).
- ⁹⁷W. Zhang, J. F. Chang, B. N. Wan, G. S. Xu, C. J. Xiao, B. Li, C. S. Xu, N. Yan, L. Wang, S. C. Liu, M. Jiang, and P. Liu, "Fast reciprocating probe system on the EAST superconducting tokamak," *Rev. Sci. Instrum.* **81**(11), 113501 (2010).
- ⁹⁸P. Liu, G. Xu, H. Wang, M. Jiang, L. Wang, W. Zhang, S. Liu, N. Yan, and S. Ding, "Reciprocating probe measurements of L-H transition in LHCD H-mode on EAST," *Plasma Sci. Technol.* **15**(7), 619–622 (2013).
- ⁹⁹J. Cheng, L. W. Yan, W. Y. Hong, K. J. Zhao, T. Lan, J. Qian, A. D. Liu, H. L. Zhao, Y. Liu, Q. W. Yang, J. Q. Dong, X. R. Duan, and Y. Liu, "Statistical characterization of blob turbulence across the separatrix in HL-2A tokamak," *Plasma Phys. Control Fusion* **52**(5), 055003 (2010).
- ¹⁰⁰O. Pan, Y. Xu, J. Cheng, J. Q. Xu, L. W. Yan, K. Yao, M. Xu, Z. B. Shi, Y. Liu, Q. W. Yang, X. R. Duan, Y. Liu, and H-2A Team, in 26th IAEA Fusion Energy Conference EX/P7-29, Kyoto, Japan (2016).
- ¹⁰¹C. Wüthrich, C. Theiler, N. Offeddu, D. Galassi, D. S. Oliveira, B. P. Duval, O. Février, T. Goufopoulos, W. Han, E. Marmar, J. L. Terry, and C. K. Tsui, "X-point and divertor filament dynamics from gas puff imaging on TCV," *Nucl. Fusion* **62**(10), 106022 (2022).
- ¹⁰²R. Kube, A. Theodorsen, O. E. Garcia, D. Brunner, B. LaBombard, and J. L. Terry, "Comparison between mirror Langmuir probe and gas-puff imaging measurements of intermittent fluctuations in the Alcator C-Mod scrape-off layer," *J. Plasma Phys.* **86**(5), 905860519 (2020).
- ¹⁰³N. Offeddu, C. Wüthrich, W. Han, C. Theiler, T. Goufopoulos, J. L. Terry, E. Marmar, A. Ravetta, and G. Van Parys, "Analysis techniques for blob properties from gas puff imaging data," *Rev. Sci. Instrum.* **94**(3), 033512 (2023).
- ¹⁰⁴S. J. Zweben, J. L. Terry, D. P. Stotler, and R. J. Maqueda, "Invited review article: Gas puff imaging diagnostics of edge plasma turbulence in magnetic fusion devices," *Rev. Sci. Instrum.* **88**(4), 041101 (2017).
- ¹⁰⁵J. A. Boedo, D. Gray, R. W. Conn, P. Luong, M. Schaffer, R. S. Ivanov, A. V. Chernilevsky, and G. van Oost, "On the harmonic technique to measure electron temperature with high time resolution," *Rev. Sci. Instrum.* **70**(7), 2997–3006 (1999).
- ¹⁰⁶M. Kočan, F. P. Gennrich, A. Kendl, and H. W. Müller, "Ion temperature fluctuations in the ASDEX Upgrade scrape-off layer," *Plasma Phys. Control Fusion* **54**(8), 085009 (2012).
- ¹⁰⁷M. Kočan, J. P. Gunn, J. Y. Pascal, G. Bonhomme, C. Fenzi, E. Gauthier, and J. L. Segui, "Edge ion-to-electron temperature ratio in the Tore Supra tokamak," *Plasma Phys. Control Fusion* **50**(12), 125009 (2008).
- ¹⁰⁸M. Komm, M. Kočan, D. Carralero, H. W. Müller, J. Stöckel, and C. Rea, in *41st EPS Conference on Plasma Physics*, P5.026 (2014).
- ¹⁰⁹S. 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 Phys. Control Fusion* **60**(10), 105001 (2018).
- ¹¹⁰D. Carralero, S. Artene, M. Bernert, G. Birkenmeier, M. Faitsch, P. Manz, P. De Marne, U. Stroth, M. Wischmeier, and E. Wolfrum, "On the role of filaments in perpendicular heat transport at the scrape-off layer," *Nucl. Fusion* **58**(9), 096015 (2018).
- ¹¹¹B. Nold, T. T. Ribeiro, M. Ramisch, Z. Huang, H. W. Müller, B. D. Scott, and U. Stroth, "Influence of temperature fluctuations on plasma turbulence investigations with Langmuir probes," *New J. Phys.* **14**, 063022 (2012).
- ¹¹²B. Shanahan, C. Killer, G. Pechstein, S. A. Henneberg, G. Fuchert, and O. Grulke, "Estimating the error in filament propagation measurement using a synthetic probe," *Plasma Phys. Control Fusion* **63**(12), 125018 (2021).
- ¹¹³H. Johnsen, H. L. Pécseli, and J. Trulsen, "Conditional eddies in plasma turbulence," *Phys. Fluids* **30**(7), 2239–2254 (1987).
- ¹¹⁴A. V. Filippas, R. D. Bengston, G.-X. Li, M. Meier, C. P. Ritz, and E. J. Powers, "Conditional analysis of floating potential fluctuations at the edge of the Texas experimental tokamak upgrade (TEXT-U)," *Phys. Plasmas* **2**(3), 839–845 (1995).
- ¹¹⁵J. Brotankova, J. Stockel, J. Horacek, J. Seidl, I. Duran, M. Hron, and G. van Oost, "Measurement of sheared flows in the edge plasma of the CASTOR tokamak," *Plasma Phys. Rep.* **35**(11), 980–986 (2009).
- ¹¹⁶L. L. Lao, H. E. St. John, Q. Peng, J. R. Ferron, E. J. Strait, T. S. Taylor, W. H. Meyer, C. Zhang, and K. I. You, "MHD equilibrium reconstruction in the DIII-D tokamak," *Fusion Sci. Technol.* **48**(2), 968–977 (2005).
- ¹¹⁷BIPM, IEC, IFCC, ILAC, ISO, IUPAC, IUPAP, and OIML, "Guide to the expression of uncertainty in measurement," JCGM 100:2008 (GUM 1995 with minor corrections), 2008, http://www.bipm.org/utls/common/documents/jcgm/JCGM_100_2008_E.pdf.
- ¹¹⁸C. Colin, P. Tamaing, P. Ghendrih, F. Schwander, and E. Serre, "Impact of a Langmuir probe on turbulence measurements in the scrape-off layer of Tokamaks," *Contrib. Plasma Phys.* **54**(4–6), 543–548 (2014).
- ¹¹⁹R. Futersack, C. Colin, P. Tamaing, G. Ciraolo, P. Ghendrih, Y. Marandet, F. Schwander, and E. Serre, "First principle modelling of interplay between Langmuir probes and plasma turbulence," *Contrib. Plasma Phys.* **56**(6–8), 575–580 (2016).
- ¹²⁰B. Nold, E. Holzhauser, H. W. Müller, L. Kammerloher, and ASDEX Upgrade Team, "Electronic issues of Langmuir probe measurements on the midplane manipulator of ASDEX upgrade," Report No. IPP 10/43, Max-Planck-Institut für Plasmaphysik, Garching, Germany, 2012.
- ¹²¹R. Jha and Y. C. Saxena, "Wavelet analysis of ADITYA edge turbulence: Evidence of non-linear interaction," *Phys. Plasmas* **3**(8), 2979–2989 (1996).

- ¹²²R. A. Moyer, R. D. Lehmer, T. E. Evans, R. W. Conn, and L. Schmitz, "Nonlinear analysis of turbulence across the L to H transition," *Plasma Phys. Control Fusion* **38**(8), 1273–1278 (1996).
- ¹²³O. E. Garcia, I. Cziegler, R. Kube, B. LaBombard, and J. L. Terry, "Burst statistics in Alcator C-Mod SOL turbulence," *J. Nucl. Mater.* **438**, S180–S183 (2013).
- ¹²⁴A. Theodorsen, *Statistical Properties of Intermittent Fluctuations in the Boundary of Fusion Plasmas* (University of Tromsø - The Arctic University of Norway, 2018).
- ¹²⁵R. Kube, A. Theodorsen, O. E. Garcia, B. Labombard, and J. L. Terry, "Fluctuation statistics in the scrape-off layer of Alcator C-Mod," *Plasma Phys. Control Fusion* **58**(5), 054001 (2016).
- ¹²⁶R. Kube, O. E. Garcia, A. Theodorsen, A. Q. Kuang, B. LaBombard, J. L. Terry, and D. Brunner, "Statistical properties of the plasma fluctuations and turbulent cross-field fluxes in the outboard mid-plane scrape-off layer of Alcator C-Mod," *Nucl. Mater. Energy* **18**, 193–200 (2019).
- ¹²⁷J. R. Myra, S. J. Zweben, and D. A. Russell, "Blob-hole correlation model for edge turbulence and comparisons with NSTX gas puff imaging data," *Plasma Phys. Control Fusion* **60**(7), 075015 (2018).
- ¹²⁸B. Nold, G. D. Conway, T. Happel, H. W. Müller, M. Ramisch, V. Rohde, and U. Stroth, "Generation of blobs and holes in the edge of the ASDEX Upgrade tokamak," *Plasma Phys. Control Fusion* **52**(6), 065005 (2010).
- ¹²⁹J. A. Boedo, D. L. Rudakov, R. A. Moyer, G. R. McKee, R. J. Colchin, M. J. Schaffer, P. G. Stangeby, W. P. West, S. L. Allen, T. E. Evans, R. J. Fonck, E. M. Hollmann, S. Krashenninikov, A. W. Leonard, W. Nevins, M. A. Mahdavi, G. D. Porter, G. R. Tynan, D. G. Whyte, and X. Xu, *Phys. Plasmas* **10**, 1670–1677 (2003).
- ¹³⁰S. J. Zweben, J. R. Myra, W. M. Davis, D. A. D'Ippolito, T. K. Gray, S. M. Kaye, B. P. Leblanc, R. J. Maqueda, D. A. Russell, and D. P. Stotler, "Blob structure and motion in the edge and SOL of NSTX," *Plasma Phys. Control Fusion* **58**(4), 044007 (2016).
- ¹³¹S. Garland, G. Fuchert, M. Ramisch, and T. Hirth, "The structure and poloidal dynamics of blob filaments in TJ-K," *Plasma Phys. Control Fusion* **58**(4), 044012 (2016).
- ¹³²D. Carralero, G. Birkenmeier, H. W. Müller, P. Manz, P. Demarne, S. H. Müller, F. Reimold, U. Stroth, M. Wischmeier, and E. Wolfrum, "An experimental investigation of the high density transition of the scrape-off layer transport in ASDEX Upgrade," *Nucl. Fusion* **54**(12), 123005 (2014).
- ¹³³W. Han, R. A. Pietersen, R. Villamor-Lora, M. Beveridge, N. Offeddu, T. Golfinopoulos, C. Theiler, J. L. Terry, E. S. Marmar, and I. Drori, "Tracking blobs in the turbulent edge plasma of a tokamak fusion device," *Sci. Rep.* **12**(1), 18142 (2022).
- ¹³⁴O. E. Garcia, R. A. Pitts, J. Horacek, J. Madsen, V. Naulin, A. H. Nielsen, and J. J. Rasmussen, "Collisionality dependent transport in TCV SOL plasmas," *Plasma Phys. Control Fusion* **49**(12B), B47–B57 (2007).
- ¹³⁵D. Wendler, G. Birkenmeier, M. Griener, R. Dux, R. Fischer, D. Stieglitz, E. Wolfrum, and U. Stroth, "Two-dimensional reconstruction of filament temperatures and densities with the thermal helium beam at ASDEX Upgrade," *Phys. Plasmas* **30**(9), 092509 (2023).

Chapter 3, with the exception of the leading paragraph, is a reprint of the material as it appears in Molesworth, Steven; Boedo, Jose; Tsui, Cedric; Perillo, Renato; Rudakov, Dmitry, “Filamentary velocity scaling validation and spin dynamics in the DIII-D tokamak,” *Physics of Plasmas* 31 (4), 042515 (1 April 2024), <https://doi.org/10.1063/5.0195025>; published by AIP Publishing under a Creative Commons Attribution (CC BY) license, which grants permission to reproduce it here. The dissertation author was the primary investigator and author of this paper.

Chapter 4

The Near-Edge Profile Prediction Framework

4.1 Framework overview

The framework is an assembled, modular workflow that carries the proven core flux matching loop out to the last closed flux surface and equips it with the additional physics the edge requires, and its structure is shown in Figure 4.1. A profile parameterization turns a handful of design variables into the kinetic profile inverse scale lengths a/L_y and the kinetic profiles themselves follow through integration. Edge source and sink modules supply the neutral and impurity physics, while a transport model returns the turbulent and neoclassical fluxes. Rotation and MHD stability components serve as modifiers in the transport calculations. The mismatch between those modeled fluxes and the volume-integrated sources forms the residual the flux match aims to drive to zero over the domain $r/a \in [0.88, 1]$, with the kinetic values imposed at the separatrix and the integrated core power entering across the inner boundary.

The edge-specific pieces are *plug-and-play*. The neutral, impurity, rotation, target, and transport modules each present a fixed interface to the solver, so the reduced models used here can be swapped for higher-fidelity alternatives without touching the surrounding loop. Following the project scope, the framework builds on the PORTALS machinery, so higher fidelity physics may be incorporated in the future.




Figure 4.1. The near-edge profile-prediction workflow. Design variables parameterize the profiles, the edge modules feed a transport-flux and a target-flux branch, and the two reconverge at the residual $R = F^{\text{tr}} - F^{\text{tgt}}$ the flux match drives to zero. Dashed outlines mark interchangeable pieces.

4.2 Profile parameterizations

Core solvers typically represent each channel by a piecewise linear a/L_y spline integrated from a fixed separatrix value, which suffices where a/L_y varies by $\mathcal{O}(1)$ across the domain. The choice is a convention rather than a requirement, so what follows contrasts against the common practice of TGYRO and PORTALS rather than all core transport solvers generically. The near-edge instead demands a parameterization that spans the full L-to-H-mode range of edge shapes, yet must still be defined by a relatively small number of parameters for efficient computation. The form adopted here blends a spline’s expressiveness with EPED’s hyperbolic tangent structure, adding an asymmetry term. Each evolved channel is described by a modified hyperbolic tangent (mtanh),

$$y(x) = A[1 - \tanh u(x)] - m(x - 1) + b, \quad u(x) = \frac{x - c}{\Delta_0[1 + \delta(x - c)]}, \quad (4.1)$$

in the normalized coordinate $x = r/a$.

Beyond the standard pedestal amplitude A , midpoint c , and width Δ_0 , Equation (4.1) carries two shape parameters. A linear offset slope m sustains the finite background gradient beneath the pedestal, and a width skewness δ lets the inboard and outboard sides of the transition take different curvature. With the constant offset b , these six parameters $\{A, \Delta_0, \delta, m, c, b\}$ reproduce both a near-featureless L-mode edge and a sharp, asymmetric H-mode pedestal.

Neither addition is an invention of this work. The standard pedestal fit has long carried an offset term to accommodate the transition from the pedestal to the core [37, 46], and it is that term the parameterization leans on. What differs is the use it is put to, because the term is retained as a free design parameter rather than fitted and set aside. A form without it returns $a/L_y \rightarrow 0$ over the inner knots, so a gradient-based solver has nothing to control there (Section 4.2.1).

A solver operating inside the LCFS has no knowledge of the SOL, so the separatrix conditions anchor the construction rather than being solved. The LCFS *values* $(n_e, T_e, T_i)_{\text{LCFS}}$ are user-provided inputs, taken from measurement or an external boundary prescription. The design variables presented to the optimizer are the inverse gradient scale lengths $z_k = a/L_y(x_k)$ at a small set of control radii just inside the separatrix, augmented by the separatrix gradient $z_{\text{LCFS}} = a/L_y(1)$ for each channel. The mtanh is anchored at the last closed flux surface by *both* boundary conditions. The fixed value y_{LCFS} and the gradient z_{LCFS} close the offset b and the background slope m analytically. The remaining four parameters $\{A, \Delta_0, \delta, c\}$ are then obtained by a bounded least squares fit to the interior knot gradients, the bounds being a feasibility box that keeps the profile monotonic and positive.

The anchored and offset mtanh reconstructs not only the pedestal but the a/L_y *peak* that often sits just inside the separatrix where, in this work, no flux matching knot samples. That peak occurs naturally from the analytic form together with the supplied boundary gradient, not out of any special treatment of z_{LCFS} . Inferring a steep, narrow peak in a/L_y is yet another deviation from core transport solvers and provides a major source of uncertainty when flux matching knots are not nearly collocated with the peak.

A corrector spline is added on top of the mtanh backbone, passing through the residuals

between fitted and target gradients at the interior knots so that a/L_y is pinned exactly to the design values there, as a flux match solve requires. Because the analytic backbone generally reproduces the knots to a few percent, this correction is small unless the solver proposes design variables far from physical.

The profile follows by integrating the gradient inward from the separatrix,

$$y(x) = y_{\text{LCFS}} \exp\left(\int_x^1 a/L_y dx'\right), \quad (4.2)$$

exactly as in the core. Imposing y_{LCFS} as a fixed input removes a fitting parameter and makes the parameter-to-profile map well-posed, while retaining z_{LCFS} as a separate handle lets the edge steepness be set independently of the interior knots. It may be held fixed or released to the solver. Both the validation runs of Chapter 5 and the studies of Chapter 6 hold it fixed at the measured value, leaving four interior knots in each of three channels as the twelve design variables. Where no measured edge profile exists, it is supplied instead by the sheared spectral filament model of Section 4.7.

For this work, those four knots sit at $r/a = 0.88, 0.91, 0.94$ and 0.97 , evenly spaced, and the layout is a compromise rather than an optimum. Three considerations fix it. The inner knot is placed at $r/a = 0.88$ so that the solver domain overlaps the region a core solver would already have converged, which is what lets the two be handed to one another. The outer knot is arbitrarily held at $r/a = 0.97$ on the basis that core transport model validity is increasingly broken as one approaches the separatrix, so moving it outward buys constraint at the price of trusting the transport model exactly where it is weakest. Even spacing between them is the neutral choice, since concentrating knots on the gradient peak presumes a pedestal width the solver is supposed to predict, but it is not required. The layout is a choice of answer and not a convergence detail, though, so Section 6.3.2 measures what it costs by solving four alternative layouts end to end on both H-modes, two of four knots and two of five, each at a narrow and a wider span.

4.2.1 Expressiveness against the alternatives

Figures 4.2 to 4.4 test the construction across the three evolved channels on measured profiles from three DIII-D discharges of Chapter 5: two H-modes at contrasting collisionality (179444, low- v^* , and 179449, high- v^*) and an L-mode (128913). The scheme labeled as this work is the production implementation, with four interior a/L_y design variables at $\rho = 0.88$ – 0.97 and the separatrix value and gradient supplied as boundary conditions rather than as a fifth knot. Every comparison scheme is given the same knot gradients and the same separatrix gradient, so the test is like for like.

The near-separatrix gradient carries the most leverage over the profile, and it has a structure the parameterization must capture. In the H-mode a/L_y peaks within the pedestal and falls to a foot at the separatrix. In the L-mode the gradients are weak across the pedestal and steepen only in a thin near-separatrix layer, with a/L_{T_e} rising toward the cool (50 eV) separatrix as $1/T_e$, so the profile information concentrates in the sliver the flux match knots barely reach.

The naive a/L_y interpolants, piecewise-linear and cubic spline, cannot represent that structure from a few knots, so the integrated profile inherits the error. They over-shoot the weak L-mode density pedestal and under-infer the H-mode gradient peak, and the cubic spline additionally oscillates. The canonical EPED tanh, fit to the *same* knots, is more expressive within the pedestal but rigid at its edges. Its symmetry point is tied to the width ($x_{\text{mid}} = 1 - w_{\text{ped}}/2$) and it carries no offset slope, so its top is flat and $a/L_y \rightarrow 0$ over the inner knots. It can honor the separatrix gradient with a narrow width or the interior gradients with a wide one, but not both.

The offset-mtanh backbone forms the correct near-separatrix gradient *shape*. Its offset term gives a/L_y the H-mode foot and the steep near-separatrix rise that neither the interpolants nor the symmetric tanh produce. Fit to the same handful of knots, with no information from the profile between them, it is the most consistently accurate of the four. Over the nine case-channels its root-mean-square profile error at the pedestal top is 7.3%, against 14, 19, and 31% for the cubic spline, the piecewise-linear interpolant, and the canonical EPED tanh. It is the most accurate

scheme in four channels and the least accurate in one, where all four agree to within a few percent, while each of the other three is the least accurate somewhere.

The dual requirement behind this test is intrinsic to the approach. A flux matching solver controls a/L_y at the flux match locations, exactly as in TGYRO and PORTALS, while the various submodels it drives consume the integrated profiles y . The parameterization must therefore both honor the prescribed a/L_y at the knots and return an accurate y , in pedestal and non-pedestal regimes alike. It is not claimed that the offset-mtanh is uniquely optimal for reconstructing y exactly or with the fewest parameters.

4.2.2 Tail inference is the binding lever

A consequence of this construction deserves emphasis, because it sets where the predictive difficulty lies. The flux matching is evaluated only inside $r/a \lesssim 0.97$, but the inverse scale length peak often sits *outside* that window, between the outermost interior knot and the separatrix, so it is not flux matched but *inferred* by the analytic form from the gradients just inside it. Where the peak falls is regime-dependent, and it lies outside the window in the density channel of both validation H-modes, at $r/a \approx 0.99$ in the high-collisionality case (Figure 4.3).

Figure 4.5 makes the leverage concrete on the high-collisionality H-mode a/L_{n_e} . Holding the boundary condition fixed and scaling the two knots at $\rho = 0.94$ and 0.97 by $\pm 50\%$ shifts the gradient carried through the inferred tail. Because the density integrates inward from the fixed separatrix value through that tail, the pedestal top density at $\rho = 0.90$ inherits a $+102\%/ -39\%$ spread, a factor of 3.3, from a change confined to the last two flux match knots. The transfer is close to proportional in logarithm, the gain between the imposed knot ratio and the resulting pedestal top ratio being 1.08, so the tail inference neither damps nor greatly amplifies an interior-knot error.

Meeting the reactor-handoff tolerance of $<10\%$ on the pedestal top density therefore demands predicting the outermost flux matched gradients to roughly the same $\pm 10\%$. That is stringent, because the ~ 0.03 -wide inferred tail carries more of the pedestal top profile than the

DIII-D 179444 (H-mode, $\nu^* \ll 1$)




Figure 4.2. Profile (left) and inverse gradient scale length (right) for n_e , T_e and T_i (rows) on the measured DIII-D 179444 low-collisionality H-mode. Reconstructions use a/L_y at the marked knots ($r/a = 0.88$ – 1) with y fixed at the separatrix (square). The SplineMtnh of this work alone forms the near-separatrix gradient structure.

DIII-D 179449 (H-mode, high ν^*)




Figure 4.3. As Figure 4.2, for the measured DIII-D 179449 high-collisionality H-mode. The density gradient peaks at $r/a \approx 0.99$, between the two outermost knots, so the interpolants miss it and under-shoot the pedestal top density while the offset-mtnh infers it.

DIII-D 128913 (L-mode / shortfall)




Figure 4.4. As Figure 4.2, for the measured DIII-D 128913 L-mode, with T_e smoothed to a physical 50eV separatrix over $r/a > 0.97$. The weak density pedestal is the hard case, every scheme over-shooting n_e at the pedestal top: SplineMtanh +6.9%, cubic +26%, piecewise-linear +43%, EPED tanh +69%.




Figure 4.5. (a) Reconstructed a/L_{n_e} and (b) the integrated density for the DIII-D 179449 high-collisionality H-mode, with $\pm 50\%$ perturbations applied only to the two knots at $r/a = 0.94$ and 0.97 (markers, double arrow). The shaded band ($0.97 < r/a < 1$) is where the transport model is not evaluated. The pedestal top inherits a $+102\% / -39\%$ spread.

flux matched region interior to it (Section 6.3.2). Tail inference, not interior knot count, is the binding lever on predictive accuracy.

The near-proportional transfer means that an accurate flux match yields an accurate inferred tail, so the tail inference adds no error of its own. It does not mean the flux match is sufficient, because the knots are only one of the two things that set the tail, the other being the imposed separatrix gradient. Scaling $a/L_{n_e}(1)$ by the same $\pm 50\%$ (Figure 4.6), with the interior knots held at their fitted values, moves the pedestal top density by $+41\% / -11\%$, a factor of 1.6.

In logarithm that is about two-fifths of the interior-knot response, so the outermost *flux matched* knot is the stronger of the two anchors and the accuracy demand falls first on the transport model at $\rho = 0.97$. The boundary gradient is not a small lever in absolute terms. It is imposed rather than solved for, it is estimator-dependent at the separatrix, and it acts on a region carrying more of the pedestal e-fold than the flux matched region does. Section 6.3.3 measures the same lever end to end on one discharge, and finds that halving the imposed density gradient alone moves the converged pedestal top pressure from 84% below the measurement to




Figure 4.6. As Figure 4.5, but perturbing the *separatrix* gradient $a/L_{n_e}(1)$ by $\pm 50\%$ (markers at $r/a = 1$) with the interior knots and the separatrix value held fixed. The pedestal top moves by $+41\%/-11\%$, about two-fifths of the interior-knot response in logarithm.

45% above.

4.3 Sources and losses

Neutrals.

The main-ion (D^0) neutral population supplies the dominant edge particle source and an electron energy sink, computed by an analytic solver run inside every iteration. The default is a two-population opacity model. Wall-recycled neutrals enter cold at the Franck–Condon energy (~ 3 eV) and are lost to both ionization and charge exchange, while each charge-exchange event converts a cold neutral into a hot one at the local ion temperature. The cold charge-exchange loss is therefore the distributed birth source of a hot population that penetrates deeper and is lost only to ionization. The cold–hot split is not prescribed but emerges from the local v_{cx}/v_{iz} balance.

Particle accounting is strict. The amplitude is fixed by setting the cold inward flux across

the LCFS equal to the recycling source rate, and the balance closes as

$$\dot{N}_{\text{source}} = \int v_{\text{iz}} n_0 dV + \Phi_{\text{inner}} + \Phi_{\text{outer}}, \quad (4.3)$$

where Φ_{inner} is hot-neutral escape across the inner boundary, which re-emerges in steady state as an outward electron flux, and Φ_{outer} is escape to the SOL. A legacy single-population solver selects, by the edge Knudsen number $\text{Kn} = \lambda_{\text{mfp}}/L_{n_e}$, between a diffusive treatment and a free-streaming kinetic one with Beer–Lambert attenuation [47],

$$n_0(r) \propto \exp[-\tau(r)], \quad \tau(r) = \int_r^{r_{\text{LCFS}}} \frac{v_{\text{iz}}(r')}{v_{\text{th}}(r')} dr'. \quad (4.4)$$

Ionization rate coefficients are taken from ADAS [48], with a Voronov analytic fall-back [49]. The volumetric ionization source $S_{\text{ion}} = n_0 v_{\text{iz}}$ enters the electron and main-ion particle targets, while each ionization removes energy from the electron channel. That cost is the bare 13.6 eV potential where the accompanying excitation radiation is represented explicitly, or a lumped ~ 40 eV otherwise. Charge exchange additionally drains ion energy, $\sim \frac{3}{2}(T_i - T_0)$ per event, with the boundary neutrals thermalized to the ion temperature at the LCFS.

Two approximations bound the module’s validity inside the LCFS. First, molecular deuterium physics is neglected, because within the separatrix the plasma dissociates molecules within a short distance of the boundary, so atomic ionization and charge exchange dominate the budget. Molecular channels are properly a divertor and SOL concern, treated by kinetic neutral codes like EIRENE [50]. Second, the source is applied as a flux surface averaged quantity, whereas the true neutral influx is strongly localized near the X-point and the main recycling regions. As with the flux surface averaged α_{MHD} of Section 2.2, the average smears a poloidal structure most pronounced where the flux surfaces themselves depart furthest from their mean.

Figure 4.7 shows the module’s behavior. The two-population model attenuates the cold Franck–Condon population near the separatrix while the charge-exchange hot population supplies




Figure 4.7. Reduced neutral module on a synthetic pedestal. (a) Neutral density penetration normalized at the separatrix, for the two-population model (solid), its cold and hot components (dotted, dash-dot), a single population at T_i (dashed), and a constant- λ_0 exponential (short dashes). (b) Pedestal opacity w_{ped}/λ_0 against pedestal top density, with the analytic scaling (gray) and the validation cases marked.

a deeper tail, and the ionization depth scales correctly with density. Absent a first-principles kinetic neutral code, this behavioral consistency together with the strict particle accounting of Equation (4.3) is the appropriate check for a reduced module, and by the plug-and-play design such a code could replace it.

Impurities.

The impurity charge-state distribution and its radiation are obtained from a steady state collisional-radiative model using the AURORA package [51]. For each charge state z the radial density n_z obeys a continuity balance between cross-field transport and ionization/recombination coupling to neighboring stages,

$$\frac{1}{V'} \frac{\partial}{\partial r} (V' \Gamma_z) = n_e (S_{z-1} n_{z-1} - (S_z + \alpha_z) n_z + \alpha_{z+1} n_{z+1}), \quad (4.5)$$

where Γ_z is the cross-field flux of stage z , $V' = dV/dr$, and S_z, α_z are the ADAS ionization and recombination coefficients [48], with charge-exchange recombination on the main-ion neutrals

retained.

The total impurity density is prescribed, carried as a fixed fraction f_Z of the predicted electron density taken from the experimental reconstruction, so that $a/L_{n_Z} = a/L_{n_e}$ holds exactly and the impurity peaking is tied to the density channel rather than solved for. This is the convention core PORTALS solves use when not evolving n_Z explicitly. Solving the impurity transport instead would import an unconstrained sensitivity: the diffusivity D_Z and pinch velocity V_Z profiles are not known independently for these discharges, yet through $a/L_{n_Z} = -aV_Z/D_Z$ they set the impurity peaking and, with the source, its magnitude. Modeling the impurity channel with naive transport coefficients or within the neoclassical and turbulence evaluations described later alters total predicted fluxes at order unity in some cases. The choice to prescribe f_Z does not remove that sensitivity but rather constrains it for better evaluation of the other components in the workflow.

Equation (4.5) fixes only the distribution across charge states at each radius. The charge-weighted total is rescaled at every radius so as to match the prescribed density, which makes the impurity source rate inert. While the prescribed V_Z is inconsequential on its own, the impurity diffusivity D_Z sets the residence time against the atomic rates and so shifts the ionization balance, altering Z_{eff} slightly.

Three target flux quantities are provided by AURORA solutions. First, the impurity and main ion line plus continuum radiation Q_{rad} are added to the electron energy target. For trace wall impurities, this is typically a small correction on the energy entering from the core. Second, the net impurity ionization

$$\sum_z z \dot{N}_z = \sum_z n_e S_z n_z - \sum_z n_e \alpha_z n_z \quad (4.6)$$

is a freed electron source entering the electron particle target. It is exact by charge conservation and so requires no trace impurity assumption. This quantity can be large relative to the electrons contributed from the main species and its radial variation depends strongly on D_Z . Because the accuracy of $\Gamma_{e,\text{tar}}$ is not the bottleneck on profile accuracy, as will be shown in Chapter 5, the

D_Z profile is again left for future work. Lastly, the charge content $\sum_{z \geq 1} z n_z$ closes the main ion density by quasineutrality $n_{\text{main}} = n_e - \sum_{z \geq 1} z n_z$ and fixes Z_{eff} , so the impurity reaches both heat channels through main-ion dilution and through the collisional exchange P_{ei} that dilution sets.

4.4 Rotation and the radial electric field

The radial electric field, and hence the $E \times B$ shear that regulates turbulence, follows from the main-ion radial force balance,

$$E_r = \frac{1}{Z_i e n_i} \frac{dp_i}{dr} + v_{\phi,i} B_\theta - v_{\theta,i} B_\phi, \quad (4.7)$$

whose three terms are the ion diamagnetic, toroidal rotation, and poloidal rotation fields. The poloidal velocity is set to its neoclassical value $v_{\theta,i} = K_{\text{neo}} B_\phi (dT_i/dr)/(Z_i e B^2)$ with $B^2 = B_\phi^2 + B_\theta^2$ and the poloidal flow coefficient K_{neo} from Sauter, Angioni and Lin-Liu [52]. The toroidal mass rotation is carried by the initial `input.gacode`, the plain text plasma state file of the GACODE suite, which holds the experimental equilibrium geometry, the species list, the kinetic profiles and the power and particle source terms on a single radial grid and is the format TGLF, NEO and TGYRO all read. That rotation is extracted once and held fixed through the solve, so the measured rotation is retained while the diamagnetic and neoclassical-poloidal contributions evolve with the predicted pressure.

The $E \times B$ rotation and its shearing rate follow as

$$\omega_0 = \frac{E_r}{R B_\theta}, \quad \gamma_E = \frac{r}{|q|} \frac{d\omega_0}{dr} \frac{a}{c_s}. \quad (4.8)$$

The diamagnetic and neoclassical poloidal contributions to E_r are regenerated from the updated pressure each iteration, so γ_E tracks the predicted profiles over those terms rather than being frozen at an initial estimate [15, 42]. The toroidal term is not evolved, so the share of the shear it carries is set by the measurement throughout. On the default analytic path every rotation quantity

Table 4.1. $E \times B$ shearing rate γ_E (in c_s/a) at $\rho = 0.94$ for the two validation discharges: the rotation module with the measured toroidal rotation retained, versus the independent experimental analysis of Haskey et al. [53]. The $v_{\phi,i} = 0$ closure (last column) is shown for contrast. It departs by $\sim 50\%$, confirming that the toroidal contribution to the shear is not negligible.

Discharge	Haskey et al. [53] Table I	module (measured v_ϕ)	module ($v_\phi = 0$)
Low-collisionality	1.24	1.32	0.48
High-collisionality	0.70	0.73	1.10

is a differentiable function of the profiles. A higher fidelity backend instead obtains ω_0 from a NEO drift-kinetic solve, which is not differentiable but is consumed by the already existing PORTALS Gaussian Process surrogate transport models.

Figure 4.8 shows the three terms at the initial state of two validation cases. The diamagnetic term sets the magnitude, the toroidal term is comparable and largely opposes it, and the neoclassical poloidal term is small throughout.

The reconstructed shear is verified against the independent experimental analysis of Haskey et al. [53], whose Table I reports γ_E at $\rho_{\text{tor}} = 0.94$ for the two H-mode discharges used in Chapter 5. With the measured toroidal rotation retained, the module reproduces these values within $\sim 6\%$ (Table 4.1). Discarding it, under the $v_{\phi,i} = 0$ closure, misses by $\sim 50\%$ in either direction. The toroidal rotation is subsonic ($M_\phi \sim 0.1$, so its *sonic* neoclassical corrections are negligible and the drift-kinetic solve runs in its low-rotation ordering), yet it can carry a non-negligible portion of the $E \times B$ shear through its steep edge gradient.

4.5 Transport

The transport flux in each channel is the sum of a turbulent and a neoclassical contribution, $F^{\text{tr}} = F^{\text{turb}} + F^{\text{neo}}$. Turbulent fluxes are computed with TGLF-SAT3 throughout this work.

4.5.1 Neoclassical fluxes

Neoclassical fluxes are computed with NEO [22], which solves the linearized drift-kinetic equation for the first-order correction to a Maxwellian background. Poloidal variation of B




Figure 4.8. Rotation module at the initial state of the L-mode (left) and low-collisionality H-mode (right) runs, read from the run tree. *Top:* the three terms of Equation (4.7) as contributions to ω_0 . *Bottom:* the shearing rate with the measured toroidal rotation retained and under the $v_{\phi,i} = 0$ closure, against the Haskey et al. [53] Table I value.

traps particles on banana orbits, and it is that orbit width, rather than the gyroradius, that sets the collisional step, raising the diffusivity roughly two orders of magnitude above the classical estimate. The distribution is never split into mean and fluctuating parts, so no turbulence enters and the two flux contributions are independent by construction.

The solve is local, and so orders out finite orbit width corrections. Those corrections are formally largest exactly where the banana width becomes comparable to the pedestal width, which is the region this work models. The neoclassical flux is comparable to turbulent flux in Q_i for specific cases at experimental profiles, but subdominant to the turbulent flux in every converged case of Chapter 5, so the approximation does not limit the predictions.

4.5.2 Quasilinear theory and the gyrofluid closure

TGLF is a quasilinear model, and the near-edge extrapolation this work tests is best understood against what quasilinear theory assumes. The turbulent flux of a transported quantity is the flux surface average of the fluctuating $E \times B$ advection, $\langle \tilde{v}_{E \times B} \tilde{g} \rangle$, with $\tilde{g}$ the fluctuating density or pressure. The *quasilinear* approximation factorizes each spectral contribution of that average into the shape of the linear eigenmode and the amplitude at which it saturates,

$$F^{\text{turb}} = \sum_{k_y} W(k_y) |\hat{\phi}(k_y)|^2, \quad W(k_y) = \frac{\langle \tilde{v}_{E \times B, k_y} \tilde{g}_{k_y}^* \rangle}{|\hat{\phi}(k_y)|^2} \Big|_{\text{linear}}, \quad (4.9)$$

where the quasilinear weight $W(k_y)$, the transport per unit fluctuation intensity, is fixed entirely by the linear response through the cross-phase between the field and the transported moment. The saturated intensity $|\hat{\phi}(k_y)|^2$ comes separately from a saturation rule (Section 4.5.3). The assumption is that the cross-phases of the linear eigenmodes survive into the saturated phase. That is a statement about fluctuation intensity rather than about scale separation, because it is the weak turbulence limit, in which the nonlinear term redistributes energy across the spectrum without reworking the phase relation inside a mode. It is expected to weaken once the fluctuation amplitude approaches the background, which is the edge condition of Section 2.3, and the factorization has been shown to break down there [12].

Both factors require the linear spectrum $\{\gamma, \omega\}(k_y)$ and eigenfunctions, which gyrokinetics cannot supply at every radius and flux match iteration. TGLF instead takes velocity moments of the linearized, gyro-averaged δf gyrokinetic equation to obtain a gyrofluid model [9], evolving density, parallel momentum, parallel and perpendicular pressure, and parallel heat flux, with separate moments for the trapped and total velocity space. The hierarchy is closed by gyro-Landau-fluid approximations that reproduce parallel Landau damping and finite-Larmor-radius resonances without resolving velocity space [54, 55]. Solved as a linear eigenvalue problem in a Gauss–Hermite ballooning basis, this returns $\{\gamma, \omega, W\}(k_y)$ three to four orders of magnitude

more cheaply than a gyrokinetic calculation, which is what makes an iterated flux match more tractable. The price is the closure approximation, whose adequacy in a pedestal Chapter 6 takes up against gyrokinetic CGYRO.

4.5.3 Saturation rules

The quasilinear weights of Equation (4.9) are set by linear physics, but the saturated intensity is the residue of the nonlinear cascade and cannot be recovered from a linear calculation, so a saturation rule must supply it. The simplest estimate is a mixing-length argument, in which an eddy of growth rate γ_k saturates when its shearing competes with its growth,

$$|\hat{\phi}(k_y)|^2 \sim \left(\frac{\gamma_k}{\langle k_\perp^2 \rangle} \right)^2 S(k_y), \quad \ell_{\text{mix}} \equiv \frac{\gamma_k}{\langle k_\perp^2 \rangle}, \quad (4.10)$$

with ℓ_{mix} the mixing length and $S(k_y)$ a spectral shape the bare estimate does not itself fix.

A TGLF saturation rule supplies $S(k_y)$, and it is *obtained by fitting* a closed-form expression to the saturated potential spectrum of a nonlinear CGYRO database run across a range of drives and geometries. It is thus the one part of the quasilinear flux that is empirical rather than derived. This work uses SAT3 throughout [10], fit to a database that includes electromagnetic modes.

The caveat that frames every edge TGLF evaluation here follows directly, because that database is drawn from the *core* rather than the pedestal [9, 10]. Applying the rule across the near-edge extrapolates its training domain into a region of steep gradients, high collisionality, and strong flow shear that the fit never saw. The framework accordingly runs TGLF in a dedicated edge configuration, uniformly across all cases, whose adequacy Chapter 5 tests and whose sensitivity to the saturation rule itself Chapter 6 examines.

4.5.4 Near-edge configuration

One setting is worth singling out, because it addresses a mixing-length pathology the steep pedestal exposes. Each mode's saturated intensity scales as a mixing length $\ell_{\text{mix}} \sim \gamma_k / \langle k_{\perp}^2 \rangle$ that diverges as $k_y \rightarrow 0$. In the core the long-wavelength growth rates are small and this limit is harmless, but in the sharp pedestal the ion-temperature-gradient branch stays unstable down to $k_y \rho_s \sim 0.07$, where the diverging mixing length inflates the lowest- k_y contribution into an unphysical flux runaway.

The physics absent from the bare closure is the zonal flow shear that caps the largest eddy scale and so imposes a spectral floor $k_{y,\text{min}}$. SAT3 carries that bound as an optional low- k_y cutoff, activated through the *sign* of the zonal flow coefficient, $\text{ALPHA_ZF} = -1$. Only the sign selects the cutoff, because the saturation amplitude enters as $|\text{ALPHA_ZF}|$, and with it active the spectral sum begins at $k_{y,\text{min}}$ rather than at the grid's lowest mode. The cutoff is retained for every case.

4.5.5 MHD ceiling

For inter-ELM H-mode prediction a peeling–ballooning ceiling prevents the transport-limited gradient from exceeding the Type-I ELM threshold. Rather than evaluating ideal MHD stability inline, the framework takes the maximum stable pedestal top pressure $p_{\text{ped}}^{\text{crit}}$ from EPED [13], in practice its neural-network surrogate EPED1NN [14], and couples it to the flux match as a stiff transport penalty. Where the pedestal top pressure overshoots the EPED value, the total modeled fluxes over the pedestal are inflated by

$$f_{\text{ELM}}(r) = 1 + \left(\max \left[0, \frac{p_{\text{ped}}}{p_{\text{ped}}^{\text{crit}}} - 1 \right] \right)^{p_{\text{ELM}}}, \quad (4.11)$$

where the exponent p_{ELM} dictates how hard the ELM overshoot penalty is applied. For the results presented in later chapters, $p_{\text{ELM}} = 10$ so as to ensure ELMs are strongly penalized if found. This artificially enhanced transport erodes the excess gradient and pushes the solution to marginal stability. At negative triangularity the penalty is inactive because the EPED1NN

surrogate is not trained for $\delta < 0$ (Section 5.6), but NT edges are known to be sub-ELM [56].

4.5.6 A minimal analytic alternative

The transport slot is interchangeable by design, so it was deemed worthwhile to test an inexpensive analytic model. A minimal turbulence model was built alongside the framework, assembling critical gradient expressions from the literature into the same three channels. It carries an ion-scale branch quenched by flow shear, an electron-scale branch, and a kinetic-ballooning branch supplied as a transport term rather than through an external stability call. The KBM branch is *not* a substitute for Section 4.5.5, since EPED fixes the pedestal by intersecting an ideal peeling–ballooning constraint with a KBM one and only the second has a counterpart here. The geometry is circular, the plasma pure deuterium, and nothing in it is fitted to the discharges of Chapter 5. A solution costs seconds. Appendix B collects the ingredients.

What such a model is good for is behavior rather than numbers, and two behaviors it returns bear on how the rest of this work should be read. With the shear term disabled, doubling the input power changes the profiles hardly at all, which is critical gradient stiffness acting as it should and a reminder that power alone is a weak lever on profiles in the presence of stiff transport. Enabling the shear at fixed power instead produces a pedestal, the ion-scale branch collapsing over the outer part of the domain, the KBM branch taking over as the limiting transport, and the pedestal stored energy rising several fold. The change appears on switching one term rather than on raising the drive, so an L–H-like transition is available to a model of this crudeness.

What that does and does not establish is worth separating. It says nothing about accuracy, because the model is not compared against a measurement anywhere and could not be trusted if it were. What it does show is that the *relationships* are reproducible from very little, since a flow shear quench of an ion-scale branch with a stiff ceiling waiting behind it is enough to generate a pedestal at fixed power and a handover between limiting modes. Producing that phenomenology is therefore not a capability that waits on transport-model fidelity, whereas getting a pedestal

height right plainly is. It also sets a modest expectation against which Chapter 5 can be read, since the TGLF-based configuration of this chapter carries both ingredients. Where the two disagree, the analytic model is the more likely of the two to be wrong, its quench acting at one representative wavenumber and omitting the toroidal rotation that Section 4.4 found carries much of the real shear. The results are accordingly stated at the level of which features appear rather than how large they are, and the shear implementation carries a defect recorded in Appendix B.4.

4.6 Flux matching

For each evolved channel c and radius r_j the transport flux is balanced against the volume-integrated net source, defining a residual

$$R_{j,c} = \frac{F_{j,c}^{\text{tr}} - F_{j,c}^{\text{target}}}{F_{j,c}^{\text{target}}}, \quad F_{j,c}^{\text{target}} = \frac{1}{V_j'} \int_0^{r_j} f_c^{\text{target}} V' dr, \quad (4.12)$$

over the channels $c \in \{Q_e, Q_i, \Gamma_e\}$. Normalizing each residual by its own target keeps channels whose targets approach zero, notably the particle flux, meaningfully represented, and is the edge analog of the convective flux and gyroBohm transformations used in core PORTALS [8]. The scalarized objective is the channel- and radius-averaged residual norm,

$$\xi = \frac{1}{N_c N_p} \|R\|_2. \quad (4.13)$$

The design variables are the interior gradient knots z_k , with the per-channel separatrix gradients z_{LCFS} held fixed at the boundary values of Section 4.2. Following PORTALS, the surrogates are fit not to raw profile values but to transformed, physically motivated dimensionless parameters, which carry the dominant turbulence dependence and aid the regression [5, 8]. The expensive transport model is replaced by Gaussian process surrogates that map the design variables to the fluxes and are refined as the optimization proceeds, so flux matching is achieved with few real transport evaluations.

The objective is partitioned deliberately. The flux misfit of Equation (4.13) is the data term on which the solver converges, while the separatrix-gradient priors enter only as a regularizer on the step and are tracked separately. The boundary treatment therefore acts as a regularized inversion rather than a reweighting of the flux match. The formulation is not tied to a particular optimizer, and can be solved by surrogate-accelerated Bayesian optimization as in PORTALS [5], or by a relaxed, damped Newton iteration as in TGYRO [7], for example.

4.7 Uncertainty quantification

A key distinction in the edge module is between two uncertainty sources that are easily conflated. The first is the surrogate posterior variance, which falls to near zero at trained points for a deterministic transport model. Injecting input uncertainty there would corrupt the training targets and blunt the surrogate, so the propagated band reaches only the objective. The second source, and the one the module propagates, is the covariance of the physical inputs at predict time.

The declared inputs in this work are the separatrix values $(n_e, T_e, T_i)_{\text{LCFS}}$ of Section 4.2 and the neutral source rate, at 10% and 20% relative uncertainty, respectively, and Figure 4.9 traces the route by which they and the two induced covariances below reach the objective. The impurity fraction and source fluxes are prescribed, so no uncertainty is injected on them. The separatrix gradients are determined from physically self-consistent relationships posed in the sheared spectral filament (SSF) model of Peret *et al.* [57]. SSF is used as a reduced model of how a/L_y depends on y , but not as a nominally correct value. Experimental separatrix gradients are used for the validation of Chapter 5, and their sensitivities are transferred as a logarithmic derivative,

$$\delta(a/L_y) = \frac{\partial \ln(a/L_y)_{\text{SSF}}}{\partial \ln T} \sigma_T (a/L_y)_{\text{nom}}, \quad (4.14)$$

where $(a/L_y)_{\text{nom}}$ is the nominal separatrix gradient in use and σ_T the relative temperature uncertainty. The SSF gradients are a near-exact power law in temperature, so the exponent

carries over to a different nominal while the prefactor would not.

The second induced covariance is the fit uncertainty of the mtanh backbone of Section 4.2. It is the parameter covariance Σ_θ over $\theta = \{A, \Delta_0, \delta, m, c, b\}$, formed from the fit Jacobian, regularized by the separatrix priors and capped at 10% per marginal. Propagating it through the corrector-projected profile Jacobian $R = \partial(a/L_y)/\partial\theta$ orders the directions by observable impact and only the leading mode per channel is retained. The corrector sets a/L_y at the interior knots, so R has near-zero rows there and the fit error band widens between knots. This source sets the widest bands in the cases examined, although the declared separatrix values and the induced gradients dominate the typical width. Knot placement is thus the binding lever on the extremes.

Let $\phi = f(\theta)$ be the map from inputs to the propagated quantity, either a profile field or a flux. To first order the delta method linearizes it,

$$\Sigma_\phi = J \Sigma_{\text{in}} J^\top, \quad J = \left. \frac{\partial f}{\partial \theta} \right|_{\theta_0}, \quad (4.15)$$

where J is evaluated at the nominal inputs θ_0 . The joint input covariance is carried as a low-rank factor $\Sigma_{\text{in}} = LL^\top$ with $L \in \mathbb{R}^{n \times k}$, and pushing each of its columns through the model gives $B = JL$ and hence

$$\Sigma_\phi = (JL)(JL)^\top = BB^\top, \quad (4.16)$$

so the covariance is assembled matrix-free from k tangents and the dense Jacobian is never formed.

Each column $b_j = J(Le_j)$ is obtained by a forward difference through the full model,

$$b_j = \frac{f(\theta_0 + hLe_j) - f(\theta_0)}{h}, \quad (4.17)$$

where h is the step in units of the input sigma. The baseline $f(\theta_0)$ is shared by every column, so a pass costs $k + 1$ real evaluations rather than the $2k$ a central difference would need. Jacobian

vector products are implemented for the differentiable nodes, but the real-model scan produced every result reported here. The pass is run once at the end of an optimization, because propagating through the full model at every optimizer step would be prohibitive.

A complication arises with the mtanh fit error columns, because the parameterizer conditions any perturbed θ before rebuilding. Samples are clipped back inside the parameter bounds and their amplitude is bisected until the reconstruction is feasible. The step is therefore shrunk until the conditioner is inactive, and the difference quotient is divided by the perturbation actually applied. Directions surviving at less than a tenth of their requested amplitude are dropped and reported, which underestimates that factor.

The transport and target fluxes are read from the same perturbed evaluations, so their correlation is preserved rather than quadrature-summed. Differencing the two factors gives the residual covariance Σ_r , which enters the objective as a Mahalanobis weighting,

$$\chi^2 = r^\top \Sigma_r^{-1} r, \quad (4.18)$$

where r is the residual vector at the flux-matching control points, tested against zero. A risk-aware, upper tail scalar is minimized alongside it. It is $\mu_J + k\sigma_J$ with $k = 2$, or the 95th percentile of a sampled objective distribution once σ_J exceeds a quarter of μ_J . The per-node standard deviations are written back into the profile fields, so every predicted quantity carries a band rather than a point value.

Bands on positive definite quantities are clipped at zero, because the parameterization itself clips y and a/L_y there. Fluxes are left unclipped, since a particle or heat flux is a physical outcome and not an artifact of the propagation. A band that reaches the floor marks where the linearization has run out of validity. For the widest cases that is much of the domain rather than a narrow tail. Quantifying the error of the linearization itself is not the intent of this exercise.




Figure 4.9. Route by which the declared separatrix values and the neutral source rate covariances reach the objective. The separatrix gradients and the mtanh fit covariance are induced. Every column is perturbed through the full chain in a single forward difference pass, and fluxes are read from the same perturbed evaluations.

Chapter 5

Validation

5.1 Validation approach

Each case study applies the framework of Chapter 4 to a measured discharge under a single, deliberately constrained protocol, so that the near-edge prediction is a test of the model rather than a fit to it. The set spans the transport-limited regimes of interest, containing an L-mode (DIII-D 128913, the “shortfall” discharge of Section 5.3), two H-modes at contrasting collisionality (Sections 5.4 and 5.5), a negative triangularity plasma (Section 5.6), and an Alcator C-Mod I-mode (Section 5.7). Each is evolved on the near-edge domain of Section 4.1, with n_e , T_e and T_i predicted at flux match knots spanning $r/a = 0.88$ – 0.97 , reconstructed by the offset-mtanh parameterization of Section 4.2, with Table 5.1 collecting the global parameters. Radii are quoted in the normalized minor radius r/a throughout, and values taken from other work in a flux coordinate are marked as such. The L-mode, negative triangularity and I-mode cases are drawn from the SMARTS reference-case collection [58], and the two H-modes are the paired-collisionality DIII-D discharges of the dedicated pedestal studies of [53].

Five discharges from two devices cover a modest span of the dimensionless space, therefore limiting what any cross-case conclusion can reach. Evaluated on the reconstructions at the pedestal top ($r/a = 0.90$), the set covers a factor of 2.5 in $\rho_\star = \rho_s/a$, from 9.1×10^{-4} in the L-mode to 2.2×10^{-3} in the I-mode, a factor of 12 in β_e , from 8×10^{-5} to 1.0×10^{-3} , and temperature ratios T_i/T_e between 0.76 and 1.6. Safety factor and magnetic shear cover

$q = 2.4\text{--}4.7$ and $\hat{s} = 2.1\text{--}3.6$. Collisionality is the one direction deliberately levered, the two H-modes being the end points of a dedicated scan spanning about an order of magnitude in pedestal top v_i^* (Section 5.4). That asymmetry matters twice over. It is what lets Section 6.4 fit a collisionality dependence at all, and it is why that fit cannot be separated from a T_i/T_e dependence covering barely a factor of two. A reactor pedestal likely sits roughly an order of magnitude below this set in ρ_* and at an even lower collisionality.

The transport model in every case is TGLF-SAT3 in its edge configuration for the turbulent flux, together with NEO for the neoclassical flux (Section 4.5). This baseline is the reference against which the alternatives developed in this work are subsequently weighed. Input uncertainties are propagated through the converged flux match by the input-covariance method of Section 4.7, so each predicted profile carries an error band. The declared covariance is 10% on each separatrix value and 20% on the neutral source rates. These are seemingly modest values yet they turn out to generate substantial uncertainties, which are discussed further in Chapter 6.

Boundary conditions and particle sources are taken from the experimental reconstruction and held fixed. The separatrix values and gradients of all three channels are frozen to the input profile (Section 4.2), and the toroidal rotation is extracted from the measured ω_0 . The main ion neutral source rate is set to the volume integral of the reconstructed edge particle source, which is the recycling term where the reconstruction provides it and otherwise the beam term (Section 4.3). This rate is *not* calibrated to the measured density, so the pedestal top density remains a predicted quantity and its agreement with experiment is a genuine test of the transport model. The impurity content is the reconstruction's own fraction f_Z carried with the predicted density (Section 4.3), so the separatrix Z_{eff} is an implied value and f_Z may already inherit errors. Table 5.2 lists the fixed boundary values and source rates. Alternative approaches to defining or solving for n_Z were considered, but were omitted from this work due to severe transport model sensitivity to n_Z and a/L_{nZ} .

Table 5.1. Global parameters of the five validation discharges: major radius R_0 , minor radius a , and the boundary elongation κ and triangularity δ . The DIII-D cases share an on-axis toroidal field $B_T \simeq 2$ T and plasma current $I_p \simeq 1$ MA; the Alcator C-Mod case runs at $B_T = 5.7$ T and $I_p = 1.1$ MA.

Case	Device	R_0 [m]	a [m]	κ	δ
L-mode (shortfall)	DIII-D	1.71	0.60	1.63	+0.44
H-mode, high- v^*	DIII-D	1.77	0.59	1.76	+0.41
H-mode, low- v^*	DIII-D	1.78	0.59	1.74	+0.41
Negative triangularity	DIII-D	1.84	0.55	1.44	-0.36
I-mode	C-Mod	0.68	0.22	1.55	+0.34

Table 5.2. Fixed inputs for each validation case (Section 5.1). The LCFS y , a/L_y , and Z_{eff} are frozen at the quantities defined by or derived from the case reconstructions. The main ion neutral source rate S_{n^0} is determined according to Section 4.3. The impurity is carbon on DIII-D and molybdenum on Alcator C-Mod. T_e at the separatrix is smoothed to a physical 50 eV for the L-mode and NT cases.

Case	LCFS value			LCFS a/L_y			S_{n^0}	Z_{eff}
	n_e [10^{19}]	T_e [eV]	T_i [eV]	n_e	T_e	T_i	[s^{-1}]	
L-mode	0.82	50	153	46	30	34	8.4×10^{20}	1.24
H, high- v^*	0.83	72	278	29	30	3	4.2×10^{20}	1.14
H, low- v^*	0.44	91	267	15	45	6	2.7×10^{20}	1.44
NT	2.17	50	125	24	90	21	3.0×10^{21}	1.42
I-mode	6.49	224	219	15	105	108	1.7×10^{21}	2.48

5.2 Results across the set

Three results emerge before the per-case detail, and all three are cross-case. The first is what the solve converges to, the second is what the resulting pedestal costs in pressure, and the third is which instabilities carry the flux, and together they frame everything that follows.

5.2.1 The flux match improves every case and closes none

The relative- L_2 flux residual falls in all five cases and reaches zero in none of them, which is the precondition for everything that follows. Figure 5.1 shows where the modeled fluxes land relative to target, in two columns per case, the flux at the initial measured profiles and the flux at the solver’s best evaluation. The reduction spans a factor of 1.3 in the negative triangularity case

to 49 in the I-mode, so the amount of “work” the solve does is itself strongly case dependent.

Where the match fails it fails in identifiable places rather than uniformly, at the outermost knot and in the particle channel. The L-mode is the one case whose particle channel is repaired in direction. It initially contains a modeled Γ_e pinch against an outward target at three of four knots, then runs outward at all four afterwards, though still over target by 10 to 59%. The heat channels exhibit varying degrees of matching across the cases and locations although no case showed $< 30\%$ of error relative to the heat flux targets at all knots.

The initial column is the more revealing half. Modeled transport at the measured profiles misses the power balance targets by up to a factor of 54 in the heat channels and 95 in the particle channel, so those profiles are nowhere near a flux matched state of the model. The direction of the miss is not a property of the regime alone as indicated by the low- v^* H-mode and the I-mode over-predictions, the L-mode under-prediction at every knot in both heat channels, and the high- v^* H-mode over-prediction through the interior while under-predicting at its outermost knot. The two cases that over-predict most severely are also the two carrying the largest ρ_* in the set (Section 5.7), an association five points cannot raise above a coincidence.

5.2.2 Pedestal top pressure

The single quantity that carries these edge predictions back to the reactor-scoping stakes of Section 1.2 is the pedestal top pressure p_{ped} . The stiff core inherits p_{ped} as its boundary condition, so a case matters chiefly insofar as it reproduces it. Table 5.3 reduces the five cases to that figure of merit, the total thermal pressure $p_{\text{ped}} = n_e(T_e + T_i)$ at $r/a = 0.90$, and it carries the central quantitative result of this work. Only the negative triangularity case lands inside the 10% tolerance that reactor scoping would require, at +4.9%. The error is otherwise large and two-sided, the L-mode running 17% high while the two H-modes fall 84 and 50% low and the I-mode 43% low. Across the set the RMS error in pedestal top pressure is 48%. Against the $P_{\text{fus}} \sim p_{\text{ped}}^2$ sensitivity of Section 1.2, an error of that size equates to a factor of two to four in inferred fusion power in either direction.




Figure 5.1. Modeled transport flux divided by its target, per channel (Q_e , Q_i , Γ_e) and per knot (r1 inner to r4 separatrix). Two columns per case give the initial (measured) profiles and the converged evaluation. Gray is a match, warm/red an over-prediction, cool/blue an under-prediction, × a raw particle flux running inward against an outward target, and Σ the column's relative- L_2 residual.

Notably, the error is not carried by the same channel in each case, and Figure 5.2 reveals the details that a single pedestal number conceals. The high- ν^* H-mode is missed chiefly in the density at -77% while its ion temperature is nearly recovered at -11% , and the low- ν^* H-mode is missed chiefly in the ion temperature at -70% against a density and electron temperature held to -15 and -19% . The I-mode is missed in the ion temperature at -53% against an electron temperature reproduced to -4% . Only the negative triangularity case is accurate in all three channels at once.

Flux matching moved the pedestal away from the measurement in every case. The error induced by the parameterizer reconstruction at the initial profiles, between $+16$ and -31% in p_{ped} , cannot be overstated, yet the converged solve worsens the p_{ped} error in all five cases. The I-mode is the sharpest instance, its reconstruction sitting within 1% of the measurement before

the solve and 43% below it after. Driving the residual down and predicting the pedestal are therefore not the same operation, which is the result this chapter hands to Chapter 6. The final residuals and pedestal errors carry a rank correlation of 0.8 over five cases, but with only five points that relationship is not distinguishable from chance.


Table 5.3. Predicted versus measured pedestal top pressure $p_{\text{ped}} = n_e(T_e + T_i)$ for the five validation cases, evaluated at $r/a = 0.90$. The relative error is the flux error normalized to the measured value, and the last column is the propagated input-covariance band of Section 4.7 on the prediction.

Case	$p_{\text{ped}}^{\text{meas}}$ [kPa]	$p_{\text{ped}}^{\text{pred}}$ [kPa]	error [%]	1σ [%]
L-mode	1.38	1.61	+16.9	13.3
H-mode, high- v^*	8.92	1.45	−83.8	29.2
H-mode, low- v^*	11.63	5.85	−49.7	13.7
Negative triangularity	6.67	6.99	+4.9	13.3
I-mode	42.17	24.16	−42.7	18.3
RMS relative error			48	

5.2.3 Turbulent fingerprints across the set

The third cross-case result is which instabilities carry the flux, read through the fingerprint paradigm of Section 2.3. Each per-case figure decomposes the transport into ion temperature gradient (ITG), trapped electron (TEM) and electron temperature gradient (ETG) branches together with the neoclassical contribution, and the pattern is largely common to the set.

TGLF returns a spectrum rather than a label, so a description of the attribution is warranted. Each wavenumber on the dominant eigenmode is weighted by its flux, so the major transporters set the branch names rather than the fastest growing roots. The sign of the mode frequency gives the propagation direction, ion-diamagnetic roots being assigned to ITG at any k_y . An electron-diamagnetic root is called ETG only where it satisfies the full electron-scale fingerprint, namely $k_y \rho_s \geq 2$, near-adiabatic ions in the sense $|Q_i|/|Q_e| \leq 0.15$ at that k_y , and a particle flux consistent with zero. Every remaining electron-direction root is called TEM, including at high k_y , because a wavenumber carrying particle flux is not ETG whatever its scale. The per-branch



gray = matched, warm = over-predicted, cool = under-predicted. $p_{\text{ped}} = n_e(T_e + T_i)$, and the column labels give its error.

Figure 5.2. Predicted pedestal top value relative to the measurement, $(\text{pred} - \text{meas})/\text{meas}$ in percent, at $r/a = 0.90$. Rows are the three predicted channels and the pressure $p_{\text{ped}} = n_e(T_e + T_i)$ they combine into, evaluated at the converged solve. Gray is a match, warm/red an over-prediction and cool/blue an under-prediction.

sums are signed and reported as computed, with no reallocation between branches, so a TGLF thermodiffusive Γ_e pinch on the electron branch appears under TEM.

Two limits belong with that reading. The labels name a spectral region rather than identify a mode in the gyrokinetic sense, since a single k_y may carry more than one root and only the dominant one is classified. The direction test also lumps the ubiquitous mode in with ITG by construction because it propagates in the *ion* diamagnetic direction [59] despite being a density-gradient-driven trapped-electron instability. The literature does not sharply distinguish the ubiquitous mode from the toroidal ITG mode in any case [60].

The common pattern is a handover with radius rather than a single dominant branch. In the high- ν^* H-mode and the negative triangularity case the ion heat flux is carried almost entirely by ITG through the interior, at 0.82 to 0.96 of Q_i at the inner knots, and passes to the trapped electron branch at the outermost knot, where ITG falls to zero and TEM carries 0.65 and 0.77. Electron-scale transport is negligible everywhere, ETG never exceeding 20% of Q_e in any

case and peaking in the interior rather than growing toward the separatrix. That decomposition qualitatively matches the published gyrokinetic picture for these pedestals [37, 53], in which ion-scale electrostatic modes dominate, neoclassical is significant but smaller, and ETG is small.

Two cases depart from the pattern in opposite directions. The L-mode and the I-mode show no ITG–TEM handover at all, the trapped electron branch instead carrying most of Q_i at most of the knots. The low- v^* H-mode converged state on the other hand loses most of its turbulent ion heat transport at the outermost knot, where neoclassical flux accounts for 94% of Q_i while the electron channel stays turbulent. This is a result of a growth in the density pedestal collocated with $E \times B$ suppression on the ion scale modes and clearly indicates residuals can be reduced by either turbulent or neoclassical transport, as expected.

In most cases, $E \times B$ shear cannot explain the handover from ITG to TEM. The shearing rate does rise across the transition in the H-modes, but the negative triangularity case shows the same handover at a shearing rate that never approaches the suppression threshold anywhere in the domain. What sets the branch, and whether the transition is physical or a property of the saturation model at steep gradient, cannot be decided from these runs.

5.2.4 Pedestal error does not track any dimensionless parameter

A natural question of a five-case set is whether the pedestal error organizes along any of the dimensionless parameters of Section 5.1. Testing that here entails correlating the signed pedestal top pressure error against each parameter in logarithm, with significance assessed by an exact permutation test over all 120 orderings of five cases.

The correlation coefficients are -0.38 for ρ_* , -0.49 for β_e , 0.75 for T_i/T_e , -0.53 for q , -0.26 for $\hat{s}$ and 0.27 for v_i^* , none of which is distinguishable from chance at any conventional level. The largest of them carries a permutation p of 0.13. Whatever sets the sign and size of the pedestal error in a given case is therefore not a smooth function of the local dimensionless state within this set.

This observation should be limited as a statement of absence of evidence rather than as

evidence of absence. Five cases sample roughly 2.7 independent directions in the six-parameter space, because the parameters are mutually correlated at up to $|r| = 0.84$, so a correlation with any one of them could not have been attributed to that parameter alone in any case. And with five points the smallest attainable two-sided permutation p is 0.017, which is above the threshold a six-parameter search would need, so this test could not have certified a positive result had one appeared. The useful reading is that no parameter stands out strongly enough to organize the set.

5.3 L-mode

The L-mode is the framework's most exposed test. Local gyrokinetic simulations of DIII-D 128913 under-predicted the power-balance energy flows by roughly a factor of seven at mid-radius [61]. That discrepancy was mostly closed with a revised saturation model [62]. The present case takes the same discharge further out in radius, with no pedestal ceiling active, so the near-edge gradients are set entirely by the transport model. Two inputs depart from the reconstruction, the separatrix electron temperature smoothed to a physical 50 eV from ≈ 35 eV and the neutral fueling fixed from a reconstructed wall-recycling source, the only case in the set for which the wall source is explicitly available.

The shortfall appears in its expected form at the measured profiles, and it has not gone away. The edge TGLF flux falls short of the power balance the same reconstruction demands by up to a factor of thirty, in both heat channels and at every knot, the ion channel running at 0.03 to 0.61 of target and the electron channel at 0.42 to 0.58. That the deficit survives under a saturation rule descended from the one that closed the original problem is a statement about where the two assessments are made rather than a contradiction of that work. That work scores predicted temperature profiles inward of $r/a = 0.98$, where pedestal top values are set by the integral of the gradient over most of the minor radius. The ratios here are local flux comparisons across the near edge. The particle channel fails differently, the modeled Γ_e running *inward* against an outward target at three of the four knots.

The solve does what such a residual invites, and Figure 5.3 shows what it costs. Steepening the gradients until the modeled flux rises to target drives the electron temperature 25% above the measurement at the pedestal top, leaving the pedestal top pressure +16.9% high while the residual falls from 0.621 to 0.194, the lowest surviving residual in the set. The particle channel is the one place the solve repairs something structural, Γ_e becoming outward at all four knots, though it lands 10 to 59% above target. The density and ion temperature move comparatively little, at +7% and under 1%, so the pressure error is carried almost entirely by T_e . The input covariance band of 13.3% on p_{ped} is comparable to that error, so the miss is not cleanly separable from input uncertainty in this case. Neoclassical transport is minor at every radius (Figure 5.4), and the ion heat flux is generally carried by TEM rather than ITG.

5.4 High-collisionality H-mode

The two H-modes probe the framework where a peeling–ballooning ceiling and a steep pedestal coexist, and isolate the collisionality dependence. They are the end points of a dedicated DIII-D scan in which the pedestal top ion collisionality was varied by about an order of magnitude, from $v_i^* \simeq 0.40$ here to $\simeq 0.040$ in Section 5.5. The two pedestal top pressures sit within 20% of each other [53]. That scan also reports a transport decomposition of the pedestal ion heat flux. Evaluated at the measured profiles at $\rho_{\text{tor}} = 0.94$, the local neoclassical flux carries 27% of the power-balance ion heat flux here and 18% in the low- v^* case, so turbulence carries the majority in both as in the previous work.

The steep pedestal here is where the edge low- k_y cutoff ($\text{ALPHA_ZF} = -1$, Section 4.5.4) does its work. With it active there is no near-separatrix runaway of Q_i , and the mode decomposition of Figure 5.6 shows ITG absent from the outermost knot altogether. The setting is therefore load-bearing exactly where the pedestal is steep. However, suppressing the runaway is not the same as matching the flux. The model over-predicts the heat channels through the interior, by factors of 1.6 to 3.0, and under-predicts at the outermost knot at 0.33 and 0.44, so a residual

built from terms of opposite sign can give the solver little coherent direction to move in.

The heat channels end largely matched, five of the eight knot-channels landing within 30%, and the density pedestal pays for it. Figure 5.5 shows the cost in the profiles. Convergence collapses the pedestal top density by -77% and the electron temperature by -46% , while the ion temperature is largely recovered (-11%). The parameterizer’s reconstruction already sat 31% below the experimental pedestal top pressure, so there is a conflation of error in this case. The propagated uncertainty band on T_e reaches 69% here, by far the widest in the set, so the electron channel of this case is also poorly determined by its own inputs.

5.5 Low-collisionality H-mode

The low- v^* counterpart to the previous section provides an interesting contrast. DIII-D 179444 carries a published CGYRO+NEO transport decomposition [53], against which the framework’s breakdown can be compared channel by channel. It is also potentially the most reactor-relevant case in the set as its collisionality approaches an ITER-relevant pedestal [63].

The modeled fluxes at the three inner knots are in some cases an order of magnitude beyond the targets. All three channels then collapse to between 0.05 and 0.14 of target at the outermost knot. A swing of nearly three orders of magnitude in the model-to-target ratio across the domain of a single discharge indicates an error dependence on radial variations of the plasma state.

The residual falls from 16.5 to 0.57, the largest surviving residual in the set, and Figure 5.7 shows what that costs. The solve brings the three inner knots to within a factor of a few, the ion channel closest at 0.74, 0.77 and 1.44 of target and the electron channel still short at 0.43 to 0.60, but it cannot move the outermost knot at all: there all three channels remain at 0.03 to 0.14 of target, essentially where they started. The pedestal that this buys is a density of 2.14 against a measured $2.51 \times 10^{19} \text{ m}^{-3}$ at the pedestal top, with the temperatures falling short by 19 and 70%, for a pedestal top pressure 49.7% below the measurement.

The mode decomposition locates the split cleanly (Figure 5.8). At the three inner knots the transport is turbulent and almost purely ion-directed, the ion-scale branch carrying 75 to 103% of every channel, the trapped-electron contribution reaching 23% only in the particle channel at the innermost knot. At the outermost knot the ion channel inverts, neoclassical transport supplying 94% of the ion heat flux, while the electron heat and particle fluxes stay largely turbulent at 19 and 26% neoclassical. The outermost interval is therefore not merely the worst matched, it is the one radius at which the turbulence model has stopped carrying the ion channel, and it is the interval Section 6.3.2 identifies as setting most of the pedestal height.

5.6 Negative triangularity

Negative triangularity accesses near-H-mode confinement, without ELMs, and with true TLPs in a regime truly above L-mode. It should therefore stress the geometric sensitivity of the transport model at $\delta < 0$. The ELM ceiling is inactive here because the EPED neural network is not trained for $\delta < 0$, and the separatrix electron temperature is smoothed to 50 eV from ≈ 5 eV as in the L-mode case.

This is the one case in the set that lands inside the reactor-scoping tolerance. Figure 5.9 shows the predicted T_e , T_i and n_e overlaying the measurement across the near edge, the pedestal top pressure reproduced to +4.9% with no channel missed by more than 18% (Table 5.3). The match itself is incomplete: the electron channel under-predicts at the two inner knots and over-predicts at the two outer, the ion channel falls to 0.53 at the outermost knot, and the residual moves only from 0.406 to 0.311, the smallest reduction in the set. So, the case that reproduces the pedestal best is also the case the solver moved least, its pedestal top pressure shifting by 1.2% between the initial reconstruction and the converged solve.

The mode decomposition of Figure 5.10 shows the composition changing across the domain rather than one branch carrying the edge. Both heat channels are ITG-dominated at the two inner knots and hand over to TEM outboard, until at $r/a = 0.97$ they are carried by

TEM almost entirely, the handover falling where the $E \times B$ shear passes through its minimum. Neoclassical transport is not negligible in the ion channel as it is in the L-mode, growing outward to roughly a quarter of Q_i at the outermost knot.

5.7 Alcator C-Mod I-mode

The I-mode is the most distinctive test in the set and the only one on a second device. Its defining feature is a temperature pedestal without a matching density pedestal. That decoupling probes whether the framework can predict a steep T_e and T_i edge atop an L-mode-like density, a configuration none of the DIII-D cases presents. It also tests cross-machine robustness at Alcator C-Mod's high field and metal wall, the impurity being molybdenum, carried as a fully ionized species in the transport model.

Figure 5.11 shows the decoupled pedestal reproduced in one of its two halves. The electron temperature is captured almost exactly, to -4% at the pedestal top, the density to -20% , the ion temperature to -53% , and the pedestal top pressure at -43% . The case therefore reproduces the electron half better than the ion half, which is the opposite of the high- v^* H-mode. The solve drives the column residual down from 16.0 to 0.325, a reduction of a factor of 49 and the largest in the set. The reconstruction reproduced p_{ped} to within 1% before the solve began, so every point of the final error was introduced by the flux match itself.

Unlike the other cases, TEM carries the majority of all three channels at every knot, as evident in Figure 5.12. The particle flux is the largest in the set, TEM-carried throughout. The $E \times B$ shear is the sharper distinction, staying under one c_s/a across all four knots and rising steeply only inside the last few percent of the domain, unlike the more inward variation seen in either H-mode. A temperature pedestal is therefore being predicted without the shear that regulates the other cases, which is consistent with the ion channel being the half the solve fails to reproduce. None of this is attributable to the change of device, because Alcator C-Mod's field is nearly triple that of DIII-D and its minor radius about a third, the two canceling in ρ_* and

leaving ρ_i as the differentiating factor as in the other cases.

DIII-D 128913 (L-mode / shortfall)




Figure 5.3. Converged near edge for the L-mode shortfall case (DIII-D 128913). *Top:* predicted (blue) against measured (gray) T_e , T_i , n_e with the input-covariance 90% band, open circles marking the flux match knots. *Middle:* the transport flux against the fixed target in each channel. *Bottom:* variance share by input source.

DIII-D 128913 (L-mode / shortfall) --- turbulent-flux mode breakdown




Figure 5.4. Decomposition of the transport flux for the L-mode case into its turbulent components and the neoclassical contribution from the TGLF spectra at the best evaluation. The top row is the $E \times B$ shear for reference against the turbulence activity. The TEM branch carries most of the ion heat flux at three of the four knots.

DIID-D 179449 (H-mode, high ν^*)




Figure 5.5. Converged near edge for the high-collisionality H-mode (DIID-D 179449), panels as in Figure 5.3. The modeled heat flux over-predicts through the interior and under-predicts at the outermost knot (Section 5.2.1), and the predicted n_e and T_e pedestals fall 77 and 46% below the measurement at the pedestal top.

DIII-D 179449 (H-mode, high ν^*) --- turbulent-flux mode breakdown




Figure 5.6. Turbulent and neoclassical flux decomposition for the high- ν^* H-mode, as in Figure 5.4. The top row is the $E \times B$ shear for reference against the turbulence activity. ITG is no longer observed at the outermost knot.

DIII-D 179444 (H-mode, low ν^*)




Figure 5.7. Converged near edge for the low-collisionality H-mode (DIII-D 179444), panels as in Figure 5.3. The solve brings the three inner knots to within a factor of a few and leaves the outermost at 0.03 to 0.14 of target, under-predicting n_e at the pedestal top by 15% while the temperatures fall short by 19 and 70%.

DIII-D 179444 (H-mode, low ν^*) --- turbulent-flux mode breakdown




Figure 5.8. Turbulent and neoclassical flux decomposition for the best evaluation of the low- ν^* H-mode. The top row is the $E \times B$ shear for reference against the turbulence activity. The three inner knots are carried almost entirely by ion-directed turbulence; at the outermost knot the ion heat flux collapses and neoclassical transport supplies 94% of what remains, while the electron channels stay turbulent.

DIII-D (negative triangularity)




Figure 5.9. Converged near edge for the negative triangularity case (DIII-D 193870), panels as in Figure 5.3. All three channels overlay the measurement. The input-covariance band for this case is computed with the SSF geometric factor $g = G_0 \rho_s / R_0$ floored to keep the LCFS a/L_y prior in its valid regime.

DIII-D (negative triangularity) --- turbulent-flux mode breakdown




Figure 5.10. Turbulent and neoclassical flux decomposition for the negative triangularity case, as in Figure 5.4. Both heat channels hand over from ITG to TEM between the inner and outer knots, and neoclassical transport retains a growing share of Q_i .

Alcator C-Mod (I-mode)




Figure 5.11. Converged near edge for the Alcator C-Mod I-mode, panels as in Figure 5.3. T_e is captured to -4% and n_e to -20% , but the decoupling fails in the ion channel, T_i being under-predicted by 53% and leaving the pedestal top pressure 43% low.




Figure 5.12. Turbulent and neoclassical flux decomposition for the I-mode, as in Figure 5.4. TEM carries the majority of all three channels at every knot, and the $E \times B$ shear stays below one c_s/a across the constrained radii.

Chapter 6

Error Analysis

6.1 Approach

Chapter 5 established some leading candidates for attributing the limits of this near-edge modeling effort, and they are not one fault in three guises. The declared input covariance, the objective together with its knot discretization, and the transport model each set their own floor on the pedestal, and only the last is a statement about the physics. The sections below bound them in that order, cheapest and most model-independent first, so that what survives at the end is attributable to the transport model rather than to the machinery around it. That the residual and the pedestal error are not necessarily coupled is what makes the decomposition necessary.

6.2 Limitations from the inputs alone

The first layer is the cheapest to bound and the most consequential for whether the workflow is usable. Each converged solve carries the input covariance of Section 5.1 propagated through it, over the quantities a user of the workflow would have to supply. The band that results is the precision the workflow could deliver if its transport model were exact.

That precision is not sufficient for the purpose, as Figure 6.1 shows in the 1σ band on the pedestal top for each case. The band on p_{ped} runs from 13.3 to 29.2%, and no channel in any case falls below 10.5%. Against the 10% per-channel tolerance that reactor scoping would require, the workflow is already outside specification before any transport model is evaluated.

Propagated through the $P_{\text{fus}} \sim p_{\text{ped}}^2$ sensitivity of Section 1.2, even the narrowest of these bands is a 27% band on inferred fusion power. A device study that took this workflow as its pedestal model would be choosing between designs it cannot distinguish.

The same band says something quite different about where the present error comes from. In four of the five cases the band is between 0.28 and 0.77 of the central pedestal error, so the misses of Table 5.3 are larger than input uncertainty alone can account for. The margin is thin at the upper end: in the L-mode the band reaches 0.77 of an already small error, so that case alone does not separate the two. The exception is instructive in the opposite direction, because the negative triangularity case, the only one inside tolerance, carries a band of 13.3%, or 2.7 times its nominal error.

Which inputs carry the bands is equally clear. The right panel of Figure 6.1 decomposes the variance by source. The three separatrix values and the separatrix gradient prior together carry 63 to 100% of it, and the parameterization fit supplies the remainder, reaching 37% in the high- ν^* H-mode and 23% in the I-mode against 1% in the negative triangularity case. Where the fit does contribute, it contributes through the radii it resolves least well between knots, which Section 6.3.2 shows to be most of the pedestal. Across the cases the dominant lever is the separatrix state, and a 10% separatrix uncertainty is already optimistic for any quantity at the last closed flux surface.

These are first-order propagations of a declared covariance through the converged solution with no transport model in them, so the claim is that the workflow's precision is inadequate on the most favorable accounting available. The mechanism appears to be the global mtanh structure and the inferred a/L_y peak in profiles with a clear pedestal, which is why the effect is not uniform across the set.




Figure 6.1. *Left:* 1σ band on n_e , T_e , T_i and p_{ped} at $r/a = 0.90$, with the 10% tolerance marked and the central |error| in p_{ped} as an open marker. *Right:* variance share by input.

6.3 The objective

The workflow assumes that reasonable kinetic profiles, inserted into the transport and target models, admit a flux matched solution. Chapter 5 showed that nominal flux matching was generally not observed: the solver minimizes competently and stops at a finite residual. The objective aggregates twelve residual components into one scalar, so once none of them is driven to zero, different combinations of non-zero components can score alike. The subsections below show that they do, that the knot discretization then selects among them, and that the separatrix gradient sets what the knots leave open.

6.3.1 Equal residuals, different pedestals

Every state the solver evaluated carries both a residual and a pedestal, so each solve leaves behind a sample of the objective's own level sets, as shown in Figure 6.2 by the vertically aligned samples. In the L-mode, which Section 5.2.1 showed achieving the best flux match in the set, two evaluations whose residuals differ by 0.9% return pedestal top pressure errors of +51.2 and +24.3%. The high- ν^* H-mode carries a pair agreeing to 1.3% in residual and differing by

25 points of pedestal, the I-mode a pair agreeing to four decimal places and differing by 14, and the low- v^* H-mode a pair matched to 0.4% sitting 17 points apart. The negative triangularity case lacked spread for relevance in this analysis.

These are evaluated states rather than converged solves, and the pairs sit above each case's best residual rather than at it. The claim is therefore the weaker one the argument needs: over the region the solver searches, the objective is uninformative about the pedestal, and no better optimizer repairs that. A global minimum that is at a truly flux matched state steps around this issue, but this requires components in the workflow to support such a solution.

Pinning the four density design variables to the reconstruction tests the same point by changing the objective instead of the starting point. It removes four of the twelve residual components and leaves the temperature channels to match against a density held at the measurement; compared on the heat channels alone, which is all the reduced problem carries, the match improves in four of the five cases by factors of 1.1 to 16, and degrades by a factor of 1.4 in the L-mode. The pedestal does not follow. The high- v^* H-mode moves from -83.8 to $+14.7\%$ and the low- v^* H-mode from -49.7 to -35.2% ; the L-mode and the I-mode barely move, from $+16.9$ to $+21.0\%$ and from -42.7 to -44.7% ; and the negative triangularity case, the one that had been inside tolerance, degrades from $+4.9$ to $+27.3\%$. The same quasilinear model is independently documented to under-predict the near-separatrix density gradient [11].

6.3.2 Most of the pedestal is never constrained

The design variables are inverse gradient scale lengths at four knots, and the profile is recovered by integrating inward from a pinned separatrix,

$$y(r) = y(1) \exp \int_r^1 \frac{a}{L_y} \frac{dr'}{a}, \quad (6.1)$$

so the map from design variables to the profile is strictly triangular and no design variable is local. The outermost knot then dominates because it carries the steepest gradients: measured on the




Figure 6.2. Every state evaluated during each production solve: relative- L_2 flux residual against signed pedestal top pressure error, marker fill light to dark with solve order, $\pm 10\%$ tolerance shaded. The connector marks the pair whose residuals agree to within 2% and whose pedestals differ most, annotated with the residual ratio and the gap. Ordinates are per panel. The negative triangularity case is omitted for want of spread.

reconstructions alone, the last interval accounts for 31 to 81% of the e-fold from the separatrix to the pedestal top, so one design variable per channel sets most of the pedestal height.

That interval lies outside the flux matched region entirely. The left panel of Figure 6.3 gives the share of the pedestal e-fold integral lying outside $r/a = 0.97$, measured inward from the innermost flux match knot at $r/a = 0.88$; on the pressure it runs from 50 to 68%. Half to two thirds of what the workflow predicts is never subject to the model being validated, which is the mechanism behind Section 6.3.1: solutions can agree on every constrained flux and still differ in the unconstrained portion of the profile.

Knot placement is therefore a choice of answer rather than a convergence detail. Four layouts, altering only the number of knots and the span they cover, were solved on both H-modes with everything else held equivalent to the base cases. The residuals reached span a factor of 1.07 on the high- ν^* case and 1.35 on the low- ν^* case, while the converged pedestal top pressure




Figure 6.3. *Left:* share of the e-fold integral $\ln[y(0.88)/y(1)]$ lying outside the outermost knot at $r/a = 0.97$, measured on the reconstructions alone. *Right:* four knot layouts solved on the two H-modes with everything else held fixed, bars giving the converged pedestal top pressure error and Σ the flux residual reached.

spans 42 and 38 percentage points respectively (right panel). No layout rescues the prediction: every one leaves both H-modes under-predicted and outside tolerance.

6.3.3 Interior gradients depend on the separatrix conditions

What the flux match leaves open, the boundary condition sets. The separatrix values and inverse gradient scale lengths are frozen to the reconstruction (Section 4.2), and because Equation (6.1) builds the profile inward from that boundary, the imposed $a/L_y(1)$ propagates into the pedestal through the parameterization alone.

The imposed electron density and temperature gradients were each set to 0.5 and 1.5 times the reconstructed value on the high- ν^* H-mode, one channel at a time, and Figure 6.4 shows the converged profiles against the measurement. The response is large and it is not confined to the channel that was moved. Halving the density gradient raises the converged pedestal top density by a factor of 5.8 and the electron temperature with it, taking the pedestal top pressure from 84% below the measurement to 45% above; increasing the same gradient by

half gives 51% below. Moving the temperature gradient instead moves the density by factors of 2.0 and 1.5 and the pressure to 67 and 70% below.

Two features matter more than the size. The nominal solve does not sit between the two perturbations of either channel but outside both brackets and worse than every endpoint, so the response is not monotonic in the imposed gradient. And all four perturbed solves reach a *lower* flux residual than the nominal, 0.26 to 0.46 against 0.47: perturbing the boundary in either direction on either channel improves the flux match and the pedestal together, which makes the nominal boundary an arbitrary point of the objective rather than a privileged one.

This is a sensitivity bracket on one discharge, not a design of experiments. Two of the three gradients were varied at two levels either side of nominal without replication, so no interaction is estimated and none could be. Each arm does change one number and only one, the separatrix values and $a/L_{Ti}(1)$ being frozen throughout.

6.4 The transport model

The bar the transport model has to clear is not agreement with the measurement. Published nonlinear CGYRO calculations on the two H-mode discharges reach roughly 60% of the power balance ion heat flux at $\rho_{\text{tor}} = 0.94$ [53], so first-principles gyrokinetics is itself short by a factor of 1.7 at the radius where Section 6.3.2 located most of the pedestal's leverage. A reduced model calibrated to nonlinear gyrokinetics inherits that gap before any error of its own is counted, and no work on it recovers what its reference does not have.

The layer under test is the saturation rule, the one semi-empirical element of the quasilinear flux [9, 10] (Section 4.5.3): the linear physics fixes the quasilinear weights, but the saturated amplitude is read from a fit to nonlinear CGYRO, and that fit is a core one. The two layers separate cleanly: the weights, growth rates and every geometry scalar are bit-identical between SAT2 and SAT3 on the same input, so only the saturated intensity differs.

A dedicated study of that layer was run over the five validation discharges at eleven

near-edge radii each, of which 54 carry a power-balance target, taken as that balance minus the local neoclassical flux. Model fluxes are TGLF's own, including the per-mode quasilinear adjustment SAT3 applies differently to the particle and energy blocks, and every flux is rebuilt from the quasilinear weights and verified against that output to one part in 10^{14} . Define the per-channel flux ratio and a normalized collisionality $\hat{v}$,

$$r_y = \frac{y^{\text{mod}}}{y^{\text{tar}}}, \quad y \in \{Q_e, Q_i, \Gamma_e\}, \quad \hat{v} = \frac{v_e}{k_{y,\text{max}} v_{\text{zf}}}, \quad (6.2)$$

where $\hat{v}$ is the electron collision rate over the zonal-flow mixing rate, both from TGLF's output.

Model and target agree in sign at all 54 points for Q_e and Q_i , and at 39 of them for Γ_e ; the fifteen disagreements are a model-inward flux against an outward target, and at fourteen of them the bare quasilinear weight still changes sign across k_y , so the particle channel is a shape failure at those radii rather than an unusable target. Over the sampled range the model error is monotonic in $\hat{v}$ in all three channels,

$$\frac{d \log_{10} r_y}{d \log_{10} \hat{v}} = -0.87 (Q_e), \quad -1.30 (Q_i), \quad -0.85 (\Gamma_e), \quad (6.3)$$

at R^2 of 0.63, 0.66 and 0.65. Panel (a) of Figure 6.5 shows the three ratios and their fitted power laws over a factor of 250 in $\hat{v}$; the heat-pair composite $R = (r_{Q_e} r_{Q_i})^{1/2}$ used by the remaining panels gives -1.08 . Every trend crosses unity, so the discharges separate along it rather than clustering on one side.

A missing intensity factor $I \sim \hat{v}^a$ produces a model-error slope of exactly $-a$, so Equation (6.3) is a direct measurement of the missing exponent, and two further routes agree. A saturation rule reconstructed from scratch against the bare quasilinear weights fits $a = 0.77 \pm 0.03$, and a plain mixing-length intensity with *zero* fitted parameters beats the existing rule by a factor of 3.5 out of sample. The particle channel supplies the third route: Γ_e carries a different quasilinear operator and a different power-balance target from Q_e but returns the same slope

to within 0.02. Three constructions recovering the same quantity with the same sign is what raises this from a correlation to a candidate root cause, and the physical reading is direct: higher collisionality damps zonal flows, weakening the shear that regulates the turbulence and raising the saturated intensity [9].

Correcting three channels with one scalar is licensed by the error covariance, whose leading principal component aligns with the common mode $(1, 1, 1)$ to a direction cosine of 0.992. Each channel prefers its own exponent, at 0.87, 1.30 and 0.85, but a single value costs under 4% of RMSE out of sample and at case level the three slopes lie within about one standard error. Because $\hat{v}$ multiplies the intensity point by point, it comes out of the k_y and mode sums exactly and can be compiled into the rule and tested live. Panel (b) of Figure 6.5 scores that rule, SAT3 multiplied by $\hat{v}^{1.025}$ and nothing else, against the same targets. Every case moves toward unity, the over-predicted low-collisionality end and the under-predicted high-collisionality end at once, which no amplitude-only change can do. Across the database the log-scatter falls from 0.81 to 0.46 dex and the fraction of points within a factor of two rises from 24 to 59%.

Three limits make this a leading hypothesis rather than a settled correction. The database holds five discharges, and 78% of the variation in $\hat{v}$ is between them rather than within, so the regression is closer to a five-point fit than the point count suggests. The ion-to-electron temperature ratio is a named alternative these data cannot separate from collisionality. The two discharges with more than a decade of internal $\hat{v}$ track it at $r = -0.98$, while two of the remaining three carry the opposite within-case sign over much shorter levers. And no nonlinear gyrokinetic validation has been performed, since only a nonlinear run constrains the saturated intensity.

The correction addresses the cross-case level and leaves the radial shape, the error Chapter 5 identified as binding, as panel (c) of Figure 6.5 shows. Within either H-mode the model-to-target ratio spans some six hundred across eleven radii before the correction and roughly twenty after, concentrated beyond $r/a \simeq 0.92$ where the two outermost flux match knots sit. Part of that is unreachable in principle, on an arithmetic rather than a fitted bound. A saturation rule

supplies one positive intensity spectrum shared by all channels, so any intensity-only correction, this one and any refit of it included, multiplies every channel alike and cannot move the contrasts between them. The common mode carries 93% of the three-channel log-error variance for the heat pair and 77% across all three, leaving floors of 0.23 and 0.39 dex against the 0.42 to 0.60 dex achieved per channel, so roughly half of what remains in the heat channels is structural. Reaching it needs a term acting on the shape of $I(k_y)$ rather than its amplitude; a factor-of-three redistribution in k_y would suffice at 51 of 54 points, but finding the covariate that prescribes it is what five discharges cannot do.

Those same radii are where the model is hardest to evaluate reliably, and that coincidence is the most damaging single fact in this chapter. A convergence study over the L-mode and both H-modes at $r/a = 0.85$ and 0.95 found four of six points converged to within 10%, and the other two are model faults rather than discretization error: one straddles the threshold at which the SAT3 branch zeroes the intensity of every wavenumber below the ion gyroradius scale, and the other carries a single eigenvalue at the top of the wavenumber grid that runs away with basis size, rising from 3.5 to 2.4×10^3 while its neighbors sit near 3.4. Sweeping the resolution controls one at a time over all five discharges at $r/a = 0.90$ and 0.95 shows why neither is repairable. The poloidal basis size never settles, drifting steadily below `NBASIS_MAX` = 20 and admitting a spurious high- k instability at and above it, and the radial grid cannot be refined at all: its dimension is fixed at `NXGRID` = 32 by a compile-time parameter, with any larger request silently reset to the default of 16 rather than clamped. Behind the residual scatter is the *selection* of the mode width rather than the eigenvalue solve: growth rates vary by under 0.3% across the entire sweep while the width chosen for the flux-carrying low- k modes varies by a factor of three, and the flux follows it. The collisionality trend survives, spanning a factor of 250 in $\hat{\nu}$ and dwarfing these shifts, but the residual radial structure of panel (c) cannot be read as wholly physical, a fourth error source absent from the input covariance of Section 6.2. An accurate near-edge quasilinear flux is limited by the stability of its own mode selection before it is limited by the physics of its saturation rule.




Figure 6.4. Converged near-edge profiles of the high- v^* H-mode with the imposed LCFS $a/L_y(1)$ set to 0.5, 1.0 and 1.5 times the reconstructed value, one channel at a time, against the measurement: density gradient moved (*left*), electron temperature gradient moved (*right*). Rows are the three predicted channels, the vertical rule marks $r/a = 0.90$, and column headers give the pedestal top pressure error. Every other boundary quantity is frozen.




Figure 6.5. (a) Channel flux ratios against normalized collisionality, 54 near-edge radii over five discharges, with a fitted power law per channel; color is the channel, marker shape the discharge, and Γ_e is plotted only where model and target agree in sign. (b) Median heat-pair error per discharge before and after the correction, evaluated in the compiled rule. (c) The same ratio against radius for the two H-modes, the shaded band marking the two outermost flux match knots.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

This thesis set out to predict self-consistent, transport-limited kinetic profiles across the near edge by flux matching. That target was not reached, and the account of why is the substance of what follows. Chapter 4 assembled the workflow: an mtanh parameterization, neutral transport, impurity charge states, rotation shear, target fluxes, neoclassical transport, and turbulent transport, with the residual between modeled and integrated fluxes minimized over the pedestal region. Chapter 5 exercised it on five discharges from two devices. It predicts a transport-limited pedestal for the negative triangularity discharge at +4.9% error in pedestal top pressure, while the remaining four are missed by 17 to 84%, two-sided, for a root-mean-square of 48%. The capability is real but case-dependent, and no property of the five cases yet predicts which side of that divide a new case falls on. Propagating the declared input covariance through the converged solves gives a 1σ band of 13 to 29% on pedestal top pressure, carried between 63 and 100% by the separatrix conditions and already exceeding the desired tolerance for reactor scoping.

Chapter 6 attributes that spread, and the first finding reframes the rest. Flux matching by optimization, posed as it is for core transport solutions, does not reliably determine the near-edge pedestal. Nominal flux matching was generally not observed, so twelve non-zero residual components combine into one scalar in more than one way: two states of the L-mode

whose residuals agree to 0.9% return pedestal top pressure errors 27 percentage points apart, and the low- v^* H-mode carries a pair matched to 0.4% whose pedestals sit 17 points apart (Section 6.3.1). The mechanism is measurable on the reconstructions alone, since 50 to 68% of the pedestal pressure e-fold lies outside the outermost flux match knot and is never constrained (Section 6.3.2). What fills that majority is discretization and boundary condition rather than data: knot placement and the imposed separatrix gradient each move the pedestal top pressure by more than the difference between the two H-mode discharges, and neither is a transport model error (Sections 6.3.2 and 6.3.3).

Transport model evaluations on the initial, experimental states appear to lack a collisionality dependence. Measured three independent ways, then verified by compiling into the quasilinear turbulence saturation rule, the empirical correction moves the ends of the dataset toward agreement at once (Section 6.4). Its sign and existence appear firm, but the magnitude is a first calibration on five discharges, so it cannot be separated from an ion-to-electron temperature ratio dependence. Without more data, this is only a well-corroborated hypothesis and still fails to address the radial structure of the flux errors.

Taken together, these findings cast doubt on the idea of carrying a local, core flux matching formulation into the pedestal. A handful of coarse knots per channel with the profile integrated inward from a pinned boundary relies on accuracy where the transport model is least trustworthy both physically and numerically and the profiles vary most nonlinearly. The workflow is therefore best read as a diagnostic of where near-edge reduced modeling breaks rather than as a pedestal predictor.

7.2 Future directions

Several pieces of near-edge physics are absent and the conclusions above are conditional on their absence: momentum is supplied rather than evolved, the closure is local at every radius, and every case is a single stationary operating point, so nothing here tests a transition, an

ELM cycle or any transient. More importantly, the findings order the work that follows them. Parameterization, knot layout and solver design are all choices made against the response the modules underneath present, so if those modules do not resemble the physics at the radii carrying the pedestal, the choices are being optimized against the wrong surface. The transport model is upstream of everything else here, and little is gained by tuning the layers above it first.

What appears necessary is a monolithic or hybrid transport model that remains valid across the pedestal region at least to the separatrix. It must be nonlocal, because the expansion the local reduction rests on erodes across the near edge as the turbulent eddy size and the ion banana width approach the gradient scale length (Section 2.3). It must remain predictive across the a/L_y peak, since that peak sits outside the outermost radius the present model is trusted at and is consequently inferred by the parameterization rather than evaluated, which is precisely the unconstrained majority of Section 6.3.2. And it must carry impurities across charge states rather than as a trace correction, because the ionization source that closes the electron particle channel depends on the charge state distribution and its radial variation. The same requirement applies to the neutrals, whose penetration is not measured on any of these discharges but inferred from a reduced model feeding the density target directly. Extending such a model through the scrape-off layer to the divertor target would go further and remove the fixed separatrix boundary altogether.

It remains unclear whether a function that encodes physical profile shapes in the solution, e.g., mtanh , is the best approach to the parameterization. Restating the solution in terms of the function's variables instead of fitting the mtanh backbone and correcting with a spline may be a worthwhile step, but remains useless without a transport model that “requests” realistic profiles. Alternatively, one may ask if any shape prior is needed at all. A flux-driven solve evolves the profile forward under the transport operator until it is stationary, so the shape is an output of the solve and nothing about it need be encoded in advance. The parameterization question is in that sense downstream of the solver question.

Appendix A

Gaussian-Process Flux Surrogates

The flux matching solver of Section 4.6 evaluates the transport response many times, and TGLF and NEO are far too expensive to call at every step. Each is therefore replaced by a Gaussian-process (GP) surrogate that maps the design variables to the fluxes, is refined as the solve proceeds, and returns a calibrated predictive uncertainty. This appendix summarizes the GP machinery as it is used here, following Rasmussen and Williams [64].

A.1 The surrogate problem

Let $\mathbf{x}$ collect the design variables that determine the local transport state at a flux matching radius, the inverse gradient scale lengths and the transformed dimensionless parameters $\{\hat{v}_{ei}, T_i/T_e, \beta_e, \omega_0/c_s\}$ on which the surrogates are trained [8]. Let $y = F(\mathbf{x})$ be a scalar gyroBohm flux, with one GP per channel. The map F is expensive, deterministic up to solver tolerance, and near-singular in the stiff region close to the turbulence threshold. A GP surrogate is cheap to evaluate, it smooths the flux response near threshold, and it supplies the predictive variance that the uncertainty quantification of Section 4.7 propagates.

A.2 Gaussian-process regression

A GP places a prior over functions, $F(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$, specified by a mean m and a covariance kernel k . Given training inputs $X = \{\mathbf{x}_i\}_{i=1}^N$ and observations $\mathbf{y}$, the posterior at




Figure A.1. The GP surrogate on a schematic stiff flux response. (a) Posterior mean and $\pm 2s$ band after four transport calls, with the interval of a/L_{T_i} over which the band admits the flux match target marked. (b) The same after four further calls placed where s was largest. (c) The ARD relevances $1/\ell_d$. All three panels are schematic.

a test point $\mathbf{x}_\star$ is Gaussian with mean and variance

$$\mu(\mathbf{x}_\star) = m(\mathbf{x}_\star) + \mathbf{k}_\star^\top (K + \sigma_n^2 I)^{-1} (\mathbf{y} - \mathbf{m}), \quad (\text{A.1})$$

$$s^2(\mathbf{x}_\star) = k(\mathbf{x}_\star, \mathbf{x}_\star) - \mathbf{k}_\star^\top (K + \sigma_n^2 I)^{-1} \mathbf{k}_\star, \quad (\text{A.2})$$

where $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$, $[\mathbf{k}_\star]_i = k(\mathbf{x}_i, \mathbf{x}_\star)$, and σ_n^2 is the observation-noise variance. The posterior mean Equation (A.1) is the surrogate flux, and the posterior variance Equation (A.2) grows away from the training data, which is the signal the active-learning solver uses when siting the next transport call.

Figure A.1 shows what that buys. With four transport calls the $\pm 2s$ band is wide enough that the flux match target is admitted over a broad interval of the driving gradient, so the residual does not yet locate a solution. Four further calls sited where s was largest collapse that interval by more than an order of magnitude, which is why the stiff threshold is resolved in far fewer evaluations than a fixed design would need.

Kernel.

A stationary automatic-relevance-determination (ARD) kernel is the default,

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp \left[-\frac{1}{2} \sum_d \frac{(x_d - x'_d)^2}{\ell_d^2} \right], \quad (\text{A.3})$$

whose per-dimension length scales ℓ_d let the surrogate learn which parameters the flux actually depends on, which is a useful diagnostic of the transport sensitivities in its own right. The near edge is not statistically stationary, because the flux is flat below threshold and stiff above it, so a single length scale is a compromise. Nonstationary kernels, either an input-dependent length scale or a monotonic warping of the gradient inputs, are an option where the stiffness is strong, at the cost of extra hyperparameters and training data.

Mean function.

A constant or low-order mean suffices when the training data are dense. A physics-informed mean, for example a critical gradient form that already encodes the threshold and stiff rise, reduces the burden on the kernel and improves extrapolation into sparsely sampled regions of the edge.

A.3 Training

The hyperparameters $\boldsymbol{\theta} = \{\sigma_f, \ell_d, \sigma_n\}$ are set by maximizing the log marginal likelihood,

$$\log p(\mathbf{y} \mid X, \boldsymbol{\theta}) = -\frac{1}{2} \tilde{\mathbf{y}}^\top (K + \sigma_n^2 I)^{-1} \tilde{\mathbf{y}} - \frac{1}{2} \log |K + \sigma_n^2 I| - \frac{N}{2} \log 2\pi, \quad (\text{A.4})$$

with $\tilde{\mathbf{y}} = \mathbf{y} - \mathbf{m}$, which balances data fit against model complexity automatically.

A subtlety specific to this application is the meaning of σ_n . The transport model is deterministic, so its *training* noise is essentially zero, and the small jitter retained is only for numerical conditioning of the Cholesky factorization, so σ_n here is not a physical uncertainty. The physical uncertainty enters separately, at predict time, and the two must be kept distinct,

because conflating a near-zero training noise with the input uncertainty would misstate the band.

A.4 Uncertainty: two separate sources

The predicted flux carries two independent uncertainties that the framework keeps apart. The first is the GP’s own predictive variance Equation (A.2), an epistemic term reflecting finite training data. It is largest where the solve has not yet sampled and it drives the active learning. The second is the propagation of input uncertainty, in the boundary conditions, source rates, and prescribed impurity content, through the surrogate at predict time. For an input with mean $\bar{\mathbf{x}}$ and covariance Σ_x , a first-order propagation gives

$$\mathbb{E}[F] \approx \mu(\bar{\mathbf{x}}), \quad \text{Var}[F] \approx s^2(\bar{\mathbf{x}}) + \nabla \mu(\bar{\mathbf{x}})^\top \Sigma_x \nabla \mu(\bar{\mathbf{x}}), \quad (\text{A.5})$$

the first term being the surrogate’s epistemic variance and the second the propagated input variance. For the squared-exponential kernel the moments can also be computed in closed form by exact moment matching, when the linearization is inadequate. Separating the three, a near-zero training noise, an epistemic s^2 from data coverage, and predict-time input propagation, is what lets the framework report a distribution over the converged profiles rather than a single point estimate (Section 4.7).

A.5 Validation of the predictive variance

A surrogate is only useful if its reported uncertainty is trustworthy. Holding out a fraction of the training data and checking how often the held-out fluxes fall within the predicted bands tests whether s^2 is over- or under-confident before the surrogate is relied upon in the flux match.

Appendix B

A Minimal Near-Edge Transport Model

The framework of Chapter 4 is built to carry a high-fidelity transport model, and it inherits that model's cost and its opacity together. A second, deliberately crude model was developed alongside it to establish what near-edge phenomenology follows from a handful of analytic critical-gradient formulae, and it is that model the qualitative results of Section 4.5.6 come from. This appendix records its ingredients at the level needed to read those results. It is a collaborative proof of concept developed with C. Holland, whose unpublished model notes are the fuller record, and the summary here is deliberately partial.

The design intent is worth stating first, because it explains the omissions. Every flux is a closed-form expression in the local plasma state, so a solution costs seconds rather than hours, and every term is inspectable. Nothing in it is fitted to the discharges of Chapter 5. What the model is asked to reproduce is therefore behavior and not numbers, and it should be read as a statement about which physics is sufficient for a given qualitative feature rather than as a prediction.

B.1 Equations and geometry

The same steady state, flux surface averaged balance of Section 2.1 is solved on $0.85 \leq x \leq 1$ with $x = r/a$, for the electron density and the electron and ion pressures,

$$\langle \mathbf{\Gamma}_e \cdot \nabla r \rangle = \frac{1}{V'} \int_0^r \langle S_{\text{wall}} \rangle V' dr, \quad (\text{B.1})$$

$$\langle \mathbf{q}_e \cdot \nabla r \rangle = \frac{1}{V'} \int_0^r [f_e \langle P_{\text{in}} \rangle - \langle P_{ei} \rangle - \langle P_{\text{rad}} \rangle] V' dr, \quad (\text{B.2})$$

$$\langle \mathbf{q}_i \cdot \nabla r \rangle = \frac{1}{V'} \int_0^r [(1 - f_e) \langle P_{\text{in}} \rangle + \langle P_{ei} \rangle] V' dr, \quad (\text{B.3})$$

with P_{in} the power crossing the inner boundary from the core, split by f_e between the channels, P_{ei} the collisional exchange, and S_{wall} a wall particle source of prescribed exponential form $S_{\text{wall}} \propto \exp[(x - 1)/\lambda_{\text{wall}}]$. The geometry is concentric circular at fixed aspect ratio with a parabolic q profile, so no equilibrium reconstruction is required and a device is specified by R_0 , a , B_T , I_p and κ alone. Fluxes are gyroBohm-normalized throughout, lengths to a and times to a/v_{ref} .

B.2 Fluxes

Each channel is the sum of a neoclassical and three turbulent contributions,

$$F = F^{\text{nc}} + F^{\text{ITG}} + F^{\text{ETG}} + F^{\text{KBM}}, \quad F \in \{\Gamma, Q_e, Q_i\}, \quad (\text{B.4})$$

with each turbulent term written in critical-gradient form through $Z(u) = \max(u, 0)$. The three branches are chosen to span the fingerprint signatures of Section 2.3, an ion-scale electrostatic branch carrying ion heat, an electron-scale branch carrying electron heat and no particles, and an electromagnetic branch carrying all three [23].

Neoclassical.

The Hinton–Hazeltine fluxes are evaluated in the fitted form collected by Belli and Candy [22], retaining the full collisionality dependence through v_{*e} and v_{*i} rather than the banana-regime limit.

Ion temperature gradient.

The ITG branch carries the model’s only flow-shear physics and is the term the results of Section 4.5.6 turn on. A linear growth rate is written against a critical gradient,

$$\gamma_{\text{lin}} = |\omega_r| Z \left(\sqrt{R/L_{T_i} - R/L_{T_i}^{\text{crit}}} \right), \quad \omega_r = -k_y \rho_s c_s / R, \quad (\text{B.5})$$

with $R/L_{T_i}^{\text{crit}}$ from the Guo–Romanelli threshold, evaluated at a single representative $k_y \rho_s = 0.3$. Flow shear enters by quenching that rate,

$$\gamma_{\text{eff}} = Z(\gamma_{\text{lin}} - \gamma_E), \quad \gamma_E = \left| \frac{dV_{*i}}{dr} \right|, \quad V_{*i} = \frac{1}{Z_i e n_e B} \frac{dp_i}{dr}, \quad (\text{B.6})$$

so the shearing rate is the gradient of the ion diamagnetic velocity and the toroidal rotation is absent by construction. That is the model’s sharpest departure from the framework of Section 4.4, where the measured toroidal term was found to carry a large share of the shear and could not be dropped, and it should be kept in view when reading anything the shear drives. Writing V_{*i} in normalized form gives $V_{*i}/v_{\text{ref}} = \rho_* p_i^{-1} dp_i/dx$, the inverse pressure rather than the inverse density setting the scale, and $\rho_* \propto \sqrt{T_i}$ sits inside the radial derivative. Both matter in a pedestal, because T_i falls by an order of magnitude across the domain. The saturated diffusivity follows a standard quasilinear estimate,

$$\chi_i^{\text{ITG}} = \frac{k_y^2}{k_\perp^4} \frac{\gamma_{\text{eff}}^3}{\omega_r^2 + \gamma_{\text{eff}}^2}, \quad (\text{B.7})$$

whose four factors are a gyroBohm scaling, a radial correlation length, a weak-to-strong transport transition, and the increase of transport with drive. The electron thermal response is $\chi_e^{\text{ITG}} =$

$f_t \chi_i^{\text{ITG}}$ with f_t the trapped fraction. The particle flux is written in the quasilinear decomposition of Angioni et al. [60], a diagonal diffusion set by a/L_{n_e} , a thermodiffusive term set by a/L_{T_e} , and a pure convection independent of any gradient, with the latter two estimated from that work's parametric dependences on T_i/T_e and collisionality. An inward particle flux is therefore available without an inward heat flux, and its sign can reverse with the turbulent regime rather than being fixed by construction.

Electron temperature gradient.

A slab ETG electron heat flux is taken from the analytic Hatch formula [45], driven by η_e above a threshold near unity. It is assigned no particle and no ion energy flux, which is the ETG fingerprint used in Section 5.2.3.

Kinetic ballooning.

The KBM branch is a stiff diffusivity switching on above a critical normalized pressure gradient,

$$D^{\text{KBM}} = C^{\text{KBM}} \frac{\rho_i^2 v_{\text{th},i}}{a} Z(\alpha - \alpha_c), \quad \alpha = q^2 \beta R_0 / L_p, \quad (\text{B.8})$$

following Guttenfelder [65], with $\chi^{\text{KBM}} = \frac{3}{2} D^{\text{KBM}}$ so that all three channels are driven together. It is what limits the pedestal once the ion-scale branch is quenched.

The correspondence with Section 4.5.5 is partial, and the part that is missing matters. EPED locates a state marginally stable to *both* ideal peeling–ballooning and kinetic ballooning, and it is the intersection of those two constraints that fixes pedestal height and width together [13]. The branch above reproduces only the second. Nothing here carries the current-driven peeling physics, the equilibrium current profile is not evolved, and no ELM cycle is represented, so a pedestal this model returns is limited by a gradient-driven transport threshold alone and may sit above where an ideal stability calculation would put it.

B.3 Boundary conditions and solution

Separatrix values and gradients are supplied rather than solved, as in Section 4.2. The values come from a two-point model [35] and the gradients from the Eich–Manz separatrix operational space [66], which sets λ_{pe} , λ_{ne} and λ_{Te} as functions of the turbulence parameter α_t and the poloidal sound radius. This is the same class of closure as the sheared spectral filament model of Section 4.7. Main-ion neutrals follow a steady state diffusive balance in the manner of Burrell [47], closed by a specified particle confinement time.

Each channel’s inverse gradient scale length is represented by an Akima spline through a handful of knots spanning the domain, chosen over a cubic spline because it does not overshoot and over a piecewise-linear representation because the flow-shear terms of Equation (B.6) require a continuous $d(a/L_y)/dx$. That requirement is what later motivated the analytic backbone of Section 4.2, which serves the same purpose with fewer degrees of freedom. The knot values are advanced by a simple relaxation on the flux mismatch,

$$(a/L_y)_j^{i+1} = (a/L_y)_j^i \left[1 + \alpha \left(1 - \frac{P_{\text{model},j}^i}{P_{\text{source},j}^i} \right) \right], \quad (\text{B.9})$$

with i the iteration, j the knot and α a relaxation rate. No Jacobian is formed, so the iteration is cheap and finds a local minimum whose identity depends on the initialization, which is the same sensitivity Section 6.3.1 measures for the full framework.

B.4 Standing limitations

Four limitations bound what the model can be asked to do. The geometry is circular, so nothing in it responds to shaping and the negative triangularity behavior of Section 5.6 is out of reach. The plasma is pure deuterium at $Z_{\text{eff}} = 1$, so there is no dilution, no impurity radiation and no counterpart to Section 4.3. The flow shear of Equation (B.6) carries the ion diamagnetic term alone, so a case whose shear is substantially toroidal is misrepresented. And the current-driven

half of the pedestal stability problem is absent altogether (Appendix B.2), so a pedestal the model returns is limited by a gradient threshold with no ideal ceiling above it.

Nothing here is compared against a measurement, which bounds the use more sharply than any of those. The model can show that a relationship between ingredients holds, and it cannot show that a number is right.

A last limitation is specific to the present implementation rather than to the formulation. The shearing rate as coded divides the ion pressure gradient by the density rather than by the pressure, and omits the radial derivative of ρ_* from dV_{*i}/dr . Both departures scale with T_i , which falls by an order of magnitude across this domain, so they distort the radial shape of γ_E and suppress it toward the separatrix. The qualitative results reported in Section 4.5.6 are stated at a level that survives this, since they concern whether a feature appears at all rather than where or how large it is, but no quantitative use should be made of the shear until it is re-derived.

Bibliography

- [1] International Energy Agency. *World Energy Outlook*. IEA, Paris. 2024. URL: <https://www.iea.org/reports/world-energy-outlook-2024> (visited on 08/04/2026).
- [2] J. P. Freidberg. *Plasma Physics and Fusion Energy*. Cambridge University Press, 2007.
- [3] J. Wesson. *Tokamaks*. 4th ed. International Series of Monographs on Physics. Oxford University Press, 2011.
- [4] N. T. Howard, P. Rodriguez-Fernandez, J. Hall, M. Muraca, A. Saltzman, A. Ho, J. C. Hillesheim, A. J. Creely, T. H. Eich, T. Body, P. B. Snyder, and C. Holland. “Performance and Transport in the ARC Tokamak”. In: *Journal of Plasma Physics* 92.3, E67 (2026). Part of the ARC Fusion Power Plant Physics Basis collection. DOI: 10.1017/S0022377826101573.
- [5] 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”. In: *Journal of Plasma Physics* (2020). ISSN: 14697807. DOI: 10.1017/S0022377820001075.
- [6] A. J. Cerfon and J. P. Freidberg. ““One size fits all” Analytic Solutions to the Grad–Shafranov Equation”. In: *Physics of Plasmas* 17.3 (Mar. 2010), p. 032502. ISSN: 1070-664X. DOI: 10.1063/1.3328818.
- [7] J. Candy, C. Holland, R. E. Waltz, M. R. Fahey, and E. Belli. “Tokamak Profile Prediction Using Direct Gyrokinetic and Neoclassical Simulation”. In: *Physics of Plasmas* 16.6 (2009). ISSN: 1070664X. DOI: 10.1063/1.3167820.
- [8] P. Rodriguez-Fernandez, N. T. Howard, A. Saltzman, S. Kantamneni, J. Candy, C. Holland, M. Balandat, S. Ament, and A. E. White. “Enhancing Predictive Capabilities in Fusion Burning Plasmas through Surrogate-Based Optimization in Core Transport Solvers”. In: *Nucl. Fusion* (2024).

- [9] G. M. Staebler, J. E. Kinsey, and R. E. Waltz. “A Theory-Based Transport Model with Comprehensive Physics”. In: *Physics of Plasmas* 14.5 (May 1, 2007), p. 055909. DOI: 10.1063/1.2436852.
- [10] H. G. Dudding, F. J. Casson, D. Dickinson, B. S. Patel, C. M. Roach, E. A. Belli, and G. M. Staebler. “A New Quasilinear Saturation Rule for Tokamak Turbulence with Application to High Beta Plasmas”. In: *Nuclear Fusion* 62.9 (Sept. 1, 2022), p. 096005. DOI: 10.1088/1741-4326/ac7a4d.
- [11] C. Angioni, T. Gamot, G. Tardini, E. Fable, T. Luda, N. Bonanomi, C. K. Kiefer, and G. M. Staebler. “Confinement Properties of L-mode Plasmas in ASDEX Upgrade and Full-Radius Predictions of the TGLF Transport Model”. In: *Nuclear Fusion* 62.6 (June 1, 2022). ISSN: 17414326. DOI: 10.1088/1741-4326/ac592b.
- [12] A. Ashourvan and J. Candy. “Breakdown of Quasilinear Theory in the Tokamak Edge”. In: *Physical Review Letters* 132.20 (May 13, 2024), p. 205101. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.132.205101.
- [13] P. B. Snyder, R. J. Groebner, A. W. Leonard, T. H. Osborne, and H. R. Wilson. “Development and Validation of a Predictive Model for the Pedestal Height”. In: *Physics of Plasmas* 16.5 (2009). ISSN: 1070664X. DOI: 10.1063/1.3122146.
- [14] O. Meneghini, S. P. Smith, P. B. Snyder, G. M. Staebler, J. Candy, E. Belli, L. Lao, M. Kostuk, T. Luce, T. Luda, J. M. Park, and F. Poli. “Self-Consistent Core-Pedestal Transport Simulations with Neural Network Accelerated Models”. In: *Nuclear Fusion* 57.8 (Aug. 1, 2017), p. 086034. DOI: 10.1088/1741-4326/aa7776.
- [15] J. F. Parisi, D. R. Hatch, P. Y. Li, J. W. Berkery, A. O. Nelson, S. M. Kaye, K. Imada, and M. Lampert. *Turbulent Transport-Limited Pedestals in Tokamaks*. May 14, 2025. DOI: 10.48550/arXiv.2505.09101. arXiv: 2505.09101 [physics]. URL: <http://arxiv.org/abs/2505.09101> (visited on 06/09/2025). Pre-published.
- [16] T. Slendebroek, A. Nelson, O. Meneghini, G. Dose, A. Ghiozzi, J. Harvey, B. Lyons, J. McClenaghan, T. Neiser, D. Weisberg, M. Yoo, E. Bursch, and C. Holland. “Exploring the Fusion Power Plant Design Space: Comparative Analysis of Positive and Negative Triangularity Tokamaks through Optimization”. In: *Nuclear Fusion* 66.2 (Feb. 1, 2026), p. 026032. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/ae27e6.
- [17] A. C. D. Hoffmann, T. N. Bernard, M. Francisquez, G. W. Hammett, A. Hakim, J. Boedo, R. Rizkallah, C. K. Tsui, and t. T. team. *Towards Fully Predictive Gyrokinetic Full-f Simulations*. Oct. 13, 2025. DOI: 10.48550/arXiv.2510.11874. arXiv: 2510.11874 [physics]. URL: <http://arxiv.org/abs/2510.11874> (visited on 11/18/2025). Pre-published.

- [18] A. Di Siena, C. Bourdelle, A. Bañón Navarro, G. Merlo, T. Görler, E. Fransson, A. Polevoi, S. H. Kim, F. Koechl, A. Loarte, E. Fable, C. Angioni, P. Mantica, and F. Jenko. *First Global Gyrokinetic Profile Predictions of ITER Burning Plasma*. Nov. 5, 2025. DOI: 10.48550/arXiv.2511.03336. arXiv: 2511.03336 [physics]. URL: <http://arxiv.org/abs/2511.03336>. Pre-published.
- [19] 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. 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”. In: *Nuclear Fusion* 47.6 (June 1, 2007). ISSN: 00295515. DOI: 10.1088/0029-5515/47/6/S02.
- [20] F. Troyon, R. Gruber, H. Saurenmann, S. Semenzato, and S. Succi. “MHD-Limits to Plasma Confinement”. In: *Plasma Physics and Controlled Fusion* 26.1A (Jan. 1, 1984), pp. 209–215. DOI: 10.1088/0741-3335/26/1A/319.
- [21] M. Greenwald. “Density Limits in Toroidal Plasmas”. In: *Plasma Physics and Controlled Fusion* 44.8 (Aug. 1, 2002), R27–R53. DOI: 10.1088/0741-3335/44/8/201.
- [22] E. A. Belli and J. Candy. “Kinetic Calculation of Neoclassical Transport Including Self-Consistent Electron and Impurity Dynamics”. In: *Plasma Physics and Controlled Fusion* 50.9 (Sept. 1, 2008), p. 095010. DOI: 10.1088/0741-3335/50/9/095010.
- [23] 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’”. In: *Nuclear Fusion* 59.9 (July 12, 2019). ISSN: 17414326. DOI: 10.1088/1741-4326/ab1fa2. arXiv: 1903.09994.
- [24] A. J. Brizard and T. S. Hahm. “Foundations of Nonlinear Gyrokinetic Theory”. In: *Reviews of Modern Physics* 79.2 (Apr. 1, 2007), pp. 421–468. DOI: 10.1103/RevModPhys.79.421.
- [25] J. Candy, E. A. Belli, and R. V. Bravenec. “A High-Accuracy Eulerian Gyrokinetic Solver for Collisional Plasmas”. In: *Journal of Computational Physics* 324 (Nov. 2016), pp. 73–93. ISSN: 0021-9991. DOI: 10.1016/j.jcp.2016.07.039.
- [26] ITER Organization. *Divertor*. ITER Organization. URL: <https://www.iter.org/machine/divertor> (visited on 08/30/2026).

- [27] D. A. D’Ippolito, J. R. Myra, and S. J. Zweben. “Convective Transport by Intermittent Blob-Filaments: Comparison of Theory and Experiment”. In: *Physics of Plasmas* 18.6 (June 2011), p. 060501. ISSN: 1070-664X. DOI: 10.1063/1.3594609.
- [28] S. I. Krasheninnikov. “On Scrape off Layer Plasma Transport”. In: *Physics Letters A* 283.5-6 (May 1, 2001), pp. 368–370. DOI: 10.1016/S0375-9601(01)00252-3.
- [29] S. C. Molesworth, J. A. Boedo, C. K. Tsui, R. Perillo, and D. L. Rudakov. “Filamentary Velocity Scaling Validation and Spin Dynamics in the DIII-D Tokamak”. In: *Physics of Plasmas* 31.4 (Apr. 2024), p. 042515. ISSN: 1070-664X. DOI: 10.1063/5.0195025.
- [30] J. W. Connor, R. J. Hastie, H. R. Wilson, and R. L. Miller. “Magnetohydrodynamic Stability of Tokamak Edge Plasmas”. In: *Physics of Plasmas* 5.7 (July 1, 1998), pp. 2687–2700. DOI: 10.1063/1.872956.
- [31] R. J. Goldston. “Heuristic Drift-Based Model of the Power Scrape-off Width in Low-Gas-Puff H-mode Tokamaks”. In: *Nuclear Fusion* 52.1 (Jan. 2012), p. 013009. ISSN: 0029-5515. DOI: 10.1088/0029-5515/52/1/013009.
- [32] 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”. In: *Nuclear Fusion* 53.9 (Sept. 2013), p. 093031. ISSN: 0029-5515. DOI: 10.1088/0029-5515/53/9/093031.
- [33] B. D.udson, M. V. Umansky, X. Q. Xu, P. B. Snyder, and H. R. Wilson. “BOUT++: A Framework for Parallel Plasma Fluid Simulations”. In: *Computer Physics Communications* 180.9 (Sept. 2009), pp. 1467–1480. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2009.03.008.
- [34] C. S. Chang, S. Ku, G. R. Tynan, R. Hager, R. M. Churchill, I. Cziegler, M. Greenwald, A. E. Hubbard, and J. W. Hughes. “Gyrokinetic Projection of the Divertor Heat-Flux Width from Present Tokamaks to ITER”. In: *Nuclear Fusion* 57.11 (Nov. 2017), p. 116023. ISSN: 0029-5515. DOI: 10.1088/1741-4326/aa7efb.
- [35] P. Stangeby. *The Plasma Boundary of Magnetic Fusion Devices*. CRC Press, Jan. 1, 2000. ISBN: 978-0-367-80148-9. DOI: 10.1201/9780367801489.
- [36] F. Nespoli, H. Bufferand, M. Valentinuzzi, N. Fedorczak, G. Ciraolo, E. Serre, Y. Marandet, R. Maurizio, H. De Oliveira, B. Labit, M. Komm, M. Faitsch, and S. Elmore. “Application of a Two-Fluid Two-Point Model to SolEdge2D-EIRENE Simulations of TCV H-mode Plasma”. In: *Nuclear Materials and Energy* 18 (Jan. 2019), pp. 29–34. ISSN: 23521791. DOI: 10.1016/j.nme.2018.11.026.

- [37] R. J. Groebner and S. Saarelma. “Elements of H-mode Pedestal Structure”. In: *Plasma Physics and Controlled Fusion* 65.7 (July 1, 2023). ISSN: 13616587. DOI: 10.1088/1361-6587/acd478.
- [38] M. Miller, J. Hughes, T. Eich, G. Tynan, P. Manz, T. Body, D. Silvagni, O. Grover, A. Hubbard, A. Cavallaro, M. Wigram, A. Kuang, S. Mordijck, B. LaBombard, J. Dunsmore, and D. Whyte. “Determination of Confinement Regime Boundaries via Separatrix Parameters on Alcator C-Mod Based on a Model for Interchange-Drift-Alfvén Turbulence”. In: *Nuclear Fusion* 65.5 (May 1, 2025), p. 052002. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/adc9c3.
- [39] S. Mordijck. “Overview of Density Pedestal Structure: Role of Fueling versus Transport”. In: *Nuclear Fusion* 60.8 (Aug. 1, 2020), p. 082006. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/ab8d04.
- [40] M. Fenstermacher, L. Baylor, E. De La Luna, M. Dunne, G. Huijsmans, A. Kirk, F. Laggner, T. Osborne, C. Paz-Soldan, S. Saarelma, P. Snyder, E. Viezzer, M. Becoulet, K. Burrell, A. Cathey, X. Chen, M. Hoelzl, J. Hughes, R. Maingi, A. Nelson, H. Urano, E. Wolfrum, X. Xu, A. Diallo, L. Frassinetti, S. Futatani, L. Gil, R. Groebner, T. Happel, S. Kim, J. King, B. Labit, P. Lang, Y. Liu, Z. Liu, R. Lunsford, G. Park, U. Sheikh, W. Suttrop, B. Vanovac, R. Wilcox, A. Wingen, and T. Zhang. “Progress in Pedestal and Edge Physics: Chapter 3 of the Special Issue: On the Path to Tokamak Burning Plasma Operation”. In: *Nuclear Fusion* 65.5 (May 1, 2025), p. 053001. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/adbf3.
- [41] F. Wagner, G. Becker, K. Behringer, D. Campbell, A. Eberhagen, W. Engelhardt, G. Fussmann, O. Gehre, J. Gernhardt, G. v. Gierke, G. Haas, M. Huang, F. Karger, M. Keilhacker, O. Klüber, M. Kornherr, K. Lackner, G. Lisitano, G. G. Lister, H. M. Mayer, D. Meisel, E. R. Müller, H. Murmann, H. Niedermeyer, W. Poschenrieder, H. Rapp, and H. Röhr. “Regime of Improved Confinement and High Beta in Neutral-Beam-Heated Divertor Discharges of the ASDEX Tokamak”. In: *Physical Review Letters* 49.19 (Nov. 1, 1982), pp. 1408–1412. DOI: 10.1103/PhysRevLett.49.1408.
- [42] K. H. Burrell. “Role of Sheared $e \times B$ Flow in Self-Organized, Improved Confinement States in Magnetized Plasmas”. In: *Physics of Plasmas* 27.6 (June 1, 2020). ISSN: 10897674. DOI: 10.1063/1.5142734.
- [43] Y. R. Martin, T. Takizuka, and ITPA CDBM H-mode Threshold Database Working Group. “Power Requirement for Accessing the H-mode in ITER”. In: *Journal of Physics: Conference Series* 123 (July 2008), p. 012033. ISSN: 1742-6596. DOI: 10.1088/1742-6596/123/1/012033.

- [44] P. H. Diamond, S.-I. Itoh, K. Itoh, and T. S. Hahm. “Zonal Flows in Plasma—a Review”. In: *Plasma Physics and Controlled Fusion* 47.5 (May 1, 2005), R35–R161. DOI: 10.1088/0741-3335/47/5/R01.
- [45] D. Hatch, M. Kotschenreuther, P.-Y. Li, B. Chapman-Oplopoiou, J. Parisi, S. Mahajan, and R. Groebner. “Modeling Electron Temperature Profiles in the Pedestal with Simple Formulas for ETG Transport”. In: *Nuclear Fusion* 64.6 (June 1, 2024), p. 066007. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/ad3972.
- [46] R. J. Groebner, T. H. Osborne, A. W. Leonard, and M. E. Fenstermacher. “Progress in Quantifying the Edge Physics of the H Mode Regime in DIII-D”. In: *Nuclear Fusion* 41.12 (Dec. 2001), pp. 1789–1802. ISSN: 0029-5515. DOI: 10.1088/0029-5515/41/12/306.
- [47] K. H. Burrell. “NEUCG: A Transport Code for Hydrogen Atoms in Cylindrical Hydrogenic Plasmas”. In: *Journal of Computational Physics* 27.1 (Apr. 1978), pp. 88–102. ISSN: 00219991. DOI: 10.1016/0021-9991(78)90027-X.
- [48] H. P. Summers. *The ADAS User Manual, Version 2.6*. University of Strathclyde. 2004. URL: <https://www.adas.ac.uk> (visited on 08/04/2026).
- [49] G. S. Voronov. “A Practical Fit Formula for Ionization Rate Coefficients of Atoms and Ions by Electron Impact: $Z = 1-28$ ”. In: *Atomic Data and Nuclear Data Tables* 65.1 (Jan. 1997), pp. 1–35. ISSN: 0092-640X. DOI: 10.1006/adnd.1997.0732.
- [50] D. Reiter, M. Baelmans, and P. Börner. “The EIRENE and B2-EIRENE Codes”. In: *Fusion Science and Technology* 47.2 (Feb. 2005), pp. 172–186. ISSN: 1536-1055. DOI: 10.13182/FST47-172.
- [51] F. Sciortino, T. Odstrčil, A. Cavallaro, S. Smith, O. Meneghini, R. Reksoatmodjo, O. Linder, J. D. Lore, N. T. Howard, E. S. Marmar, and S. Mordijck. “Modeling of Particle Transport, Neutrals and Radiation in Magnetically-Confining Plasmas with Aurora”. In: *Plasma Physics and Controlled Fusion* 63.11 (Nov. 1, 2021), p. 112001. ISSN: 0741-3335, 1361-6587. DOI: 10.1088/1361-6587/ac2890. arXiv: 2106.04528 [physics].
- [52] O. Sauter, C. Angioni, and Y. R. Lin-Liu. “Neoclassical conductivity and bootstrap current formulas for general axisymmetric equilibria and arbitrary collisionality regime”. In: *Physics of Plasmas* 6.7 (1999), pp. 2834–2839. DOI: 10.1063/1.873240.
- [53] S. R. Haskey, A. Ashourvan, S. Banerjee, K. Barada, E. A. Belli, A. Bortolon, J. Candy, J. Chen, C. Chrystal, B. A. Grierson, R. J. Groebner, F. M. Laggner, M. Knolker, G. J. Kramer, M. R. Major, G. McKee, G. M. Staebler, Z. Yan, and M. A. Van Zeeland. “Ion Thermal Transport in the H-mode Edge Transport Barrier on DIII-D”. In: *Physics of Plasmas* 29.1 (Jan. 1, 2022). ISSN: 10897674. DOI: 10.1063/5.0072155.

- [54] G. W. Hammett and F. W. Perkins. “Fluid Moment Models for Landau Damping with Application to the Ion-Temperature-Gradient Instability”. In: *Physical Review Letters* 64.25 (June 1990), pp. 3019–3022. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.64.3019.
- [55] G. M. Staebler, J. E. Kinsey, and R. E. Waltz. “Gyro-Landau Fluid Equations for Trapped and Passing Particles”. In: *Physics of Plasmas* 12.10 (Oct. 2005), p. 102508. ISSN: 1070-664X. DOI: 10.1063/1.2044587.
- [56] A. O. Nelson, L. Schmitz, T. Cote, J. F. Parisi, S. Stewart, C. Paz-Soldan, K. E. Thome, M. E. Austin, F. Scotti, J. L. Barr, A. Hyatt, N. Leuthold, A. Marinoni, T. Neiser, T. Osborne, N. Richner, A. S. Welanders, W. P. Wehner, R. Wilcox, T. M. Wilks, and J. Yang. “Characterization of the ELM-free Negative Triangularity Edge on DIII-D”. In: *Plasma Physics and Controlled Fusion* 66.10 (Sept. 5, 2024), p. 105014. DOI: 10.1088/1361-6587/ad6a83.
- [57] M. Peret, J. Guterl, T. N. Bernard, F. D. Halpern, and N. Fedorczak. “Predictive Turbulence-Driven Flux Model of Scrape-off Layer Widths across Confinement Regimes in Tokamaks”. In: *Nuclear Fusion* 65.5 (May 1, 2025), p. 056043. ISSN: 0029-5515, 1741-4326. DOI: 10.1088/1741-4326/adcf37.
- [58] SciDAC SMARTS Project. *SMARTS Reference Cases*. 2025. URL: <https://github.com/scidac/SMARTS-reference-cases> (visited on 07/17/2026).
- [59] B. Coppi and F. Pegoraro. “Theory of the “Ubiquitous” Mode”. In: *Nuclear Fusion* 17.5 (1977), p. 969. DOI: 10.1088/0029-5515/17/5/009.
- [60] C. Angioni, E. Fable, M. Greenwald, M. Maslov, A. G. Peeters, H. Takenaga, and H. Weisen. “Particle Transport in Tokamak Plasmas, Theory and Experiment”. In: *Plasma Physics and Controlled Fusion* 51.12 (Nov. 10, 2009), p. 124017. DOI: 10.1088/0741-3335/51/12/124017.
- [61] C. Holland, A. E. White, G. R. McKee, M. W. Shafer, J. Candy, R. E. Waltz, L. Schmitz, and G. R. Tynan. “Implementation and Application of Two Synthetic Diagnostics for Validating Simulations of Core Tokamak Turbulence”. In: *Physics of Plasmas* 16.5 (2009). ISSN: 1070664X. DOI: 10.1063/1.3085792.
- [62] G. M. Staebler, J. M. Park, E. Hassan, C. Angioni, E. Fable, C. Bourdelle, J. E. Kinsey, C. Holland, E. A. Belli, T. Neiser, J. Candy, and R. E. Waltz. “Successful Prediction of Tokamak Transport in the L-mode Regime”. In: *Nuclear Fusion* 64.8 (July 2, 2024), p. 085002. ISSN: 0029-5515. DOI: 10.1088/1741-4326/ad5a1b.
- [63] L. Frassinetti, D. King, S. Saarelma, B. Labit, S. Blackmore, D. Keeling, C. Perez von Thun, C. Giroud, S. Wiesen, A. Kappatou, N. Vianello, E. Alessi, M. Brix, I. S. Carvalho, P. Carvalho, A. Chomiczewska, J. S. Elmore, M. Fontdecaba, E. Giovannozzi, J. Harrison,

- S. Hendersson, D. Kos, E. Kowalska, K. Imada, M. Lennholm, M. Maslov, A. Meigs, S. Menmuir, A. Merle, R. B. Morales, H. Nyström, E. Pawelec, G. Pucella, R. Scannell, D. Silvagni, A. Stagni, P. Ryan, and H. J. Sun. “Effect of Density on Peeling Limited Pedestals in JET-ILW, MAST-U and TCV and Implications for ITER Pedestal Predictions”. In: *Nuclear Fusion* 66.5 (Apr. 22, 2026), p. 056028. DOI: 10.1088/1741-4326/ae5c1c.
- [64] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press, 2006. ISBN: 026218253X.
- [65] W. Guttenfelder, J. Candy, S. M. Kaye, W. M. Nevins, E. Wang, R. E. Bell, G. W. Hammett, B. P. LeBlanc, D. R. Mikkelsen, and H. Yuh. “Electromagnetic Transport from Microtearing Mode Turbulence”. In: *Physical Review Letters* 106.15 (Apr. 14, 2011), p. 155004. DOI: 10.1103/PhysRevLett.106.155004.
- [66] T. Eich and P. Manz. “The Separatrix Operational Space of ASDEX Upgrade Due to Interchange-Drift-Alfvén Turbulence”. In: *Nuclear Fusion* 61.8 (Aug. 1, 2021). ISSN: 17414326. DOI: 10.1088/1741-4326/ac0412.