

A Bayesian Framework for Multi-Diagnostic Inference of Tokamak Divertor States

Daniel Michael Greenhouse

Doctor of Philosophy

University of York

School of Physics, Engineering and Technology

September 2025

Abstract

A critical challenge for future nuclear fusion tokamak power plants is the controlled management of heat and particle exhaust. The divertor, a specialised region within the tokamak, is responsible for managing this exhaust. However, the divertor plasma exhibits complex, multi-dimensional dynamics involving various species. No single diagnostic can provide the required comprehensive experimental insight.

This thesis outlines and assesses developments to a divertor multi-instrument Bayesian analysis system (D-MIBAS) designed to infer, given available measured experimental data, which plasma states were plausible. These plasma states (two-dimensional profiles of electron temperature, electron density, and neutral hydrogen density) offer a deep insight into divertor plasma dynamics and enable direct comparison between interpretive simulations and experiments. D-MIBAS uses a *mesh-based* approach, spatially aligned to the magnetic equilibrium, to provide a natural way to describe the plasma state and bring together different diagnostics and prior physics knowledge within a single, consistent framework.

Through synthetic and experimental validation using the MAST-U tokamak, this work demonstrates the benefit of combining multiple diagnostics and embedding physics information within the D-MIBAS mesh-based framework. This represents a significant advance: a physics-informed machine learning framework to reliably infer divertor plasma states in a non-machine, non-scenario specific manner. Ultimately, this can inform both the design and control of future fusion tokamak exhaust systems.

Author's Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as references.

Chapter 4 has led to the following publication:

D Greenhouse et al. *Two-dimensional inference of divertor plasma characteristics: advancements to a multi-instrument Bayesian analysis system*. 2025. DOI 10.1088/1361-6587/adab1b .

Signature:

A handwritten signature in black ink, appearing to read 'D Greenhouse', written in a cursive style.

Date: September 14, 2025

“To assume nothing is to assume something.”

Dedication

This work is dedicated to the ongoing pursuit of controlled nuclear fusion energy. Controlled nuclear fusion has the potential to provide abundant, low-climate-impact power worldwide, significantly contributing to the end of energy poverty. While the research is immensely fascinating, we must not lose sight of this crucial ambition.

Acknowledgements

The work presented in this thesis is supported by the patience, insights, and kindness of many generous people. I will always be immensely grateful to all of you.

It has been extremely enjoyable, informative, and inspiring to be part of the divertor physics integrated data analysis group at MAST-U. Led by Professor Bruce Lipschultz, alongside Dr. Chris Bowman, Dr. Kevin Verhaegh, and (soon to be Dr.) Nicola Lonigro, the discussions have been, and continue to be, simply fascinating. I would like to express my immense gratitude to the group; their guidance and collaboration enabled the success of this work. The framework, D-MIBAS, outlined in this thesis builds on the pioneering work of Dr. Bowman and the team in:

C Bowman et al. *Development and simulation of multi-diagnostic Bayesian analysis for 2D inference of divertor plasma characteristics* 2020. DOI 10.1088/1361-6587/ab759b.

The basis for the Balmer-line emission models incorporating molecular contributions was produced by Dr. Verhaegh. The model stems from BaSPMI analysis introduced in:

K Verhaegh et al. *A novel hydrogenic spectroscopic technique for inferring the role of plasma-molecule interaction on power and particle balance during detached conditions* 2020. DOI 10.1088/1361-6587/abd4c0.

I am also grateful for valuable contributions to the team over the years from Dr. James Harrison, Dr. Ben Dudson, Professor Chris Ridgers, Thomas van den Biggelaar, and Xander Pope.

This work relies on MAST-U experimental data (operation and diagnostics), SOLPS-ITER simulations, and collisional-radiative models. These contributions were respectively provided by the MAST Upgrade operations and diagnostics team, alongside the Eurofusion Tokamak Exploitation Team; the SOLPS-ITER community, notably David Moulton, Omkar Myatra, and Alexandre Fil; and the Open-ADAS and

Yacora teams. I'd like to also thank Dr. A. Perek, Dr. T. A. Wijkamp, and Dr. J. Karhunen, whose alternative integrated data analysis methods offered valuable points of comparison and discussion.

I would like to thank those who have given me the freedom to pursue my passion. This ongoing journey has included many significant steps: the Ogden Trust, which made my early studies possible; exposure to Bayesian inference and software engineering during my undergraduate degree at the University of Birmingham and the National University of Singapore; and my introduction to Gaussian processes through a project with Dr. J. Leddy at Tech-X Corporation during my PhD. I would especially like to thank the Fusion CDT programme for providing such an exciting project and an immensely enjoyable environment for early-career researchers.

It seems that the combination of unwavering support, humour, and positivity is almost sufficient to maintain sanity during a PhD. Thank you to my friends and family for providing this. My beautiful family, thank you for reminding me to not forget to "smell the roses" and that there is apparently more to life than fusion. My incredible partner, Vicky, thank you for sharing this crazy PhD journey with me. I also wish to thank all the amazing friends that I have made along the way. Nick, the Culham football and cheeky triathlon or two really kept me going.

Finally, I would like to acknowledge the physicists, engineers, and social scientists who have worked, and continue to work, tirelessly to advance the feasibility of controlled nuclear fusion. I remain confident that together we will make the future a brighter place for all.

Executive Summary

Nuclear fusion tokamak power plants offer a viable solution to alleviating global energy poverty. Achieving this future requires advances that overcome both physics and engineering challenges. One of the most significant challenges is the controlled exhaust of heat and particles. Numerous developments have been made in the divertor, the region responsible for handling this exhaust. However, as critical decisions are made regarding the final designs of future tokamak power plants, experimental insight into divertor behaviour remains both essential and difficult. While predictive models and various diagnostics are available, no diagnostic can directly measure the key plasma characteristics across the full divertor which complicates model validation. A validated and flexible integrated data analysis approach is therefore essential to deepen our understanding of divertor behaviour and to support the design of reliable tokamak power plants.

The work presented in this thesis develops an integrated data analysis system that combines multiple divertor diagnostics with established physics insights into a single framework. This framework, the Divertor Multi-Instrument Bayesian Analysis System (D-MIBAS), infers two-dimensional profiles of key plasma parameters in the divertor region: electron temperature (T_e), electron density (n_e), and neutral atomic hydrogen density (n_D). It is not tied to any specific machine or scenario, enabling application across present-day tokamaks and providing a transferable framework for future power plants with differing diagnostic coverage. D-MIBAS infers the key plasma parameters with inherent uncertainty quantification by addressing the question: which plasma states plausibly explain the measured diagnostic data?

The central thesis of this work is that the answer lies in a unified *basis* for inference. That basis is a spatially distributed two-dimensional mesh aligned with the divertor's magnetic equilibrium. This approach allows multiple diagnostics and sophisticated physics-informed spatial information (such as the monotonicity of the electron temperature along a magnetic 'flux surface') to be incorporated into a single inference problem, thereby reducing the range of plausible plasma states. While this

incurs additional computational cost, the work outlined in this thesis demonstrates its feasibility and applicability for two MAST Upgrade ‘low-confinement mode’ plasma discharges with differing divertor scenarios. Further validation is required in ‘high-confinement mode’ plasmas and ‘attached’ divertor scenarios, and recommendations are made for the inclusion of additional diagnostics, such as coherence imaging spectroscopy and x-point imaging data, to reduce uncertainties in the inferred plasma state.

D-MIBAS makes use of filtered spectroscopy of Balmer line emission, arising from electronic transitions in hydrogen, which has long been a key diagnostic in fusion divertor research. It is *non-invasive*, provides wide-area coverage of the divertor plasma, and is well-suited for application in future power plants. Hot, ionising plasmas (> 5 eV electron temperature) emit a different spectral hue compared to colder (< 1 eV electron temperature), recombining plasmas. This change arises from variations in plasma conditions that alter the relative intensity of Balmer transitions. Filtered imaging and spectroscopic systems can isolate these transitions, enabling high-resolution, real-time mapping of divertor behaviour across wide spatial domains. However, accurate modelling of Balmer emission requires accounting for both atomic and molecular processes. This leads to an *ill-posed* problem, in which many plasma states are consistent with the same emission pattern, preventing Balmer data alone from constraining reliable profiles of T_e , n_e , and n_D .

Emission data has been incorporated in alternative approaches to integrated data analysis for the divertor, but these adopt a *cell-based* inference. In this approach, diagnostic measurements are pre-processed and distributed into independent spatial cells. While computationally efficient, this method prevents the consistent combination of multiple diagnostics covering different regions and the inclusion of physics-informed spatial information. Each cell is inferred in isolation, typically leaving the problem ill-posed and yielding large uncertainties in the inferred plasma parameters.

This thesis demonstrates a step-change through the development and application of a *mesh-based* Bayesian inference framework. By constructing a two-dimensional mesh, aligned to the magnetic equilibrium, multiple diagnostics and physics-based spatial priors can be combined within a single inference problem. These physics-based spatial priors, such as anisotropic smoothness of plasma fields, are guided by known physics and *learned* from various SOLPS-ITER simulations. The resulting physics-informed machine learning framework maximises the information extracted from a limited set of diagnostics and constrains the range of plausible plasma states,

significantly reducing uncertainties compared to cell-based methods. By reducing uncertainty through physics-informed, mesh-based inference, this framework facilitates the interpretation of experimental data, validates predictive models, and ultimately supports the design of reliable fusion power plant divertors.

Contents

Abstract	ii
Executive Summary	x
List of Figures	xxi
List of Tables	xxiii
List of Abbreviations	xxiv
List of Symbols	xxvi
1 Introduction	1
1.1 The Tokamak	2
1.1.1 Fusion Plasma Requirements	3
1.1.2 Plasma Facing Components	4
1.1.3 SOL Plasma in the PFC Vicinity	5
1.1.4 Processes Occurring in the SOL Plasma	7
1.1.5 Plasma Species	7
1.2 Thesis Overview	8
1.2.1 Thesis Methodology	8
1.2.2 Thesis Structure	9
1.2.3 Contributions of this Work	11
2 Background: The Tokamak Divertor	13
2.1 Atomic and Molecular Physics	13
2.1.1 Photon Emission	16
2.2 The Divertor	17
2.2.1 Divertor Regimes	17
2.2.1.1 Sheath Limited Regime	17

2.2.1.2	Conduction Limited Regime	18
2.2.1.3	Detached Regime	18
2.2.2	Divertor Configuration	19
2.2.2.1	Connection Length	21
2.2.2.2	Grazing Angle	21
2.2.2.3	Total Flux Expansion	22
2.2.2.4	Baffled, Closed Configuration	22
2.2.3	Analytic Models of the Scrape-Off Layer	23
2.2.3.1	Perpendicular Profile: Eich Function	23
2.2.3.2	Parallel Profile: The Two-Point and Lengyel Models	24
2.3	Insights into the Tokamak Divertor Plasma	26
2.3.1	Desired Quantities	27
2.3.2	Plasma Edge Modelling	29
2.3.3	Diagnostics	30
2.4	Integrated Data Analysis	31
2.4.1	Characterising the Posterior Distribution	33
2.4.1.1	Low-Dimensional Problems: Grids	33
2.4.1.2	High-Dimensional Problems: MCMC	34
2.4.1.3	Handling Hyperparameters	36
2.4.1.4	Bayesian Inference	36
2.4.2	Bayesian Inference in Fusion Research	37
2.4.2.1	Combining Diagnostics	37
2.4.2.2	Utilising Spatial Information: Gaussian Processes . .	38
2.4.3	Bayesian Inference in the Divertor Region	39
2.5	Summary: Current Understanding of the Tokamak Divertor	43
3	Bayesian Inference using Balmer Emission Lines	45
3.1	Model for Molecular Contributions to Balmer Line Emission	46
3.1.1	Fulcher Band Model	48
3.2	Balmer Line Ratios	49
3.3	Bayesian Inference with Balmer Lines	51
3.3.1	Balmer Line Ability to Infer the Whole Plasma State	54
3.3.1.1	Electron Temperature	54
3.3.1.2	Electron Density	56
3.3.1.3	Atomic Neutral Density	56
3.3.1.4	Q mol.	56

3.4	Discussion	56
3.4.1	Inclusion of Additional Information	56
3.4.2	Inferred Ion Source and Sink Rates	58
3.5	Summary	58
3.5.1	Implications for Using Balmer Lines in Divertor Data Analysis	59
4	Synthetic Verification	61
4.1	The D-MIBAS Method	62
4.1.1	A Field-Aligned Inference Mesh	62
4.1.2	The Posterior Distribution	63
4.1.3	Likelihoods	63
4.1.3.1	Emission Models	65
4.1.4	Priors	65
4.1.4.1	Spatially-Independent and -Dependent Priors	66
4.1.5	Inference Parameters	66
4.1.6	Characterising the Posterior Distribution	67
4.2	Synthetic Data Setup	68
4.3	Results	70
4.3.1	Synthetic Diagnostic Test Results of the Full D-MIBAS analysis	70
4.3.2	Performance of the Different Posterior Distributions	72
4.3.2.1	Importance of Including PMI in the Balmer Emission Model	72
4.3.2.2	Importance of Spatially-Dependent Priors	73
4.3.2.3	Importance of Helium Emission	73
4.4	Discussion	74
4.4.1	Limitations of the Current D-MIBAS and Possible Improvements	74
4.4.1.1	Consideration for Experimental Data Application: Forward Model Inaccuracies	74
4.4.1.2	Improving D-MIBAS Analysis Through Additional Diagnostics	76
4.4.1.3	Derived Quantities	77
4.4.1.4	Performance of HMC	78
4.5	Summary and Conclusions	79

5	D-MIBAS for Experimental Data	81
5.1	Mega Amp Spherical Tokamak Upgrade	82
5.2	Comprehensive Diagnostic Forward Models	83
5.2.1	Full Mesh Coverage	83
5.2.2	Toroidal Asymmetries	84
5.2.3	Included Diagnostics	84
5.2.3.1	Thomson Scattering	84
5.2.3.2	Langmuir Probes	85
5.2.3.3	Multi Wavelength Imaging Camera	87
5.2.3.4	Spectroscopy	89
5.3	Prior Probability Functions	89
5.3.1	Calibrations	90
5.3.2	Gaussian Process Priors	90
5.3.2.1	Hyperparameter Choice	91
5.3.3	Monotonicity Priors	94
5.3.4	Upstream priors	95
5.3.5	Parameter Bounds	96
5.4	Covariance Matrix Approximation	96
5.5	Computational Cost	97
5.6	Summary	98
6	D-MIBAS Inference of a MAST-U Super-X Divertor	100
6.1	Discharge #46860	100
6.1.1	Verification, Validation and Uncertainty Quantification	101
6.1.1.1	Diagnostic Predictions	101
6.1.1.2	Experimental Comparison to Alternative Methods	103
6.1.1.3	Modelling Comparison	105
6.1.1.4	Inference Uncertainty	107
6.2	Results: Inferred Plasma State	107
6.3	Discussion	110
6.3.1	Progression of the Detached Divertor State	110
6.3.2	Changing Target Conditions	113
6.4	Summary and Conclusions	115

7	D-MIBAS Inference in Alternative Divertor Configurations	117
7.1	Discharge #46895	117
7.1.1	Interpretive Modelling	118
7.2	Results	119
7.3	Discussion	121
7.3.1	Diagnostic Agreement	121
7.3.2	Derived Fields	124
7.4	Summary	126
8	Conclusion and Outlook	127
8.1	On Furthering Exhaust Understanding	129
8.2	On Advising Future Integrated Data Analysis Efforts	130
8.3	On D-MIBAS's Role in Fusion Power Plants	131
	Bibliography	132
	Appendix	141
A	Mass Matrix Construction for Hamiltonian Monte Carlo	142
A.1	Adaptive Estimation	142
A.2	MAP Curvature Approximation	143
A.3	<i>A Priori</i> Construction	144
B	Settling Aspects of the Bayesian Inference Approach of this Work	146
B.1	Non Informative Prior Choice	146
B.2	Likelihood Degrees of Freedom	148
B.3	Likelihood Scale	148
B.4	Comparing Inferred Parameter Estimation Methods	149
C	Helium Emission Models	150
D	Priors for Chapter 4	151
D.1	Prior Distributions	151
D.2	Prior Details	152

List of Figures

1.1	Schematic of a tokamak.	
	a) Magnetic fields in a tokamak. Adapted from Figure 7 of [81].	
	b) Cross section displaying magnetic flux surfaces in the tokamak core.	3
2.1	Reaction rates of collated hydrogenic atomic and molecular processes using ADAS[90] and Yacora[134] at $T_e n_D = 10^{20} eV m^{-3}$ and $T_e n_{D_2} = 10^{19} eV m^{-3}$. Shading transparency corresponds to electron density. Additional shading of MAD and MAI at high electron densities indicative of uncertainties surrounding the molecular rates.	
	Upward/downward arrows indicate the impact on reaction rates due to: a ten times increase/decrease in n_D for atomic ionisation (red); and a ten times increase in n_{D_2} for MAD (purple), MAR (green) and MAI (cyan). MAD rates include MAR and MAI rates.	16
2.2	Poloidal cross-section of the MAST-U tokamak divertor.	
	a) Divertor configurations in the MAST-U tokamak divertor. Adapted from Figure 1 of [122].	
	b) Schematic of a flux-tube in a Super-X divertor projected onto the 2D poloidal cross section.	20
2.3	The Lengyel model. Qualitative example solution to equations (2.1) to (2.6) of [62].	
	Adapted from Figure 11 of [62].	26
2.4	Flux tube element obeying particle balance (illustration of Gauss's theorem). S_+ (S_-) represents the rate of particle sources (sinks) per unit volume.	28

2.5	Representations of a 2D posterior distribution with colour reflecting the posterior probability.	
	(a) The underlying posterior distribution to be deduced. Marginalised probability distributions shown with shading reflecting 95% highest density interval (HDI).	
	(b) The posterior distribution found through a grid based method.	
	(c) The posterior distribution found through an adaptive-grid based method.	
	(d) The posterior distribution found with a Hamiltonian Markov Chain Monte Carlo (starting in lower left corner with ‘burn in’ included). . .	34
2.6	Experimental measurements of two-dimensional electron temperature and electron density profiles as available in the literature. Please see original papers [12, 47, 63, 85], figures displayed here for visual comparison only.	42
2.7	Qualitative summary of a detached divertor. Inspired by Figure 4 of [114]. Dashed line indicates the separatrix.	44
3.1	Photon emissivity of the Fulcher band. Emissivities at $T_e n_{D_2} = 10^{19} \text{eV m}^{-3}$. Upward/downward arrows show impact on the peak due to a ten times increase/decrease in $T_e n_{D_2}$. . .	49
3.2	Photon emissivities of Balmer lines $3 \rightarrow 2$ (D_α) to $6 \rightarrow 2$ (D_δ) at $T_e n_D = 10^{20} \text{eV m}^{-3}$ and $T_e n_{D_2} = 10^{19} \text{eV m}^{-3}$. Upward/downward arrows show impact on the peaks due to a ten times increase/decrease in $T_e n_D$ (dashed arrow) and $T_e n_{D_2}$ (solid arrow).	50
3.3	Balmer emission ratio of $3 \rightarrow 2$ (D_α) to $6 \rightarrow 2$ (D_δ) with different molecular densities ($T_e n_D = 10^{19} \text{eV m}^{-3}$).	51
3.4	Inferred marginalised electron temperature probability distributions through Bayesian inference with different Balmer line combinations. Upper row (a-c): D_α, D_δ Lower row (d-f): $D_\alpha, D_\beta, D_\gamma, D_\delta$ Samples refer to multiple trials of the inference, each with different synthetic data owing to the expected 10% uncertainty on measurement data.	53

3.5	Qualitative pictogram of typical inference capability (error, accuracy and uncertainty) for key plasma conditions based on differing information included in the inference. Large, green circles with solid rings indicate effective inference. Error, uncertainty and % inside correspond to inference error, inference uncertainty and inference accuracy respectively as defined in Section 3.3. Plasma conditions A, B, C, D and E roughly correspond to regions in Figure 2.7 and are defined in the prose. Specific quantities represented in the pictogram available on request.	55
4.1	The inference mesh and the diagnostic coverage in the MAST-U divertor. The magnetic geometry is obtained from SOLPS-ITER simulations of the MAST-U Super-X divertor.	63
4.2	Activity diagram for characterising the posterior distribution in Bayesian inference.	67
4.3	Comparison of ‘true’ fields used to generate synthetic data (a) and the inferred MAP estimate (b). (a) Plasma fields for two synthetic cases from SOLPS-ITER. (I) The detached case where the T_e (a.I.i.) drops below 5eV before it reaches the target. (II) The attached case in which the n_e (a.II.ii.) peaks at the target. (b) Maximum a posteriori (MAP) estimate inference of scenarios displayed in Figure 4.3(a) using posterior i. MAP was found using the generational algorithm outlined in Appendix A of [9].	69
4.4	Comparison of true and inferred (MAP) parameter values using posterior ifor the detached case. The 95% highest density interval (HDI) for the inferred parameters are displayed. The percentage of true data points lying within the 95% HDI are 95%, 88%, 63%, 91%, 95%, 63% for the T_e , n_e , n_D , $Q_{\text{mol.}}$, n_{He^0} and $n_{\text{He}^{1+}}$ respectively.	71
4.5	The synthetic camera data for Balmer lines D_α (3→2, red) and D_δ (6→2, lilac) and their forward model predictions based on the MAP estimate of the four posterior distributions outlined in Section 4.3.	72
4.6	Performance of posteriors i-iv across fields of interest for the detached case.	72

4.7	Comparison of inferred parameters with the ‘true’ values (Figure 4.3Ia) for the detached case. The parameters follow a surface of constant poloidal flux, $\psi = 1$ (the separatrix). Solid lines indicate the inferred (MAP) parameters and the shaded region indicate the 95% HDI. . . .	75
4.8	Comparison of true and inferred (using posterior i) static electron pressure and parallel heat flux values for the detached case. Evaluated for parameters perpendicular to surfaces of constant poloidal magnetic flux at a magnetic connection length of around 3 metres to the target . .	78
4.9	Comparison of true (a) and inferred (b) (using posterior i) dominant emission regions for the D_γ Balmer line for the detached case.	78
4.10	Visualisation of error (a) and uncertainty (b) for the detached case using posterior (i).	79
5.1	Inference mesh (a) and diagnostic coverage (b) for experimental use. Diagnostic coverage describes: spectroscopy - solid lines; multi-wavelength imaging - shaded region; Thomson scattering - hollow circles; and Langmuir probes - filled triangles.	83
5.2	MWI D_β channel image processing. (a.i.) Single camera viewing chord comprised of 7 evenly spaced individual chords (a.ii.) Camera viewing chord, for pixel (293, 97), through MAST-U lower divertor chamber. (b) Camera chord passing through an inference mesh. Geometry matrix contribution for each mesh vertex for chord (293, 97). (c) D_β MWI brightness image. Chord (293, 97) circled. (d) Relative uncertainty (uncertainty / brightness) for D_β MWI brightness image.	88
5.3	Sample functions from different Gaussian Process priors using the Matérn and Rational Quadratic kernels given in equations (5.17) and 5.19 with arbitrary scale, varying lengthscale (ℓ), and varying shaping parameters (α and ν). Dashed curves show <i>latent</i> GP samples $f(x)$ and solid curves show the corresponding <i>physical</i> fields $\exp(f(x))$. Input defined over a domain from 0 to L with lengthscales, ℓ , given relative to L. . .	93

5.4	Covariance matrix approximation based on the covariance kernels of equation (5.16). Parameter number represents position in parameter vector θ (see equation (4.6)), with numbers assigned sequentially across the rows and columns of the inference mesh. Labels A to D described in the prose.	97
6.1	Comparison between diagnostic data and forward model at the MAP inference.	102
6.2	Comparison of alternative analysis procedures (bottom row) to results recreated from the D-MIBAS inferred plasma state (top row). Bottom-left: A SART inversion from techniques given in [133] Bottom-centre: Coherence-imaging spectroscopy, recreated from Figure 16 of [72]. Bottom-right: Spectroscopic analysis through the BaSPMI method [112], recreated from Figure 12 of [89]. Note: shaded region corresponds to 95% HDI in D-MIBAS results (top-right), but 67% equal-tailed HDI in BaSPMI Analysis (bottom-right).	104
6.3	D-MIBAS inferred plasma state at 600 ms (37% Greenwald fraction) compared to SOLPS-ITER interpretive modelling of discharge #46860 .	106
6.4	Percentage uncertainty of inferred parameters at 600 ms (37% Greenwald fraction). 95% HDI used to quantify the uncertainty.	108
6.5	Inferred plasma states throughout discharge #46860. a) 27% Greenwald fraction. b) 37% Greenwald fraction. c) 45% Greenwald fraction. Magenta dash represents the 5 eV electron temperature front. Shaded region represents lack of diagnostic coverage.	108
6.6	Derived fields from the inferred plasma states throughout discharge #46860. a) 27% Greenwald fraction, 450 ms. b) 37% Greenwald fraction, 600 ms. c) 45% Greenwald fraction, 750 ms. Magenta line represents the 5 eV electron temperature front. Shaded region represents lack of diagnostic coverage.	110

6.7	Inferred and derived fields following the separatrix ($\psi=1$). Shaded region indicates 95% highest density interval. Particle flux density calculated using equation 2.19 and the assumption of negligible cross-field transport.	111
6.8	D-MIBAS inferred plasma state during the core density ramp discharge #46860.	114
7.1	Inference mesh (a) and diagnostic coverage (b) for experimental use. Diagnostic coverage describes: spectroscopy - solid lines; multi-wavelength imaging - shaded region; Thomson scattering - hollow circles; and Langmuir probes - filled triangles.	118
7.2	D-MIBAS inferred fields for discharge #46895. Top row (a): Conventional divertor (CD), 400 ms. Middle row (b): Elongated divertor (ED), 600 ms. Bottom row (c): Super-X divertor (SXD), 800 ms.	121
7.3	Interpretive SOLPS-ITER modelling of CD, ED, and SXD MAST-U divertor configurations presented as part of [122] and available at [113]. Top row (a): Conventional divertor (CD). Middle row (b): Elongated divertor (ED). Bottom row (c): Super-X divertor (SXD).	121
7.4	Comparison between diagnostic data and forward model at the MAP inference.	123
7.5	Derived fields from the inferred plasma states throughout discharge #46895. Top row (a): Conventional divertor (CD). Middle row (b): Elongated divertor (ED). Bottom row (c): Super-X divertor (SXD).	125
7.6	Interpretive SOLPS-ITER modelling of CD, ED, and SXD MAST-U divertor configurations presented as part of [122] and available at [113]. Top row (a): Conventional divertor (CD). Middle row (b): Elongated divertor (ED). Bottom row (c): Super-X divertor (SXD).	125

B.1	Inferred marginalised electron temperature probability distributions through Bayesian inference with Balmer lines D_α , D_β , D_γ , and D_δ . Upper row (a-c): Jeffreys Prior included Lower row (d-f): Fulcher data included Samples refer to multiple trials of the inference, each with different synthetic data owing to the expected 10% and 20% uncertainty on Balmer line and Fulcher measurement data respectively.	147
-----	---	-----

List of Tables

2.1	Dominant process occurring in the scrape-off layer. Species X and Y include intrinsic impurities, extrinsic impurities, helium and hydrogen. Excited species (*) subsequently relax through photon emission (in recombination, it is possible for species listed as excited to immediately enter the ground state in which case a photon is immediately released).	14
2.2	Hydrogenic processes occurring in the scrape-off layer and their impact on particle number. Processes involving D^- and D_3^+ are omitted.	15
5.1	Empirical Bayes selection of hyperparameters for various plasma fields in the MAST-U divertor. Where two values are provided for the length-scale, ℓ , an anisotropic kernel is used with the first element corresponding to the direction orthogonal to surfaces of constant poloidal	94
5.2	Monotonicity priors. $\partial_{s_{\parallel}}$ denotes the derivative, found via finite difference across neighbouring mesh vertices, parallel to the surface of constant poloidal magnetic flux. ‘symlog’ denotes a symmetric log transform, i.e. $f_{symlog}(x) = \frac{x \log_b(b+ x)}{b+ x }$ where b is the logarithm base (in this case, Euler’s number, e).	95
5.3	Upstream priors. $H(\theta) = 1$ for mesh vertices within 30 cm of the magnetic X-point, and 0 otherwise. $\text{sgn}(\theta) = 1$ for $\psi(\theta) < 1$ (mesh vertices in the private-flux region) and -1 otherwise. $\partial_{s_{\perp}}$ and $\partial_{s_{\parallel}}$ denote the spatial derivative perpendicular and parallel to surfaces of constant poloidal magnetic flux respectively. These derivatives were found via finite difference across neighbouring mesh vertices. P_{max} is the external heating power supplied to the plasma.	96
5.4	Parameter bounds used in the inference.	96
D.1	The spatially independent priors implemented within D-MIBAS for Chapter 4. When the argument deploys an upper and lower limit, the scale parameters are respectively listed.	152

D.2	The spatially dependent priors implemented within D-MIBAS for Chapter 4. Where multiple fields F are stated, the scale parameters are, respectively, listed. Derivatives $\partial_{s_{\parallel}}$, $\partial_{s_{\theta,\parallel}}$ and $\partial_{s_{\psi_N}}$ are with respect to distance parallel to total flux surfaces, distance parallel to surfaces of constant poloidal magnetic flux, and distance perpendicular to surfaces of constant poloidal magnetic flux, respectively. Gradients are approximated by means of finite differences with neighbouring mesh vertices.	154
-----	---	-----

List of Abbreviations

ADAS	Atomic Data and Analysis Structure
ASDEX-U	Axially Symmetric Divertor EXperiment Upgrade
BaSPMI	Balmer Spectroscopy technique for Plasma–Molecule Interaction
CX	Charge Exchange
CFR	Common Flux Region
CIS	Coherence Imaging Spectroscopy
D-MIBAS	Divertor Multi Instrument Bayesian Analysis System
DMS	Divertor Monitoring Spectroscopy
DTS	Divertor Thomson Scattering
ED	Elongated Divertor
EIE	Electron Impact Excitation
EIR	Electron-Ion Recombination
ELM	Edge Localised Mode
GP	Gaussian Process
HDI	Highest Density Interval
HMC	Hamiltonian Monte Carlo
IDA	Integrated Data Analysis
JET	Joint European Torus
LCFS	Last Closed Flux Surface
MAD	Molecular Activated Dissociation
MAI	Molecular Activated Ionisation
MANTIS	Multispectral Advanced Narrowband Tokamak Imaging System
MAP	Maximum a posteriori
MAR	Molecular Activated Recombination
MAST-U	Mega Amp Spherical Tokamak Upgrade
MCMC	Markov Chain Monte Carlo
MWI	Multi-Wavelength Imaging
OD	Open Divertor
PEC	Photon Emissivity Coefficient
PFC	Plasma-Facing Component
PFR	Private Flux Region
PMI	Plasma-Molecule Interaction
RQ	Rational Quadratic kernel

SART	Simultaneous Algebraic Reconstruction Technique
SOL	Scrape-Off Layer
SXD	Super-X Divertor
TCV	Tokamak à Configuration Variable
UQ	Uncertainty Quantification

List of Symbols

Constants:

e	Elementary charge, 1.602×10^{-19}	C
ϵ	Ionisation potential: atomic hydrogen, 13.6; molecular hydrogen, ≈ 15.8 .	eV
$\kappa_{0,e}$	Parallel electron thermal conductivity coefficient (Spitzer–Härm), 2000	$\text{W m}^{-1} \text{eV}^{-7/2}$

Common subscripts:

0	Separatrix
A	Area
atm	Atomic
dep	Deposited
e	Electron
eff	Effective
exc	Excitation
i	Ion
int	Integral
lwr	Lower
m	Outer midplane
mol	Molecular
q	Heat flux
rec	Recombination
se	Sheath edge
sh	Sheath
t	Target
u	Upstream (between X-point and mid-plane)
upr	Upper
V	Volume

x	At the magnetic X-point
\parallel	Parallel to magnetic field
\perp	Perpendicular to magnetic field
+	Positive charge
-	Negative charge

Inferred plasma parameters:

b_i^j	Background emissivity for line i and diagnostic j	photons $\text{m}^{-3} \text{s}^{-1}$
c_i^j	Absolute brightness calibration factor for line i and diagnostic j	
n_{D}	Neutral atomic hydrogen density	m^{-3}
n_e	Electron density	m^{-3}
n_{He^0}	Neutral atomic helium density	m^{-3}
$n_{\text{He}^{1+}}$	Singly charged atomic helium density	m^{-3}
T_e	Electron temperature	eV
$Q_{\text{mol.}}$	Ratio of Balmer α emission from molecular to atomic contributions (eq. 3.3)	
β	Langmuir probe nuisance parameter (eq. 5.7)	

Plasma state quantities

B	Magnetic field strength	T
\mathbf{B}	Magnetic field vector	T
c_s	Ion sound speed	m s^{-1}
I_p	Plasma current	A
D_{\perp}	Cross-field diffusion coefficient	m^2s^{-1}
P	Power	W
p	Pressure	eV m^{-3}
n	Number density	m^{-3}
q	Heat flux density	W m^{-2}
T	Temperature	eV
v	Particle velocity	m s^{-1}
Z_{eff}	Ion effective charge state	
λ_{int}	Integral heat-flux width	m
λ_q	Heat-flux decay length	m
λ_{SOL}	Scrape-off layer width	m
Γ	Particle flux	$\text{m}^{-2}\text{s}^{-1}$

Geometric and magnetic factors

a	Plasma minor radius	m
$F_{R,m}$	Total flux expansion (outer-midplane reference)	
$F_{R,x}$	Total flux expansion (X-point reference)	
$F_{X,m}$	Poloidal flux expansion (outer-midplane reference)	
$F_{X,x}$	Poloidal flux expansion (X-point reference)	
L_{\parallel}	Connection length (midplane to target along \mathbf{B})	m
n_G	Greenwald density limit	m^{-3}
R	Major radius (distance from torus axis)	m
R_0	Plasma major radius (at magnetic axis)	m
S	Cross-field spreading parameter	m
s_{\perp}	Perpendicular (to field line) distance coordinate	m
s_{\parallel}	Parallel (to magnetic field line) distance coordinate	m
$s_{\text{pol.}}$	Poloidal distance coordinate	m
\bar{s}	Displacement from separatrix	m
Z	Height	m
α	Grazing angle of field line to surface	rad
ϕ	Toroidal angle	rad
ψ	Normalised poloidal flux (0 at axis, 1 at separatrix)	

Two-point model parameters

$f_{\text{cond.}}$	Conductive heat-flux fraction
$f_{\text{mom.}}$	Momentum retention factor
$f_{\text{pow.}}$	Fraction of power loss due to volumetric processes
p	Total pressure
γ	Adiabatic index
γ_{sh}	Sheath heat transmission factor

Emission / atomic physics parameters

A_{pq}	Einstein coefficient for transition $p \rightarrow q$	photons s^{-1}
$PEC_{p \rightarrow q}$	Photon emissivity coefficient	photons $\text{m}^3 \text{s}^{-1}$
R_p	Population coefficient for level p	m^3
$\varepsilon_{n \rightarrow 2}$	Emissivity of $n \rightarrow 2$ transition	photons $\text{m}^{-3} \text{s}^{-1}$

Inference notation

\mathcal{D}	Measured data
\mathbf{f}	Function values sampled from GP
F	Field
ℓ	GP length-scale parameter
\mathcal{L}	Natural log of the probability
$\mathcal{M}(\theta)$	Forward model mapping parameters to data
P	Probability
V	The number of mesh vertices.
\mathbf{X}	Inputs
α	Hyperparameters / shape parameter for Rational Quadratic kernel
κ	Shape parameter
μ	Mean
ν	Degrees of freedom (Student's t , Matérn kernel)
Σ	Covariance
θ	Inference parameters (plasma state)

Chapter 1

Introduction

In 2020, at least one in seven humans were found to live in energy poverty¹ [74]. There will be yet further pain to the world's ecosystem without the immediate cessation of electricity generation through means which directly contribute to climate change. The need to transition the world's means of electricity generation must come alongside cheaper, more available and more reliable electricity worldwide. The new means of electricity generation must be without geographical advantage to avoid the risk of exploitation. The achievement of nuclear fusion on Earth, as a means of mass electricity generation, offers a potential pathway forward for humanity.

The coming fleet of nuclear fusion power plants makes a compelling argument to share a substantial fraction of humankind's future energy demands. The proposed fuels are derived from abundant resources, extracted from seawater and bred from lithium, and fusion does not produce long-lived radioactive waste. However, the successful realisation of nuclear fusion power plants still requires significant challenges to be overcome. The common choice for current power plant design fuels is the fusion of deuterium and tritium into a neutron and helium nucleus which share 17.6 MeV of released energy. To become feasible, this reaction requires temperatures in excess of 100 million degrees Celsius and substantial pressures [1]. The deuterium and tritium therefore must exist in a highly ionised state of matter: plasma.

Controlled nuclear fusion requires the management of high-pressure plasma. Multiple strategies are being investigated to confine this plasma. One such strategy, the tokamak, uses magnetic fields to trap the plasma in a torus-shaped vessel. However, the plasma will damage any material with which it comes into contact and multiple processes lead the plasma to transport from the plasma centre to the vessel edge and subsequent surrounding material. The magnetic fields can be configured to handle

¹Energy poverty is defined to be unable to use electricity either through lack of access or cost.

the plasma's interaction with the surrounding material wall (plasma-facing components, PFCs) in a specialised region known as the divertor. Being able to diagnose the interaction between the plasma and surrounding vessel is paramount for realising controlled nuclear fusion power plants on Earth. However, diagnosing such a harsh environment is extremely complex.

This work focuses on the development of advanced methodology for diagnosing this critical region. Fully appreciating the various complex processes that occur in the divertor demands minimally a two-dimensional view of the divertor plasma. In this work, we outline and test a methodology to infer two-dimensional cross sections of key plasma characteristics to sufficient spatial resolution, accuracy, and precision in order to give an experimental insight into this critical region of a tokamak. Having the capability to do so with a non-device, non-scenario specific method is of great benefit to the design of this critical region of a future fusion power plant.

1.1 The Tokamak

The torus-shaped tokamak uses toroidal and poloidal magnetic fields to confine the plasma. As shown in Figure 1.1(a), the resultant magnetic field, with magnetic flux density \mathbf{B} , spirals around the torus with the charged plasma species gyrating around the resultant magnetic field, with a resultant plasma current, I_p . Surfaces can exist for which sets of magnetic field lines never intersect the surface. The scalar ψ can be used to describe these surfaces such that

$$\mathbf{B} \cdot \nabla \psi = 0. \tag{1.1}$$

These surfaces of constant magnetic flux (magnetic flux surfaces) can themselves be toroidal in shape such that they wrap around a closed volume, Figure 1.1(b). However, at the end of this 'core' region of a tokamak exists the last closed flux surface (LCFS) beyond which subsequent flux surfaces are open (ultimately being ended by surrounding wall material). This region of open flux surfaces defines the 'scrape-off-layer' (SOL) region of a tokamak. ψ denotes the normalised poloidal flux, defined as the magnetic flux enclosed between the magnetic axis and a given flux surface. It is normalised such that $\psi = 0$ at the magnetic axis, $\psi = 1$ at the LCFS, and $\psi > 1$ is the SOL region. Numerous processes, such as collisions, magnetic drifts, and turbulent fluctuations, can cause a charged particle to leave the magnetic field line it initially gyrates around. The resulting cross-field transport allows particles to move radially outward from the confined core plasma into the scrape-off layer (SOL),




Figure 1.1: Schematic of a tokamak.

a) Magnetic fields in a tokamak. Adapted from Figure 7 of [81].

b) Cross section displaying magnetic flux surfaces in the tokamak core.

where they can eventually interact with plasma-facing components and surrounding wall materials [131].

1.1.1 Fusion Plasma Requirements

Power plants require the fusion power generated to be sufficient to sustain the heat of the plasma. This requires the fusion power to exceed power losses from the plasma through particles leaving the plasma and radiation from the plasma [60]. A useful metric to quantify the conditions required for sufficient fusion power is the ‘fusion triple product’ of the electron density, n_e , electron temperature, T_e , and confinement time, τ_E . Future power plants require a high product of these reactants to satisfy the ‘Lawson criterion’ [60].

Challenges to a high triple product in tokamaks revolve around achieving substantial confinement times. The magnetohydrodynamics (MHD) in the plasma are prone to instabilities, resulting in a loss of confinement. Tokamaks require high core temperatures and cold wall temperatures. However, MHD instabilities are triggered by both the cold species from the wall entering the core plasma and the sharp temperature gradients required to keep plasma near the wall cold. Increasing the line-averaged core density above the Greenwald limit,

$$n_G = \frac{I_p}{\pi a^2} \times 10^{14}, \quad (1.2)$$

where a is the minor radius of the plasma, is found to typically lead to MHD instabilities [37]. This provides an empirical operational boundary for tokamaks, with the

line-averaged core density often expressed as a fraction of the Greenwald limit.

The SOL plasma is populated by high temperature and density plasma that is transported from the core and across the LCFS. Consequently, the static electron pressure of the SOL plasma ($n_e T_e$) is closely related to the plasma pressure of the core. The ASDEX tokamak discovered a ‘high-confinement’ (H-mode; as opposed to the low-confinement, L-mode) operating regime where, beyond a threshold power, an energy transport barrier is formed [125]. This barrier leads to the formation of a pedestal in the vicinity of the LCFS, producing a sharp increase in core plasma temperature and density relative to the SOL. However, it also leads to a narrow scrape-off layer width, meaning that unmitigated heat fluxes leaving the core plasma can substantially exceed the engineering limits of the surrounding wall material.

In the SOL plasma, a highly favourable ‘detached’ scenario has also been discovered. Here, the hot plasma (exiting the plasma core upstream) interacts with neutral and impurity species away from the divertor target walls. This can lead to multiple orders of magnitude reduction in peak heat and particle fluxes to target walls, as well as a reduction in plasma temperature in the wall’s vicinity, which is required for future power plant devices to meet engineering limits [88, 128]. Detachment is achieved through power, particle, and momentum losses (within the plasma, before they reach material surfaces), due to interactions between the plasma and neutral [65, 121], impurity [7, 42, 92], and molecular [112, 114, 116, 117, 118, 120] species.

1.1.2 Plasma Facing Components

The desire to dissipate power, particles, and momentum within the plasma before they reach material surfaces arises from the need for plasma-facing components (PFCs) to retain sufficient mechanical strength in the extreme thermal environment. Divertor target PFCs must be able to conduct away the incident heat flux to avoid melting. However, localised high heat flux densities are common which can lead to localised melting. The PFC is also susceptible to the sputtering of impurities into the plasma. Ions can transfer their kinetic energy to PFC atoms and dislodge them through physical sputtering. Some PFC’s (e.g. Carbon) are also susceptible to high volumes of chemical sputtering in which PFC atoms react with ions to form compounds which are subsequently released into the plasma. The extent of sputtering depends on both the size of incident ion flux and the energy carried by these ions (their temperature). Thus, there is an engineering requirement to limit the peak heat flux density, particle flux to the target, and the temperature of the ions near the target.

1.1.3 SOL Plasma in the PFC Vicinity

Plasma, transported from the core to the scrape-off layer, predominantly flows along flux surfaces to a ‘target’ PFC wall. To meet material constraints, the plasma at the target must simultaneously have a manageable heat flux density q_t , particle flux Γ_t , and temperature T_t . However, fundamental plasma behaviour constrains these quantities.

In a plasma, each charged particle is surrounded by a cloud of opposite charges that screens its electric field. The characteristic scale of this Debye shielding is the Debye length, which in fusion-relevant plasmas is typically very small. Therefore, macroscopically, plasmas are quasi-neutral, meaning the electron density approximately balances the ion charge density,

$$n_e \approx \sum_i Z_i n_i. \quad (1.3)$$

For a hydrogenic plasma, the effective plasma charge, Z , is one, and therefore $n_e \approx n_i$. Deviations from neutrality are confined to scales shorter than the Debye length.

Close to the wall, the lighter, more mobile electrons initially reach the wall more rapidly, generating a potential difference between the plasma and the wall. The resulting electric field repels subsequent electrons and attracts positive ions, generating a ‘sheath’ which balances the flux of ions and electrons to the wall. The substantial electric field in the sheath is screened from the plasma within a few Debye lengths.

To ensure that the flux of ions and electrons to the wall is matched (a *stable* sheath), ions must enter the sheath edge (se) with sufficient average velocity, $v_{i,se}$. This condition is known as the Bohm criterion and, as generalised by Riemann in [95], is expressed as

$$v_{i,se} \geq c_s = \sqrt{\frac{(T_e + \gamma T_i) e}{m_i}}, \quad (1.4)$$

where c_s is the ion sound speed, e is the elementary charge, m_i is the ion mass, and T_i and T_e are the ion and electron temperatures respectively. γ , the adiabatic index, is dependent on the plasma flow [100]: one for 1D isothermal flow and three for 1D adiabatic flow.

Since the sheath is thin, there is limited opportunity for significant perpendicular transport within it. Therefore, the particle flux entering the sheath edge approximates the particle flux to the target wall,

$$\Gamma_t = \Gamma_{se} = n_{se} v_{i,se}, \quad (1.5)$$

where n_{se} is the electron (\approx ion) density at the sheath edge. Momentum conservation dictates that the total pressure along a flux surface, p_{\parallel} , remains constant in the absence of sinks. Assuming quasi-neutrality and a negligible contribution from dynamic electron pressure,

$$p_{\parallel} = \underbrace{n_e (T_e + T_i)}_{\text{static}} + \underbrace{m_i n_e v_{i,\parallel}^2}_{\text{dynamic}}, \quad (1.6)$$

where $v_{i,\parallel}$ is the ion velocity parallel to a flux surface.

Therefore, the acceleration of ions to the sound speed demands that the density at the sheath edge is not the same as the density upstream of the sheath edge. Assuming isothermal conditions, the density is typically considered to halve during acceleration to the sound speed [100]:

$$\Gamma_t \gtrsim \frac{1}{2} n_e c_s. \quad (1.7)$$

This flux of ions transfer a certain amount of power to the ‘target’ wall. Considering flux at a normal to a surface, this total heat flux delivered to the target from the plasma² is,

$$q_t = \Gamma_t \left(\underbrace{\gamma_{\text{sh}} T_t}_{\text{kinetic}} + \underbrace{\epsilon}_{\text{potential}} \right), \quad (1.8)$$

where T_t is the temperature at the target and ϵ is the potential energy carried by each ion (13.6 eV for a hydrogenic ion and an additional 2.5 eV for molecular dissociation). The sheath heat transmission coefficient, γ_{sh} , represents the dynamics of the sheath that can increase the kinetic energy of the ions. It is commonly taken that $\gamma_{\text{sh}} \approx 7$ [100].

Lowering the target temperature reduces the heat flux density. At temperatures below approximately 2 eV, the heat flux received by the target is dominated by the potential energy carried by ions, so the target particle flux, $\propto n_e \sqrt{T_e}$, must also be reduced. However, due to momentum conservation (see equation (1.6)), lowering the temperature increases the density and hence the particle flux. This intrinsic coupling between target temperature, particle flux, and heat flux density must be broken for the divertor to sustain both high upstream plasma pressure and target conditions compatible with plasma-facing components. Plasma detachment offers a promising route to do this [100].

²Plasma-facing components also receive heat flux from radiated photons and energetic neutral particles.

1.1.4 Processes Occurring in the SOL Plasma

Particles exiting the core plasma can undergo numerous processes in the SOL before reaching the wall. Ions striking the wall neutralise (surface recombination) and may re-enter the plasma as neutrals. These neutrals can then be excited or ionised by hot plasma species, or undergo charge exchange with ions. In addition, ions can capture free electrons within the plasma itself (volumetric recombination). Aside from surface recombination, these atomic and molecular processes (discussed in Section 2.1) directly influence the divertor balance:

- Power loss through line radiation associated with excitation, ionisation, and recombination, which reduces q_t ;
- Particle loss through volumetric recombination, which decreases the ion population and hence the Γ_t ;
- Momentum loss through charge exchange and elastic scattering with neutrals, which acts as a pressure sink from the plasma permitting low T_t without leading to high densities and corresponding high particle fluxes.

Together, these processes underpin the onset of detachment. As further discussed in Section 2.2.1.3, the removal of power, particles, and momentum from the divertor plasma through radiation, recombination, charge exchange, and elastic scattering causes the plasma temperature near the target to decrease. This shifts the ionisation source upstream and leads to a roll-over and eventual reduction in target ion flux and heat flux. In this detached state, the plasma-facing components receive substantially lower heat and particle fluxes.

1.1.5 Plasma Species

The plasma in many planned fusion power plants is composed of deuterium and tritium fuel. These produce neutrons (neutral and thus not constrained by the magnetic fields) and helium nuclei. To obtain detachment, ‘seeded’ impurity species such as nitrogen or argon may be added to the plasma SOL [79, 92]. ‘Intrinsic’ impurities from the surrounding wall material will also enter the plasma. At the low temperatures associated with detachment, all of these species can exist as neutral atoms or even as molecules.

For ease of study, the experiments considered in this work contain only deuterium *fuel* (which, in the achieved experimental temperatures, do not fuse in large amounts).

Helium can be added to the plasma to replicate its presence in power plant-like conditions and corresponding diagnostic potential. However, the experiments used in this work produced an insufficient helium emission camera signal. The experiments in this work use carbon walls; however, these are not planned to be used in future power plants. Consequently, diagnostics that make use of carbon *impurities* in the plasma are not considered for diagnostic purposes in this work.

1.2 Thesis Overview

This work focuses on developing an integrated data analysis system to enhance our understanding of the tokamak divertor plasma. Improved characterisation of the divertor has two key benefits. Firstly, it supports and refines our understanding of important divertor processes, such as impurity radiation, ionisation, recombination, charge exchange, momentum loss, and the onset of detachment, thereby facilitating the validation of both high- and low-fidelity models. This validation is crucial for the design of future tokamak divertors for power plants. Secondly, the system enables the identification of essential quantities to track in the divertor, such as strike point positions and detachment front location, and can be used to determine the minimum set of diagnostics required to infer these quantities. This supports the development of real-time control schemes for future power plants by identifying suitable markers for key divertor states and the reduced diagnostic set needed to monitor them.

1.2.1 Thesis Methodology

The diagnosis of the complex divertor region is challenging. It requires at least a two-dimensional description of multiple plasma parameters. However, it has a sparse set of diagnostics which either provide direct but highly localised measurements or wide-view but indirect measurements of plasma parameters. Therefore, sophisticated integrated data analysis (IDA) methods are required which our work addresses through the the Divertor Multi Instrument Bayesian Analysis System D-MIBAS [9, 36]. D-MIBAS utilises Bayesian inference for IDA, enabling the incorporation of multiple diagnostics and expected physics while providing inherent uncertainty quantification. We identify both aspects as crucial for developing accurate, low-uncertainty inferences of key plasma characteristics in the divertor region.

The diagnostics considered in this work include Langmuir probes, Thomson scattering systems, and filtered spectroscopy systems. The filtered spectroscopy systems investigated the hydrogen Balmer series, helium transitions, and the Fulcher band.

For the Balmer series, an advanced emission model is used which accommodates contributions from hydrogen molecules that can form in the tokamak divertor region [36, 112].

D-MIBAS parameterises the plasma state through a two-dimensional inference *mesh*, aligned to the divertor magnetic equilibrium. This provides a unified spatial basis from which models to multiple diagnostics, each with different views, can be constructed. It also enables the natural inclusion of spatially dependent physical expectations into the Bayesian inference. For example, the electron temperature is expected to decrease monotonically along flux surfaces and the electron density is expected to vary smoothly across space. In this work, such *priors* are not imposed but instead *learned* from simulated plasma states. This allows D-MIBAS to quantify, probabilistically, the extent to which these spatial dependencies hold. Plasma fields are found to exhibit correlations well described by Gaussian processes, with characteristic lengthscales and amplitudes inferred for subsequent use. The result is a physics-informed machine learning framework capable of inferring divertor plasma states in a machine- and scenario-independent manner.

1.2.2 Thesis Structure

The thesis begins by exploring the integrated data analysis problem that D-MIBAS must solve to extract divertor plasma parameters essential for understanding heat and particle exhaust. Chapter 2 reviews the fundamental physics of the divertor covering analytic models that underpin the physics-based information (*priors*) included in D-MIBAS. Chapter 2 outlines the atomic and molecular processes governing divertor behaviour and Balmer line emission, and surveys the available diagnostics, interpretive models, and integrated data analysis strategies. The chapter establishes a benchmark for the current understanding of the tokamak divertor and concludes that two-dimensional descriptions of T_e , n_e , n_D , and an explicit treatment of molecular processes alongside some highlighted assumptions, are adequate for providing comprehensive experimental insights to validate this understanding.

Chapter 3 adapts and uses a four-parameter model of Balmer line emission, presented in [112], that captures both atomic and molecular processes. The model is used to explore Bayesian inference for integrated data analysis with Balmer line emission data. It demonstrates that the data can differentiate cold recombining plasmas from other plasma states. However, for all other plasma states, many different plasma parameter combinations could explain the measured data resulting in large inferred uncertainties. We conclude that additional information regarding whether the plasma's

electron temperature is above or below 3 eV is essential to reduce inferred plasma state uncertainties to acceptable levels, thereby highlighting the necessity for further data. This data can either be provided through additional diagnostics or physics-informed priors.

In Chapter 4, a synthetic framework is developed to evaluate inference strategies in the Mega Amp Spherical Tokamak Upgrade (MAST-U) tokamak divertor. D-MIBAS is introduced as a mesh-based Bayesian inference system capable of including multiple diagnostics and incorporating sophisticated physics-informed spatial priors. The analysis highlights the importance of including molecular processes in Balmer line emission models to prevent biased inferences. It demonstrates that physics-based spatial priors, such as the expected monotonic behaviour of electron temperature, can deliver the required information highlighted in Chapter 3, enabling D-MIBAS to produce low inference uncertainties. The chapter further verifies the computational feasibility of a mesh-based approach with thousands of inference parameters and shows that including additional diagnostic models can further minimise inference uncertainty, even at the expense of increasing the number of inference parameters. This chapter is published in:

D Greenhouse et al. *Two-dimensional inference of divertor plasma characteristics: advancements to a multi-instrument Bayesian analysis system*. 2025. DOI 10.1088/1361-6587/adab1b

Chapter 5 transitions the framework toward experimental data. It focuses on MAST-U, which offers a rich set of diagnostics including filtered imaging, spectroscopic systems, Langmuir probes, and Thomson scattering. Experimental analysis introduces new challenges, particularly calibration uncertainties. To address these, the diagnostic models of Chapter 4 are made more flexible, at the cost of increased inference uncertainty. This is mitigated by introducing Gaussian process priors that model plasma fields as samples from smooth distributions governed by interpretable hyperparameters, such as correlation length and amplitude. This approach is enabled by the mesh-based strategy of D-MIBAS. Through the analysis of diverse divertor scenarios, a single set of hyperparameters is proposed to be sufficiently flexible while restricting inference uncertainty.

In Chapter 6, D-MIBAS is applied to a MAST-U experimental discharge with a divertor configuration known to contain both atomic and molecular processes (the Super-X divertor). The analysis covers multiple time points during a core density ramp. Verification and validation efforts are made to contextualise the inferred results.

These efforts evaluate the consistency of the diagnostic models used in the inference, compare the results with alternative diagnostic interpretations, and align them with interpretive SOLPS-ITER modelling. Strong agreement is observed and the results are found to be consistent with our current understanding of the divertor plasma during a core density ramp. These results confirm that D-MIBAS can reproduce key divertor behaviours and offers a powerful tool for studying the divertor.

Chapter 7 assesses the generality of the approach by applying D-MIBAS, with a fixed set of hyperparameters, to a range of divertor configurations (conventional, elongated, and Super-X geometries). The framework performs robustly across these scenarios. However, in regions with limited diagnostic coverage, sensitivity to the chosen hyperparameters remains a practical consideration. These results highlight D-MIBAS as a flexible and reusable analysis tool for other divertor experiments.

1.2.3 Contributions of this Work

This thesis recognises the rich information content of spectroscopic data, particularly the hydrogenic Balmer line series, and its relevance both to present-day tokamaks and to future power-plant scenarios. Building on the mesh-based approach for integrated data analysis introduced in [9] and adapting the Balmer emission model to include molecular contributions from [112], this work develops the Divertor Multi-Instrument Bayesian Analysis System (D-MIBAS). The central result is that experimental inference of the divertor plasma state, with uncertainties sufficiently constrained to reveal robust physical patterns, is achievable with the mesh-based nature of D-MIBAS. This has been verified and validated on MAST-U for an L-mode detached plasma.

A key finding is that information distinguishing whether the plasma is above or below ~ 3 eV is crucial for reducing inference uncertainties when using Balmer spectroscopy. This threshold arises because ionising plasmas and molecularly activated recombining plasmas produce similar emission signatures, while atomically recombining plasmas are more easily distinguished. Such information may be provided by complementary diagnostics or by physics-informed, spatially dependent prior knowledge, both of which can be naturally incorporated in D-MIBAS's mesh-based framework.

A consistent probabilistic approach has been developed for incorporating physics-informed priors. This is achieved by analysing a suite of SOLPS-ITER simulated divertor plasma states and *learning* hyperparameters that govern how prior information is supplied. For example, while a monotonic decrease of electron temperature along flux surfaces toward the target is generally expected, it does not always hold

due to localised plasma effects or inaccuracies in the magnetic equilibrium reconstruction. To account for this, deviations were found to follow a Cauchy distribution; most violations were small, but larger deviations were possible. A corresponding Cauchy probability penalty, with a learned scale hyperparameter, was therefore introduced into the inference. Similarly, while plasma fields are expected to be smooth, they are not flat; in this work, they are modelled as samples from Gaussian processes characterised by spatial correlations dictated by amplitudes, length scales, and noise levels. These hyperparameters were learned from SOLPS-ITER simulations alongside physics-based intuition, such as the known anisotropy of charged-species dynamics along versus across magnetic field lines. In D-MIBAS, these learned hyperparameters provide a physics-informed probabilistic framework for inference.

Early comparisons indicate consistency with predictive models such as SOLPS-ITER and the Detachment Location Sensitivity model. Further validation across a broader set of conditions (H-mode plasmas, attached states, alternative tokamaks, diverse diagnostic sets) is required. This will benefit from better utilisation of existing measurements (e.g. reduced calibration uncertainties, exploitation of Stark broadening in spectroscopic data, and coherence imaging spectroscopy) and from extended diagnostic coverage (e.g. filtered imaging of the X-point region).

Together, these developments establish a systematic approach to compensating for sparse diagnostic coverage in order to provide low-uncertainty inference of the tokamak divertor plasma state. This can be used to inform the design of next-generation fusion devices.

Chapter 2

Background: The Tokamak Divertor

In this chapter, the current understanding of the divertor region is outlined. We consider the quantities that are desired to be known in order to provide a comprehensive description of the divertor. Methods and techniques for integrated data analysis are discussed, including their application in nuclear fusion.

2.1 Atomic and Molecular Physics

The interaction between plasma and neutral species plays a key role in establishing the divertor operating regime and, consequently, the resulting heat and particle loads experienced by plasma-facing components. Atomic and molecular processes govern exchange between ions, electrons, and neutrals and lead to power loss from the plasma via emission of photons. As outlined in Table 2.1, the various volumetric processes can be broken into: ionisation processes (liberation of a bound electron to the plasma continuum); excitation processes (excitation of a bound electron to a higher energy level); charge exchange processes (transfer of an electron from a hot/cold species to a cold/hot species); and recombination processes (bounding an electron from the plasma continuum). The photons subsequently released from these processes typically escape the plasma (in optically dense plasmas, photons can be re-absorbed in the plasma through radiative excitation and photoionisation). Critically, the escaped photons provide volumetric power losses from the plasma. The isotropic radiation spreads this power throughout the divertor chamber, reducing the peak heat flux to the target. Additionally, from a diagnostic view point, these photons contain substantial information as to the plasma state.

Tokamak divertors have been found to contain substantial molecular deuterium populations [49, 116]. Key reactions with molecular deuterium and their impact on particle number are shown in Table 2.2. We utilise changes in particle numbers to define the generic processes: EIR (blue) reduces the electron (ion) number with no change in molecular number; EIE (red) increases the electron (ion) number with no change in molecular number; MAD (purple) processes reduce the molecular number; MAR (green) reduces electron number through MAD; and MAI (cyan) increases the electron (ion) number through MAD.

Table 2.1: Dominant process occurring in the scrape-off layer. Species X and Y include intrinsic impurities, extrinsic impurities, helium and hydrogen. Excited species (*) subsequently relax through photon emission (in recombination, it is possible for species listed as excited to immediately enter the ground state in which case a photon is immediately released).

Reaction		Type	
$e^- + X^{n+}$	$\rightarrow 2e^- + X^{(n+1)+}$	Ionisation	(electron impact)
$X^{n+} + Y^{m+}$	$\rightarrow e^- + X^{(n+1)+} + Y^{m+}$	Ionisation	(ion impact)
$\gamma + X^{n+}$	$\rightarrow 2e^- + X^{(n+1)+}$	Ionisation	(photoionisation)
$e^- + X^{n+}$	$\rightarrow e^- + X^{*n+}$	Excitation	(electron impact)
$X^{n+} + Y^{m+}$	$\rightarrow X^{n+} + Y^{*m+}$	Excitation	(ion impact)
$\gamma + X^{n+}$	$\rightarrow e^- + X^{*(n+1)+}$	Excitation	(radiative)
$X^{n+} + Y^{m+}$	$\rightarrow X^{(n-1)+} + Y^{*(m+1)+}$	Charge Exchange	
$e^- + X^{n+}$	$\rightarrow X^{*(n-1)+}$	Recombination	(two-body, radiative)
$e^- + X^{n+}$	$\rightarrow X^{** (n-1)+} \rightarrow X^{*(n-1)+} + \gamma$	Recombination	(two-body, dielectronic)
$2e^- + X^{n+}$	$\rightarrow e^- + X^{*(n-1)+}$	Recombination	(three-body)

The SOL plasma is comprised of numerous species with varying ionisation and excitation states. The relative populations of different ionisation states and energy levels depend on the balance between these plasma processes. The reaction cross-sections of these processes vary with the energies of the involved species, imparting a temperature dependence to the relative ionisation population levels. In low-density plasmas, generated excited states may radiatively decay spontaneously before collisional charge-exchange processes occur. However, in higher-density plasmas, free

Table 2.2: Hydrogenic processes occurring in the scrape-off layer and their impact on particle number. Processes involving D^- and D_3^+ are omitted.

Reaction 1		Reaction 2		ΔD^+	ΔD	ΔD_2
$e^- + D^+$	$\rightarrow D$			-1	+1	0
$2e^- + D^+$	$\rightarrow e^- + D$			-1	+1	0
$D_2 + D^+$	$\rightarrow D + D_2^+$	$e^- + D_2^+$	$\rightarrow 2D$	-1	0	-1
$e^- + D_2$	$\rightarrow e^- + 2D$			0	+2	-1
$D_2 + D^+$	$\rightarrow D + D_2^+$	$e^- + D_2^+$	$\rightarrow e^- + D + D^+$	0	+2	-1
$e^- + D_2$	$\rightarrow 2e^- + D_2^+$	$e^- + D_2^+$	$\rightarrow 2D$	0	+2	-1
$e^- + D$	$\rightarrow 2e^- + D^+$			+1	-1	0
$D + D^+$	$\rightarrow e^- + 2D^+$			+1	-1	0
$D_2 + D^+$	$\rightarrow D + D_2^+$	$e^- + D_2^+$	$\rightarrow 2e^- + 2D^+$	+1	0	-1
$e^- + D_2$	$\rightarrow 2e^- + D + D^+$			+1	+1	-1
$e^- + D_2$	$\rightarrow 2e^- + D_2^+$	$e^- + D_2^+$	$\rightarrow e^- + D + D^+$	+1	+1	-1
$e^- + D_2$	$\rightarrow 2e^- + D_2^+$	$e^- + D_2^+$	$\rightarrow 2e^- + 2D^+$	+2	0	-1

electrons can contribute to recombination (as in three-body recombination) and de-excitation (with the de-excitation energy being transferred to an additional electron rather than emitted as a photon). Consequently, the relative populations of ionisation states and energy level populations also depend on the plasma's electron density.

The local temperature and electron density can be used to model the population densities of the energy levels of various species present. This is performed by collisional-radiative models such as ADAS¹ [90], which are widely used for atomic modelling in fusion. Based on these population densities, collisional-radiative models predict the volumes and densities required to provide sufficient power and manage ion losses through ionisation and recombination processes, essential for detachment. Figure 2.1 illustrates these interconnected processes under divertor-relevant conditions. There is a strong temperature dependence on the dominance of each process, with an increase in reaction rate as electron density rises. Additionally, as indicated by the arrows, the EIE rate scales approximately linearly with atomic neutral density, n_D , while MAD, MAR and MAI rates each scale approximately linearly with molecular density, n_{D_2} .

¹ADAS is a *generalised* collisional radiative model which includes components such as metastables [102].

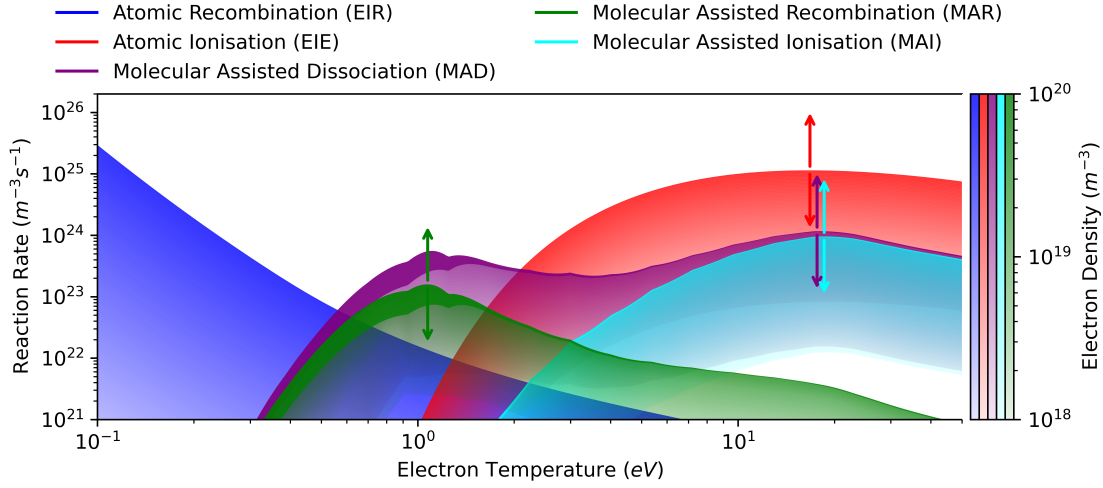


Figure 2.1: Reaction rates of collated hydrogenic atomic and molecular processes using ADAS[90] and Yacora[134] at $T_e n_D = 10^{20} \text{ eV m}^{-3}$ and $T_e n_{D_2} = 10^{19} \text{ eV m}^{-3}$. Shading transparency corresponds to electron density. Additional shading of MAD and MAI at high electron densities indicative of uncertainties surrounding the molecular rates. Upward/downward arrows indicate the impact on reaction rates due to: a ten times increase/decrease in n_D for atomic ionisation (red); and a ten times increase in n_{D_2} for MAD (purple), MAR (green) and MAI (cyan). MAD rates include MAR and MAI rates.

2.1.1 Photon Emission

Collisional-radiative models offer significant diagnostic capabilities by linking measurable photon emissions to critical plasma state parameters, such as temperature and density. Collisional-radiative models like ADAS utilise photon emissivity coefficients (PECs) for transitions between energy states p and q , such that

$$PEC_{p \rightarrow q}(n_e, T_e) = A_{pq} R_p(n_e, T_e) = A_{pq} \frac{n_p}{n_e \sum_i n_i}. \quad (2.1)$$

A_{pq} is the Einstein coefficient for the transition between energy states p and q . R_p is the population coefficient for state p , which includes the population density of state p , n_p , the electron density n_e , and the sum of the densities of all excited levels i of the species.

Hydrogenic $n \rightarrow 2$ transitions, known as the ‘Balmer series’, emit photons in the visible spectrum. The radiative emissivity of these photons can be modelled through,

$$\varepsilon_{n \rightarrow 2} = \sum_{\rho} n_{\rho, X} n_{\rho, Y} PEC_{n \rightarrow 2}^{\rho}(T_e, n_e). \quad (2.2)$$

ρ represents the processes leading to an excited state, n . $n_{\rho,X}$, and $n_{\rho,Y}$ denote the densities of the relevant species involved in the process.

Crucially, the PECs corresponding to different n Balmer transitions have varying dependencies on T_e and n_e for atomic and molecular recombination and excitation processes. As shown in Figure 2.8 of [119], higher n Balmer lines exhibit a greater relative ratio of recombination processes to excitation processes compared to lower n Balmer lines. Consequently, ratios of Balmer lines can be used to determine the relative ratio of recombination to excitation, as discussed in Section 7.2 of [119]. Similarly, MAR processes predominantly populate lower energy levels and thus drive a higher ratio of lower n Balmer lines to higher n Balmer lines. Given the T_e and n_e dependencies in atomic and molecular recombination and excitation processes, these line ratios offer insights into the plasma state. This is investigated in more detail in Chapter 3.

2.2 The Divertor

In this section we introduce the divertor as a means of heat and particle exhaust. Divertor plasmas have been observed to exist in different regimes which are governed by distinct dynamics. These are outlined such that our inferred results can be qualitatively compared to them. We outline how the divertor can be configured to theoretically achieve optimal target conditions. These configurations are actively being researched in tokamaks such as MAST-U, and their impact on the plasma state is what we wish to be able to infer. Simplified analytic models describe the divertor plasma state, and they are outlined in this chapter so that our inferred plasma states can be compared with them.

2.2.1 Divertor Regimes

The divertor operating regime and target conditions are strongly coupled, with each influencing the other. A useful distinction between divertor regimes is the degree of plasma recycling, whereby plasma ions undergo surface recombination at material surfaces, return to the plasma as neutral particles, and are subsequently re-ionised within the divertor plasma.

2.2.1.1 Sheath Limited Regime

In the low-recycling regime, divertor dynamics are dominated by the plasma sheath (see Section 1.1.3). With limited ionisation in the divertor plasma, there is little

local particle source and the plasma density remains approximately constant along the SOL. At low densities, the ionisation mean free path of recycled neutrals can exceed the SOL length, limiting plasma-neutral interactions within the divertor. Thus, the sheath-limited, low-recycling regime provides little mechanism for substantial volumetric momentum or power loss. Consequently, this regime does not achieve a substantial reduction in target particle flux, heat flux, or temperature.

2.2.1.2 Conduction Limited Regime

If the amount of recycling increases beyond the sheath-limited regime, the target electron density can increase approximately quadratically with increasing upstream density [100]. Ionisation of recycled neutrals introduces additional volumetric power losses within the divertor, allowing a significant temperature drop to develop between the upstream plasma and the divertor target. Although the increased recycling reduces the target temperature, momentum losses remain weak and the plasma pressure is therefore approximately conserved along the SOL. Consequently, the reduction in target temperature must be accompanied by an increase in target density. From equation (1.5), the target ion flux scales as $\Gamma_t \propto n_e T_e^{1/2}$. Therefore, the increased target density leads to high particle fluxes to the target despite the lower temperatures. Thus, the conduction-limited, high-recycling regime can achieve low target temperatures while retaining high target particle fluxes.

2.2.1.3 Detached Regime

The detached regime enables both a reduced target temperature and particle flux. As described in Section 2.1, hydrogenic ionisation processes reduce sharply at temperatures $\lesssim 5eV$. Therefore, if the target temperature can reach such temperatures, high neutral populations can propagate away from the target and be sustained by volumetric recombination. This increase in neutral species can increase the neutral drag and momentum transfer from the plasma. Consequently, the upstream- and downstream-pressure relationship is broken. This allows the target temperature to decrease further without requiring an associated increase in target density and particle flux. Atomic recombination processes dominate over ionisation processes at temperatures $\lesssim 1eV$. Consequently, with considerable volumetric recombination present, the target electron density can start to decrease giving a ‘roll-over’ in ion flux to the target. In such a regime, the plasma is said to be ‘detached’ from the target.

This detached regime has been widely observed in experiments; however, it appears that molecules also play a role in the form of molecular-activated dissociation

(MAD). It has been observed [116] that during detachment, a region of ionising plasma first moves away from the target (due to the neutral populations propagating further from the target when the temperature remains low). Between this ionisation front and the target exists a region in which deuterium molecules can dissociate and assist in the neutralisation of deuterium ions (molecular activated recombination, MAR). MAR reactions allow additional power to be radiated from the plasma, lowering the electron temperature such that electron-ion recombination (EIR) can take place. EIR radiates additional power from the plasma, further reducing the plasma temperature and the number of charged particles streaming along field lines and striking the target.

Similar results have been observed in L-mode detachment at JET [49]. [49] describes that when the strike point electron temperature exceeds 2 eV, there is substantial electron impact ionisation (EIE) under high recycling conditions in which ions neutralise at the wall (surface-recombination), return as neutrals and are subsequently re-ionised generating a charged-particle source. When the strike point temperature drops (to approximately between 0.9 and 2 eV), atomic ionisation reduces but the presence of molecular deuterium brings molecular activated ionisation (MAI) alongside the volumetric-recombination processes MAR and EIR. This results in a plasma that is neither significantly net recombining nor net ionising. A further reduction in the strike point temperature, <0.9 eV, brings dominating EIR with diminishing MAR and MAI creating a net-particle sink. This suggests an intrinsic relationship between molecular-activated dissociation (MAD) processes and the change of target conditions.

2.2.2 Divertor Configuration

The magnetic configuration can be adapted to direct the SOL such that it reaches the ‘target’ PFC in a region optimised for heat and particle exhaust. This is achieved in a divertor as shown in Figure 2.2(a). By bringing the poloidal magnetic field to zero (a magnetic ‘X-point’), the magnetic field lines can change direction and be diverted away from the core plasma. The X-point splits the plasma into: common flux region(s) (CFR) with magnetic field lines that reach the mid-plane such that plasma entering the SOL from the core can readily stream throughout the CFR; and private flux region(s) (PFR) which is only populated by plasma transported across flux surfaces from the CFR. The ‘separatrix’ describes $\psi = 1$ which specifies both the LCFS and the transition between the PFR and CFR.

Manipulation of the divertor configuration has numerous advantages in handling the interaction between the plasma and PFC. The tokamak MAST-U has adaptable

divertor magnetic fields such that it is able to generate different configurations. In Figure 2.2(a), three such configurations are shown: the ‘open’ divertor (OD); the ‘elongated’ divertor (ED); and the ‘Super-X’ divertor (SXD).




Figure 2.2: Poloidal cross-section of the MAST-U tokamak divertor.

a) Divertor configurations in the MAST-U tokamak divertor. Adapted from Figure 1 of [122].

b) Schematic of a flux-tube in a Super-X divertor projected onto the 2D poloidal cross section.

These configurations explore different critical aspects of the divertor that can mitigate heat and particle fluxes to the target wall (see Figure 2.2(b)). Some of these aspects are conveniently described by the ‘flux-tube’ construct: a theoretical region bounded by magnetic field lines, within which charged species predominantly flow parallel to the magnetic field and their numbers are approximately conserved. Changes in the magnetic field can alter the shape of this tube. Importantly, increased poloidal flux expansion (due to a reduction in the poloidal magnetic field) widens the flux tube in the poloidal cross section, but does not alter the cross-sectional area of the flux tube unlike total flux expansion (see Section 2.2.2.3).

As described in Section 2.2.1.3, beyond a threshold for the loss of power, particles, and momentum within the plasma itself, the detached operating regime occurs. This favourable regime is characterised by low target temperatures (below a few electron volts) and is accompanied by reduced heat and particle fluxes to the target wall. This threshold can be surpassed through increased upstream densities and impurity seeding; however, these can lead to poor core plasma performance due to, respectively, exceeding the Greenwald limit or core energy loss from impurity contamination in the core plasma. As discussed in this section, changes in the divertor configuration can

lower the required threshold for detachment access. Analytic studies [17, 67] have shown that the nature of the magnetic field in the divertor can be used to control the location at which the divertor plasma transitions to the low temperatures synonymous with detachment. Control of this point is important for optimal divertor performance.

2.2.2.1 Connection Length

The toroidal magnetic field, which is dominant relative to the poloidal magnetic field, causes flux tubes to wrap around the tokamak torus approximately one to five times before intersecting material surfaces. The length of the magnetic field along these flux tubes from the tokamak mid-plane to the interception with the target PFC (the ‘strike point’) is the ‘magnetic connection length’, L_{\parallel} .

An increase in toroidal magnetic field strength, relative to the poloidal magnetic field, can increase the magnetic connection length and consequently the distance travelled by charged species before striking the target. This gives greater possibility for key processes within the plasma (i.e. *volumetric*): excitation, ionisation, charge exchange, and volumetric recombination. These processes remove power, particles, and momentum from the plasma, thereby reducing the target heat flux, ion flux, and temperature for given upstream conditions.

2.2.2.2 Grazing Angle

The magnetic field can be adapted so that the flux tubes meet the PFC at a shallow ‘grazing angle’, α . Consequently, particle and heat fluxes held in the flux tube do not strike the target along the target area’s normal. Thus, the parallel heat and particle flux being deposited to the PFC is reduced through,

$$\Gamma_{t,dep.} = \Gamma_t \sin \alpha, \quad (2.3)$$

and,

$$q_{t,dep.} = q_t \sin \alpha. \quad (2.4)$$

Because the toroidal magnetic field is substantially greater than the poloidal magnetic field, the flux tube shown in Figure 2.2(b) meets the target wall at a shallow angle in both the poloidal and the toroidal plane. Thus, α can be less than five degrees², giving a reduction of more than ten times in the deposited particle and heat flux.

²Engineering limitations typically prevent α from being reduced below around two to three degrees[136].

2.2.2.3 Total Flux Expansion

As the magnetic field strength, $B = |\mathbf{B}|$, reduces, the area of the cross sections of the flux tubes must increase. This increase in area, ‘total flux expansion’, is given through,

$$F_{R,m} = \frac{B_m}{B(s_{\parallel})} \quad (2.5)$$

where B_m is the magnetic field strength at the outer midplane and s_{\parallel} is the distance along the magnetic field line from the outer midplane. Total flux expansion can also be defined relative to the magnetic x-point (rather than the outer midplane) which we refer to as $F_{R,x}$.

The particle fluxes of charged species are conserved throughout each flux tube when there are no ion sources, ion sinks, or perpendicular ion transport. It follows that, in the absence of sinks, the particle flux density scales with $F_{R,m}^{-1}$. The toroidal magnetic field scales approximately with the reciprocal of the major radius. Consequently, total flux expansion is readily achieved by diverting the plasma to a larger major radius, as is performed in the Super-X divertor. Theoretically, Super-X divertors lead to reduced particle and heat flux densities to the target.

The influence of total flux expansion on the Super-X divertor plasma has been investigated through SOLPS-ITER modelling in [77]. As indicated by equation (1.7), a reduction in the particle flux density arriving at the target implies a lower target density, a lower target temperature, or both. The different effects depend on the existing divertor regime (see Section 2.2.1) and can affect the transition between these regimes. For example, as predicted in [77], increased total flux expansion can lower the threshold to access the desirable detached regime.

2.2.2.4 Baffled, Closed Configuration

Studies on TCV showed that increased total flux expansion alone did not reduce the detachment threshold as expected [110]. Plasma–surface interactions generate hydrogenic neutral species (through surface recombination) and impurity species (predominantly carbon in the carbon-armoured TCV and MAST-U divertors), both of which play a key role in detachment. They do so by removing plasma momentum (via charge exchange with ions and molecular reactions) and plasma energy (via ionisation and radiative power losses). For these processes to effectively reduce the target fluxes, the charge-exchange and ionisation mean free paths must be comparable to or shorter than the divertor leg length, ensuring that momentum and power loss processes occur predominantly within the divertor volume [30]. Closed, baffled divertor geometries

improve neutral confinement and increase neutral residence time, thereby increasing neutral pressure and reducing charge-exchange and ionisation mean free paths. As emphasised in Section 4.6. of [54], closed divertor geometries thus promote higher neutral pressures and facilitate reliable access to detachment.

2.2.3 Analytic Models of the Scrape-Off Layer

Analytic models offer simplified descriptions of heat, particle, and momentum transport through the scrape-off layer (SOL) plasma. They capture fundamental relationships which are valuable for informing the physics-based prior knowledge used in D-MIBAS. At the same time, the probabilistic reconstructions produced by D-MIBAS can be used to test the consistency and validity of these analytic models against experimental data.

The SOL can be considered to have a characteristic width, λ_{SOL} , and length L_{\parallel} . Perpendicular to the magnetic field lines ('cross-field'), the plasma spreads radially outward through turbulent transport and diffusive processes. The density-profile width of the SOL, λ_{SOL} , is approximated by

$$\lambda_{SOL} = \sqrt{\frac{D_{\perp} L_{\parallel}}{c_s}}. \quad (2.6)$$

D_{\perp} defines the cross-field diffusion of particles. Whilst numerous processes contribute to the cross-field spreading of the plasma, D_{\perp} is often defined from Fick's law, such that the cross-field flux of particles is given by,

$$\Gamma_{\perp} = -D_{\perp} \frac{dn_e}{ds_{\perp}}, \quad (2.7)$$

where s_{\perp} tracks the cross-field displacement.

2.2.3.1 Perpendicular Profile: Eich Function

A total power of P_{SOL} can be considered to enter at the outer midplane. This power is carried conductively along magnetic field lines with heat flux density, $q_{\parallel,m}$, and is spread over a characteristic width at the midplane, $\lambda_{q,m}$, such that

$$q_{\parallel,m} \approx \frac{P_{SOL}}{2\pi R_m \lambda_{q,m}} \quad (2.8)$$

where R_m is the major radius at the midplane.

At the midplane, the heat flux entering the SOL, can be modelled to exponentially decay cross-field (toward larger ψ) with characteristic lengthscale $\lambda_{q,m}$. As the heat

flux progresses toward the target, diffusive processes further spread the heat flux cross-field. This is parameterised by the Eich-function [23], which models the parallel heat-flux density reaching the target by:

$$q_{\parallel}(\bar{s}) = \frac{q_{\parallel,m}}{2} \frac{1}{F_{R,m}} \exp \left(\left(\frac{S}{2\lambda_{q,m}F_{R,m}} \right)^2 - \frac{\bar{s}}{\lambda_{q,m}F_{R,m}} \right) \operatorname{erfc} \left(\frac{S}{2\lambda_{q,m}F_{R,m}} - \frac{\bar{s}}{S} \right) + q_{\parallel,BG}. \quad (2.9)$$

S gives the typical length scale for the cross-field diffusive processes in the divertor and $q_{\parallel,m}$ and $q_{\parallel,BG}$ are respectively the mid-plane and background parallel heat flux densities. \bar{s} is the displacement from the separatrix at the target (which is where the $F_{R,m}$ is evaluated) such that $\bar{s} = s_{\perp} - s_0$ with s_0 being the separatrix position at the target (the strike point). In this form of the Eich-function, $\lambda_{q,m}$ is defined at the midplane, with the characteristic heat flux decay width at the target approximately $\lambda_{q,m}F_{R,m}$.

The shape of the target heat flux profile is conveniently described with respect to the maximum heat flux density at the target, $q_{\parallel,max.}$, via the integral decay length:

$$\lambda_{int} = \frac{\int q_{\parallel}(s_{\perp}) - q_{BG} ds_{\perp}}{q_{\parallel,max.} F_{R,m}}. \quad (2.10)$$

Here, q_{BG} is the background heat flux. The integral decay length accounts for both the diffusive broadening S and the exponential fall-off mapped to the target, $\lambda_{q,m}F_{R,m}$. A higher value of λ_{int} is advantageous, as it spreads the heat flux over a larger area.

The Eich function can be used to illustrate how divertor magnetic topology shapes the heat-flux profile at the target, with characteristic profiles shown in Figure 2 of [23]. An increase in total flux expansion spreads the heat flux profile over a larger distance from the separatrix, thereby reducing the peak heat flux. Similarly, increasing the magnetic connection length or raising the divertor neutral pressure enhances volumetric and diffusive processes, which increase the power-spreading parameter S and further broaden the profile. From equation (2.4), a smaller grazing angle also directly reduces the deposited heat flux via the geometric projection factor. These effects are not mutually exclusive and, when combined with additional volumetric losses from detachment, can yield orders-of-magnitude reductions in the peak heat load to the target [5, 65].

2.2.3.2 Parallel Profile: The Two-Point and Lengyel Models

The SOL's length can be simplified into two points: the upstream, u , and the target, t , (at the sheath entrance). The exact position of the upstream is an arbitrary position

between the X-point and mid-plane (since plasma parameters are not expected to vary significantly between the X-point and mid-plane). As discussed in Section 5.2 of [100], an analysis of particle, pressure, and power balance leads to a set of coupled equations known as the ‘two-point model’. These relate the upstream and target conditions:

$$n_u T_u = 2n_t T_t, \quad (2.11)$$

$$T_u^{7/2} = T_t^{7/2} + \frac{7}{2} \frac{q_{\parallel,t} L_{\parallel}}{\kappa_{0,e}}, \quad (2.12)$$

and,

$$q_{\parallel} = \gamma n_t T_t c_{s,t}. \quad (2.13)$$

$\kappa_{0,e}$ is the electrical conductivity of the plasma.

The presence of neutrals and impurity species in the scrape-off layer (SOL) facilitates momentum and power transfer from the plasma. Ionisation and recombination processes generate, respectively, ion sources and sinks throughout the SOL, promoting plasma flows and enhancing heat convection. The ‘extended two-point model’ captures these through: the fraction of momentum lost from the plasma, $1 - f_{\text{mom.}}$, due to neutral drag (plasma viscosity) and recombination processes; the fraction of heat flux density carried by convection, $1 - f_{\text{cond.}}$; and the fraction of power losses from volumetric processes, $f_{\text{pow.}}$.

$$n_u T_u = \frac{1}{f_{\text{mom.}}} 2n_t T_t, \quad (2.14)$$

$$T_u^{7/2} = T_t^{7/2} + \frac{7}{2} \frac{f_{\text{cond.}} q_{\parallel,t} L_{\parallel}}{\kappa_{0,e}}, \quad (2.15)$$

and,

$$q_{\parallel} = \frac{\gamma n_t T_t c_{s,t}}{1 - f_{\text{pow.}}}. \quad (2.16)$$

The extended two-point model gives insight into the relationship between upstream and target conditions. The Lengyel model [62] (which is built on the Goedheer models found in [35]), informs how these parameters may develop between the upstream and target. Through power and particle balance (including volumetric processes ionisation, charge exchange and recombination) a series of coupled differential equations are derived (see equations (2.1) to (2.6) of [62]). Solutions to these equations provide plausible profiles of key plasma characteristics along a SOL flux tube. A qualitative example of a solution to these equations is shown in Figure 2.3 (adapted from Figure 11 of [62]).

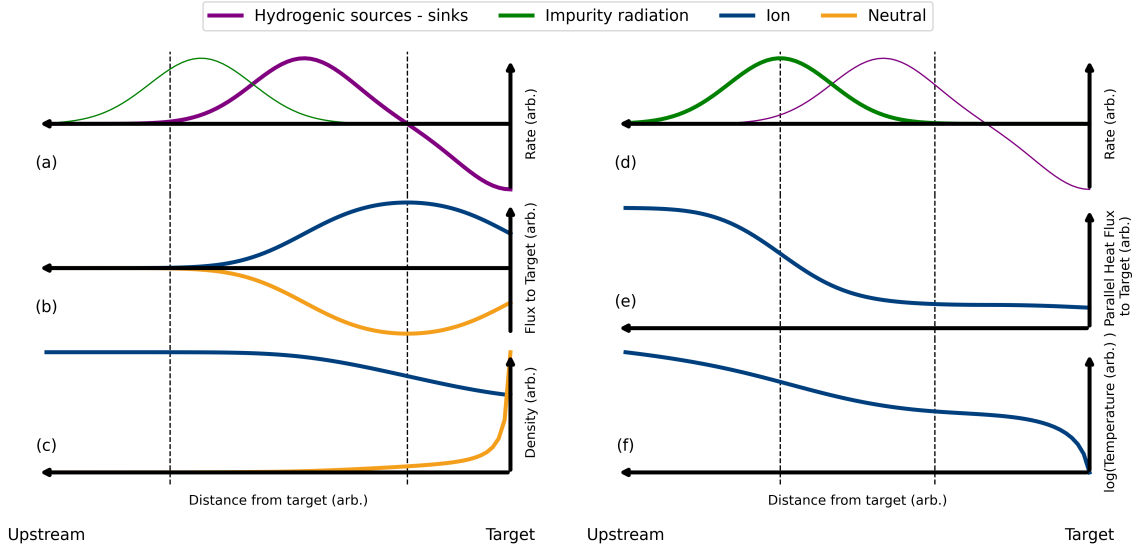


Figure 2.3: The Lengyel model. Qualitative example solution to equations (2.1) to (2.6) of [62].

Adapted from Figure 11 of [62].

In Figure 2.3(a), we observe a characteristic form for hydrogenic sources and sinks in a detached divertor state: ionisation (hydrogen ion source) dominates away from the target while recombination (hydrogen ion sink) dominates closer to the target. This results in a ‘roll-over’ in ion flux to the target (Figure 2.3(b)); from particle number conservation, the ion flux increases due to the net ion source before reducing due to the net ion sink. The inverse is true for the neutral hydrogen flux. As seen in Figure 2.3(c) the neutral population is expected to be negligible upstream (which is the driving cause for the lack of ionisation far upstream) and large close to the target. The non-zero ion flux is accompanied by a slight reduction in the ion density (equivalent to the electron density in a hydrogen-dominant plasma). For plasmas with substantial impurity densities, it is expected that the losses from impurity radiation (Figure 2.3(d)) drive the majority of the reduction in parallel heat flux (Figure 2.3(e)). This reduction in heat flux is accompanied by a reduction in plasma temperature (Figure 2.3(f)). The substantial hydrogenic recombination close to the target is also expected to lead to a reduction in temperature.

2.3 Insights into the Tokamak Divertor Plasma

The analytic models of Section 2.2.3 offer a simplified view of plasma characteristics in the scrape-off layer. However, understanding detachment’s complexities necessi-

tates a two-dimensional (and potentially three-dimensional) representation of various plasma species. To achieve this, high-fidelity models and comprehensive experimental data are essential. Current modelling capabilities, however, exhibit known inaccuracies and uncertainties. Despite the availability of multiple diagnostics, none provide direct, two-dimensional measurements of the plasma characteristics. Consequently, validating simulations against experimental data can be complex.

In this section, we outline how key quantities for describing the divertor plasma can be derived from two-dimensional descriptions of T_e , n_e , n_D , and molecular hydrogen density, n_{D_2} . We outline how these quantities are found in high-fidelity models and which diagnostics give insight into them.

2.3.1 Desired Quantities

The tokamak exhaust problem necessitates that target ion particle flux, heat flux density, and temperature are minimised. Therefore, we aim to find the particle flux, Γ_i , heat flux density, q , and temperature throughout the divertor's two-dimensional poloidal cross-section. Furthermore, the plasma's momentum (pressure) is desired to provide a further understanding of the detached divertor state.

In a hydrogen-dominant plasma, the effective charge of the plasma, Z_{eff} , indicates that the densities of electrons, e , and ions, i , are equal ($n_e \approx n_i$). Consequently, the ion particle flux is given by,

$$\Gamma_i \approx \Gamma_{i,\parallel} \approx n_e v_{i,\parallel} \approx \Gamma_e. \quad (2.17)$$

$v_{i,\parallel}$ is the ion velocity parallel to magnetic field lines.

The heat flux comprises conductive and convective components. The conductive component is predominantly carried by electrons, while the convective component is primarily carried by ions. Under the assumption that $T_e = T_i$, the heat flux is given by,

$$q \approx q_{\parallel} \approx \underbrace{-\kappa T_e^{5/2} \nabla_{\parallel} T_e}_{\text{Conducted}} + \underbrace{\left(5T_e + \frac{1}{2} m_i v_{i,\parallel}^2\right) \Gamma_{i,\parallel}}_{\text{Convected}}, \quad (2.18)$$

where $\nabla_{\parallel} T_e$ is the gradient of the electron temperature along a magnetic field line, and κ is the thermal conductivity coefficient of the plasma.

The particle balance can be employed to determine the ion velocity, as outlined in Appendix C of [114]. Consider a volume of plasma shaped like the frustum of a cone (a flux tube element) as illustrated in Figure 2.4. Using particle balance,

$$\underbrace{\int_{A_1} \Gamma_{\parallel, \text{out}}}_{\text{Downstream}} = \underbrace{\int_{A_0} \Gamma_{\parallel, \text{in}}}_{\text{Upstream}} + \underbrace{\int_V (S_+ - S_-)}_{\text{Particle Sources - Sinks}} + \underbrace{\int_{A_2} \Gamma_{\perp, \text{in (net)}}}_{\text{Cross-field}}, \quad (2.19)$$




Figure 2.4: Flux tube element obeying particle balance (illustration of Gauss's theorem). S_+ (S_-) represents the rate of particle sources (sinks) per unit volume.

where quantities are as defined in Figure 2.4. Ion sources and sinks, $S_{+/-}$, can be calculated as functions of temperature and density through the collisional-radiative models used to generate Figure 2.1 (see Section 2.1). The target particle flux is derived from equation (1.7). Therefore, with one additional assumption, equation (2.19) can be used to find the parallel particle flux throughout the divertor chamber as a function of electron temperature and electron, atomic neutral, and molecular densities (necessary to identify additional ion sources and sinks from MAR and MAI processes). This assumption can either be that the upstream parallel ion flux at the midplane is negligible, or, as is assumed in this work, that the cross-field particle flux is negligible. The assumption of one enables the derivation of the other.

The plasma pressure consists of static and dynamic components. Assuming that the bulk flow velocity is predominantly parallel and that neutral, molecular, and dynamic electron contributions are minimal, the total plasma pressure can be calculated using:

$$p \approx \underbrace{n_e T_e + n_i T_i}_{\text{Static}} + \underbrace{m_i n_i v_{i,\parallel}^2}_{\text{Dynamic}}. \quad (2.20)$$

In thermodynamic equilibrium, $T_e = T_i$. However, thermodynamic equilibrium is not always maintained in the divertor; the presence of currents and magnetic gradients can result in different electron and ion temperatures.

In conclusion, a comprehensive understanding of the divertor, including plasma pressure, particle and heat fluxes and temperature, can be derived as functions of: electron temperature T_e ; electron density n_e ; atomic neutral hydrogen density n_D ; and molecular density n_{D_2} . Therefore, these are the quantities sought by D-MIBAS in this thesis. To minimise assumptions, separate measurements of ion and electron temperatures, as well as direct measurements of electron velocity, are beneficial. Additionally, to enhance understanding of neutrals, measurements of neutral temperature

and flow velocity would be advantageous. However, these quantities fall outside the scope of this work.

2.3.2 Plasma Edge Modelling

High-fidelity, multi-physics modelling of the scrape-off layer offers a great insight into the divertor plasma. The modelling impacts D-MIBAS in two ways. Firstly, it provides a clear point of comparison and validation. Secondly, it can be used to inform some of the prior (see Section 2.4 assumptions used).

The scrape-off layer plasma is complex to model due to the various species, with neutral and charged particles exhibiting different dynamics. The widely used SOLPS-ITER [132] combines a fluid solver (B2.5) for charged species movement with a Monte-Carlo kinetic code (EIRENE) for neutral species transport [53, 94]. It incorporates the rates of various atomic and molecular volumetric processes (Section 2.1) as well as plasma-surface interactions. SOLPS-ITER offers a comprehensive description of the densities, temperatures, and velocities of the different species. These parameters enable the calculation of particle and heat fluxes to the target walls.

SOLPS-ITER has demonstrated detached divertor states [129] and informed the design of tokamak divertors [56, 88] while investigating the likely operating scenarios of future planned devices [101]. However, significant uncertainties exist regarding the rates of various atomic and molecular processes, as well as the extent of transport and drifts. To address these uncertainties, a method is needed for easy comparison between SOLPS-ITER simulations and experimental data.

The ADAS, AMJUEL, and YACORA atomic and molecular rate coefficients used by SOLPS-ITER are known to contain non-negligible uncertainties arising from several sources, including incomplete or sparse experimental benchmark data, cross-section calculation and collisional-radiative model assumptions, and incomplete datasets requiring extrapolation beyond validated temperature and density regimes [6, 134]. Recent studies have demonstrated that discrepancies in commonly used AMJUEL molecular rates can significantly alter predicted detachment behaviour in SOLPS-ITER, motivating ongoing efforts to replace or supplement these rates with updated YACORA-derived collisional-radiative data [11]. These uncertainties directly affect the source and sink terms governing ionisation, recombination, excitation, charge exchange, and molecular interactions within edge plasma simulations.

Beyond atomic and molecular data limitations, SOLPS-ITER also adopts a time-averaged fluid description of the plasma, meaning that turbulent fluctuations, intermittent transport events, and kinetic non-equilibrium effects are not fully captured.

Consequently, SOLPS-ITER primarily provides a mean-field representation of divertor conditions rather than a complete fluctuating plasma picture. This limitation is particularly relevant when interpreting diagnostics sensitive to transient or filamentary transport processes. Alternative edge modelling frameworks, such as HERMES [22], are under active development to better incorporate aspects of turbulence and plasma fluctuations.

As a result, SOLPS-ITER's predictive capability remains partially limited. This has two primary consequences for the present work. Firstly, when SOLPS-ITER simulations are used as synthetic ground truth for diagnostic or inferential validation, discrepancies between simulation and experimental reality are expected due to imperfect rate data and broader model assumptions. Secondly, when SOLPS-ITER-derived plasma states are used to inform priors within the integrated inference framework, these uncertainties may propagate systematic biases into posterior estimates if not appropriately accounted for.

2.3.3 Diagnostics

There are multiple diagnostics routinely used in tokamak divertors. To enhance the understanding of processes occurring in the divertor, diagnostics that provide good spatial resolution and 2D descriptions of key plasma characteristics without complex, error-prone analysis are desirable but currently non-existent. For example, bolometry systems [25] provide information on the total radiated power from the plasma; however, due to the presence of many radiating species in a divertor, specifics of the plasma state are difficult to extract from total power information. Infrared thermography cameras can monitor the heat flux delivered to the target (a key characteristic of plasma); however, their analysis requires sophisticated methods [4].

Cameras, spectrally filtered for different atomic transitions (such as MANTIS at TCV [86] and the MWI at MAST-U [28]), oriented both tangentially and radially through the plasma, provide substantial information across the entire divertor region. Tomographic inversion is commonly performed to retrieve two-dimensional poloidal cross-sections of localised emissivity from various atomic transitions [133]. Various transitions occur under characteristic, informative plasma conditions. For instance, the low temperature tail of the Fulcher band (further discussed in section 3.1.1) is often associated with the detachment front, allowing these cameras to track its movement [52].

Coherence imaging spectroscopy (CIS) can be combined with spectrally-filtered cameras [21, 72, 97]. CIS utilises phase delays to interrogate emission line broadening

mechanisms, allowing for the extraction of two-dimensional measurements of impurity ion velocity, electron density, and neutral temperature. The broadening mechanisms necessary for extracting these plasma characteristics dominate in different regimes. For example, Stark broadening provides information on electron density; however, Doppler effects prevail at high temperatures, which increases the uncertainty of electron density measurements [72]. Additionally, complications may arise from impurity contamination and background emission.

Spectroscopy systems with multiple lines of sight are routinely fanned through a poloidal plane. The line-integrated emission information across a wide range of wavelengths can be interpreted to provide one-dimensional insight into plasma characteristics from the divertor entrance to the divertor target [112]. High-resolution spectroscopy can interrogate the broadening of individual atomic transitions and extract density and temperature information [83, 98].

Thomson scattering systems routinely recover measurements of electron temperature and density at discrete spatial locations along the laser beam path. By adjusting the divertor's magnetic equilibrium in relation to these localised measurements through a process known as 'strike-point sweeping', a two-dimensional image can be constructed [3, 12, 13], as shown in Figure 2.6a. However, these techniques require constant upstream conditions and accurate magnetic equilibrium information. These restrictions can limit spatial resolution and, as discussed in [12], lead to calibration challenges. The DIII-D tokamak employs a mirror system to redirect the laser beam throughout a shot, generating a 2D Thomson scattering system [96]. However, the time taken to redirect the beam and collect the scattered light for each path limits the overall system frequency to just 5 Hz.

Langmuir probes, ubiquitous in current tokamaks, are installed along tiles throughout the divertor. They measure the ion current to the target walls. Variants [18], along with advanced analysis, can yield measurements of ion and electron temperatures, as well as insights into electron density and heat flux [100]. However, the harsher conditions in future power-plant devices necessitate alternative methods to study heat and particle fluxes to the target walls.

2.4 Integrated Data Analysis

With multiple diagnostics contributing to the overall picture, integrated data analysis is essential for developing a comprehensive description of the plasma state in the divertor. Data analysis can be approached using either a *frequentist* or *Bayesian*

methodology. These frameworks differ in their interpretation of the measured data, \mathcal{D} , and the corresponding set of parameters, θ , which define the underlying *latent* state of the system.

The frequentist approach treats underlying parameters as fixed but unknown quantities. The central question becomes: what is the likelihood of observing the data given these parameters, $P(\mathcal{D}|\theta)$? Measured data are considered a sample from a broader population and inherently contain uncertainty due to sampling variability. Consequently, large datasets are often required to reduce this variability and produce reliable estimates. However, acquiring such extensive datasets is not feasible in many practical situations. One strategy to increase the effective size of the dataset is to combine data from multiple diagnostics. However, heterogeneous diagnostics often rely on different assumptions and measurement models, meaning their samples may not be drawn from the same underlying population. This significantly complicates frequentist integrated data analysis and may lead to biased or misleading interpretations.

The Bayesian approach inverts the problem by treating the measured data as fixed and the underlying parameters as random variables. The key question therefore becomes: what is the probability of a particular set of parameters (i.e. the plasma state) given the measured data, $P(\theta|\mathcal{D})$? This formulation accommodates the inclusion of multiple diagnostics by considering the joint probability of the plasma state based on all available measurements. The advantage of treating θ as a random variable is clear in this context: instead of identifying a single *true* parameter value (which may be poorly defined or conflicted between inconsistent datasets) the goal is to infer which parameter values are more plausible than others. This inversion is achieved through Bayes' theorem,

$$\underbrace{P(\theta|\mathcal{D})}_{\text{posterior}} = \frac{\underbrace{P(\mathcal{D}|\theta)}_{\text{likelihood}} \times \underbrace{P(\theta)}_{\text{prior}}}{\underbrace{P(\mathcal{D})}_{\text{evidence}}}. \quad (2.21)$$

A choice must be made regarding the probability distribution used to evaluate the likelihood. Each choice can introduce subjectivity into the analysis. Consequently, it is beneficial to define generic functions that are parameterised by hyperparameters α . These hyperparameters can be treated as non-fixed, random variables, thereby removing the subjectivity associated with the choice. A common generic likelihood

distribution is the Student's t-distribution,

$$P(\mathcal{D}|\mu, s, \nu) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)(\nu\pi)^{1/2}s} \left(1 + \frac{1}{\nu} \left(\frac{\mathcal{D} - \mu}{s}\right)^2\right)^{-\frac{\nu+1}{2}}, \quad (2.22)$$

centred at μ with hyperparameters s and ν , which represent the scale and degrees of freedom, respectively. Here, Γ denotes the gamma function. In the limits $\nu = 1$ and $\nu \rightarrow \infty$, this distribution becomes the Cauchy and Gaussian distributions, respectively.

2.4.1 Characterising the Posterior Distribution

Each free parameter (i.e., non-fixed, random variable) in the inference problem adds a dimension to the posterior probability distribution. The goal of inference is to identify the most probable parameters and quantify the associated uncertainty. These correspond to the mode of the posterior (the maximum a posteriori (MAP) estimate) and the highest density interval (HDI), which contains the most probable $X\%$ of the mass of the posterior probability distribution. HDI is typically desired for each parameter individually. To obtain it, the multidimensional posterior is marginalised by integrating over all parameters except the one of interest, resulting in a one-dimensional marginal distribution from which the HDI can be extracted.

Typically, analytic solutions for the MAP and HDI are unavailable, necessitating numerical methods to characterise the posterior distribution. Consider a posterior involving two continuous parameters, as shown in Figure 2.5.a. Computing probabilities for all combinations of parameter values would require an infinite number of evaluations of Bayes' theorem (equation (2.21)). Efficient strategies are therefore needed to sample or evaluate only the most informative regions of the posterior. Optimal strategies sufficiently characterise the distribution while minimising computational cost. The efficiency of different methods typically depends on the dimensionality of the parameter space.

2.4.1.1 Low-Dimensional Problems: Grids

The 'grid-based method' conducts analysis at sequential parameter combinations arranged over a grid, as illustrated in Figure 2.5.b. Subsequently, a spline can be fitted for a more comprehensive approximation of the posterior distribution. To ensure a suitable approximation, each dimension must have a minimum number of grid points distributed within the expected parameter range limits. Consequently, the total number of grid points increases exponentially with the number of parameters.

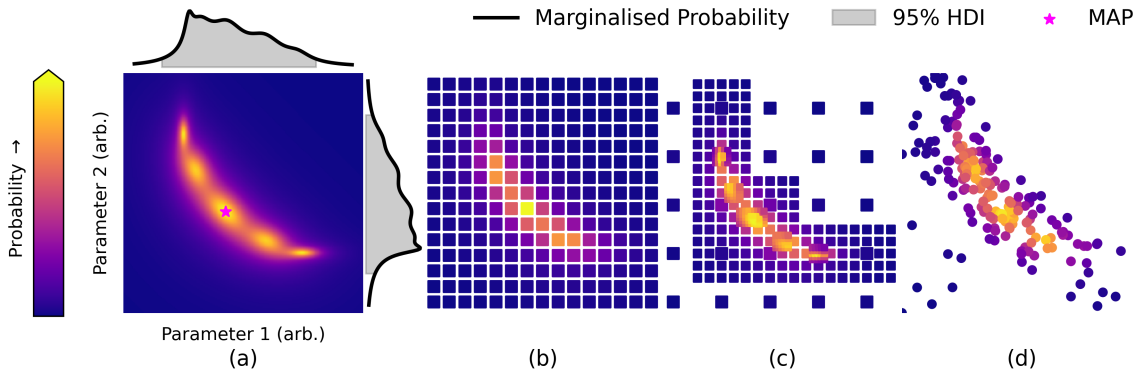


Figure 2.5: Representations of a 2D posterior distribution with colour reflecting the posterior probability.

(a) The underlying posterior distribution to be deduced. Marginalised probability distributions shown with shading reflecting 95% highest density interval (HDI).

(b) The posterior distribution found through a grid based method.

(c) The posterior distribution found through an adaptive-grid based method.

(d) The posterior distribution found with a Hamiltonian Markov Chain Monte Carlo (starting in lower left corner with ‘burn in’ included).

Moreover, as the number of parameters rises, a larger proportion of these grid points falls within low probability regions. In two dimensions, parameter ranges (each of width w) should be positioned around a central region of expected high probability, resembling a square circumscribing a circle. Thus, $\frac{1}{4}\pi w^2$ is significant out of the total w^2 . For k dimensions, the volume of the hypersphere encompassing expected high probability relative to the circumscribing hypercube is $(\pi/4)^{k/2} \left(\Gamma\left(\frac{k}{2} + 1\right)\right)^{-1}$. For $k = 6$, with just 32 grid points required per dimension, over a billion posterior evaluations are necessary, of which less than 10% are anticipated to fall within the *useful* high probability region. For $k = 8$ (and 32 grid points per dimension) this number exceeds a trillion, with fewer than 2% being useful.

It becomes highly inefficient to use grid-based methods with more than five parameters. Advanced adaptive grid designs, which dynamically adjust grid resolution, mitigate some issues related to diminishing useful space in higher dimensions [10, 58, 89]. An example of such an adaptive grid-based method is illustrated in Figure 2.5c. However, these methods still face difficulties when the number of parameters exceeds approximately eight.

2.4.1.2 High-Dimensional Problems: MCMC

To address high-dimensional inference problems, Markov Chain Monte Carlo (MCMC) methods are commonly used. In MCMC, each state represents a specific combina-

tion of parameters. New states are proposed and accepted or rejected based on a defined acceptance algorithm. Over time, the chain of accepted states can reach a pseudo-stationary regime, where the density of samples approximates the posterior distribution. This behaviour is illustrated in Figure 2.5.d.

The efficiency of a Markov Chain Monte Carlo (MCMC) algorithm heavily relies on the intelligent proposal of new states. Hamiltonian MCMC (HMC) [80] enhances efficiency by treating the posterior distribution as a *landscape* through which a *ball* traverses. The particle is assigned random momentum and evolves through the posterior via Hamiltonian dynamics. At the end of a defined simulation time, its trajectory is halted, and a Metropolis-Hastings acceptance step determines whether the proposed state is accepted [40]. This method improves sampling efficiency by proposing states that are more likely to be accepted. However, each trajectory necessitates multiple evaluations of the posterior to simulate the dynamics.

Further efficiency gains can be achieved by incorporating an estimate of the posterior covariance into the HMC algorithm through the use of a ‘mass matrix’. Within HMC, the mass matrix defines the kinetic energy term and therefore controls how efficiently parameter space is explored. HMC performs most efficiently when parameters are approximately uncorrelated and span similar characteristic scales; however, in reality, posterior distributions often exhibit strong correlations and large differences in parameter sensitivity. The mass matrix therefore acts as a preconditioner, approximately rescaling parameter space according to the local curvature structure of the posterior distribution. In directions of high posterior variance, the particle can move more rapidly, while motion is reduced along tightly constrained directions. Similarly, correlated momentum components can be introduced for strongly correlated parameters, enabling more efficient exploration of the posterior distribution.

The mass matrix may be approximated using three approaches:

- Adaptive Estimation: adaptively learned during sampling;
- MAP Curvature Approximation: estimated from the local curvature near the maximum a posteriori (MAP) solution;
- *A Priori* Construction: specified *a priori* using prior knowledge of the parameter structure.

Their determination is detailed in Appendix A.

2.4.1.3 Handling Hyperparameters

The *full Bayesian* approach treats hyperparameters as free random variables within the posterior distribution. Subsequent marginalisation of these hyperparameters ensures that all reasonable hyperparameter values are considered. Although Hamiltonian MCMC can manage high-dimensional inference problems, a large number of free parameters may still impede the accurate characterisation of the posterior distribution.

The *empirical Bayesian* approach to hyperparameters, treats them as fixed, data-informed quantities. Bayes' theorem then takes the form,

$$P(\theta|\mathcal{D}, \alpha) = \frac{P(\mathcal{D}|\theta, \alpha)P(\theta|\alpha)}{P(\mathcal{D}|\alpha)} \quad (2.23)$$

where α denotes the hyperparameters. Two common strategies are used to determine these values. In *Type II maximum likelihood*, the hyperparameters are chosen to maximise the evidence, $P(\mathcal{D}|\alpha)$. However, this approach necessitates computationally expensive marginalisation over all model parameters.

In the data-driven empirical Bayesian approach employed in this work, pre-existing simulation data is utilised to identify hyperparameter values that generalise across multiple scenarios. For instance, latent quantities that inform the structure or constraints of the prior (e.g., whether a parameter is strictly positive) are inferred from the simulation data. The hyperparameters are subsequently fitted to align with the empirical distribution of these inferred values. This ensures that the hyperparameters reflect the statistical behaviour observed across representative datasets and minimises introduced bias.

2.4.1.4 Bayesian Inference

In most cases, diagnostic systems measure quantities that are indirectly related to the parameters of interest. Consequently, a forward model, \mathcal{M} , is required to connect the parameters of interest, θ , to the measured data. In this context, the mean, μ , from equation (2.22) becomes $\mathcal{M}(\theta)$. Recognising that the evidence in equation (2.21) serves as a normalisation constant that does not affect the MAP or HDI, equation (2.21) can be re-written as,

$$\underbrace{P(\mathcal{M}(\theta)|\mathcal{D}, \alpha)}_{\text{posterior}} \propto \underbrace{P(\mathcal{D}|\mathcal{M}(\theta), \alpha)}_{\text{likelihood}} \times \underbrace{P(\theta|\alpha)}_{\text{prior}}. \quad (2.24)$$

In the context of our work, θ represents multiple parameters describing the ‘plasma state’ and \mathcal{M} represents the multiple forward models required to predict the multiple diagnostic responses. The parameters providing the modal posterior probability (MAP) is taken to be the inferred plasma state (informally referred to as the inferred parameters/inference). For each parameter, the highest density interval (HDI) containing 95% of the marginalised probability distribution is taken to be the uncertainty in the inference.

2.4.2 Bayesian Inference in Fusion Research

The forward models dictate the number of parameters required for Bayesian inference. Certain diagnostics have routine procedures to extract parameters of interest. Thus, simple forward models can be developed to relate the measured quantity to the parameter of interest. However, many diagnostics measure quantities that exhibit complex dependencies on these parameters. The value of a measured quantity may arise from multiple combinations of parameters. Consequently, the forward model of the measured quantity cannot be easily inverted. This constitutes an ‘ill-posed’ problem.

To alleviate ill-posed problems, additional diagnostics may be incorporated into the inference. However, certain diagnostics necessitate numerous parameters to achieve an accurate forward model. As discussed in Section 2.4.1, the complex MCMC method is essential when the number of parameters exceeds approximately eight. This highlights a significant divergence in the community, with some researchers opting for grid-based methods by first simplifying diagnostic data, enabling the use of forward models with a reduced number of parameters.

2.4.2.1 Combining Diagnostics

A source of ill-posed problems in data analysis is unknown factors, such as calibrations. In [33], a new Thomson Scattering system with various unknown factors was analysed using Bayesian inference. The nuisance parameters were marginalised, resulting in considerable joint uncertainties in the inferences of electron temperature and electron density. In [32], the researchers continued the Thomson Scattering analysis by including additional heterogeneous diagnostics to create an integrated data analysis system. The additional information from these diagnostics helped to constrain the joint uncertainties in the inferences of electron temperature and electron

density. Ultimately, the inclusion of additional diagnostics helped to *learn* the various calibration factors in the new diagnostic. Subsequent analysis could then be performed with fixed calibration factors, rendering the problem well-posed.

Combining diagnostics is commonly employed to infer electron temperature, electron density, and magnetic equilibrium profiles in the core region using the MINERVA framework [103, 105, 107, 108]. In an overview [84], the authors discuss the necessary assumptions for inference in the MINERVA framework, indicating that prior assumptions (included via the prior in equation (2.21)) must be made to prevent the uncertainties in the inferred parameters becoming too large. Too large uncertainties prevents the extraction of meaningful results and render the inference insignificant. We refer to this as ‘intractable inference’. As mentioned in [106], a key source of prior assumptions lies in the spatial structure of the plasma, which can transform intractable inference into tractable inference.

2.4.2.2 Utilising Spatial Information: Gaussian Processes

Gaussian process (GP) regression is a non-parametric Bayesian technique that *learns* correlations between points in the input domain and produces a distribution over possible outputs. Each input \mathbf{X} , within a continuous and potentially multidimensional domain, is associated with an output value $f(\mathbf{X})$. These outputs are modelled as random variables, so that the entire output space forms a multivariate distribution. A Gaussian process is a stochastic process in which all such random variables are jointly Gaussian distributed, yielding a Gaussian multivariate distribution over the domain.

The Gaussian processes encode correlations through a covariance, $\Sigma(\mathbf{X}, \mathbf{X}')$, defined by a kernel function relating inputs \mathbf{X} and \mathbf{X}' . The kernel function depends on hyperparameters (such as lengthscales and amplitudes). Combining the resulting covariance with a mean function, $\mu(\mathbf{X})$, a GP prior on the function values is expressed as

$$P(\mathbf{f}) = \mathcal{GP}(\mu(\mathbf{X}), \Sigma(\mathbf{X}, \mathbf{X}')). \quad (2.25)$$

Given the observed data \mathcal{D} , Bayes’ theorem updates the prior distribution to yield the posterior distribution of functions.

$$P(\mathbf{f}|\mathcal{D}) = \frac{P(\mathcal{D}|\mathbf{f})P(\mathbf{f})}{P(\mathcal{D})}. \quad (2.26)$$

GP regression has been applied in fusion research, for example, to fit noisy core pressure profiles from Thomson scattering [61]. By exploiting spatial correlations

and optimising hyperparameters, GP regression can reconstruct latent, noise-reduced pressure profiles. More generally, Gaussian processes can encode spatial information in Bayesian inference. [106] demonstrated this by applying GP priors in tomographic inversion, transforming diagnostic measurements into an unobserved latent field. The latent field $F(\theta)$, parameterised by inference parameters θ , represents the function values \mathbf{f} at inputs \mathbf{X} . The GP prior then takes the form

$$P(\theta) \propto \exp\left(-\frac{1}{2}(F(\theta) - \mu)^T \Sigma^{-1}(F(\theta) - \mu)\right), \quad (2.27)$$

providing spatial covariance information that can render otherwise intractable inference problems tractable.

2.4.3 Bayesian Inference in the Divertor Region

Despite numerous diagnostics, the divertor region poses a substantial challenge for tractable inference. This complexity arises from the multitude of species and the limited symmetry present. Consequently, numerous free parameters are needed to develop comprehensive forward models for the various diagnostics and the overall plasma state. There are many combinations of these free parameters that can explain the diagnostic measurements within each diagnostic's uncertainty. The result is an extremely ill-posed problem.

Cameras, spectrally filtered for various visible atomic transitions (such as MANTIS at TCV [86] and MWI at MAST-U [28]), oriented tangentially and radially through the divertor plasma, provide substantial information across the entire divertor region. Tomographic inversion retrieves two-dimensional poloidal cross-sections of local emissivity from different atomic transitions [133]. This enables the division of emissivity data into 'data-points' for each spatial location ('cell') within the poloidal cross-section. Consequently, the overall inference problem can be segmented into multiple independent inference problems, one for each spatial data point (cell). As only localised emission needs to be described, each independent inference problem requires fewer free parameters. Therefore, this 'cell-based' method permits the quicker, grid-based inference method. However, tomographic inversion is ill-posed, leading to additional uncertainty and spurious 'inversion artefacts'. To avoid this, an alternative to pre-inverting the data is to provide a comprehensive forward model for the non-inverted camera data. This approach, however, necessitates that free parameters for all spatial locations be jointly considered in the inference. As such, the more expensive MCMC inference method is required.

At TCV [85], JET [48], and ASDEX-U [2], the cell-based approach of pre-inverting the camera data has been adopted. According to equation (3.4), the Balmer series emissivity due to atomic recombination and electron-ion excitation can be modelled using T_e , n_e and n_D . By using line ratios between two Balmer lines as the ‘data’, unknown calibrations can be omitted. Consequently, only three free parameters are needed for inference at discrete spatial locations; these inferences can be combined to recover 2D profiles of electron temperature, electron density, and neutral atomic hydrogen density. As shown in [85], the quick, cell-based inference method produced results displayed in Figure 2.6b. [85] recognises that Thomson scattering T_e and n_e data can also be included in the inference without adding extra free parameters. This data significantly reduced the uncertainty in the inferred T_e and n_D ; however, it can only be included at locations where Thomson scattering data is available. [48] notes that line-integrated spectroscopy data of the Balmer series also uses the same parameters. However, modelling the spectroscopy data requires parameters from multiple spatial locations, making it impossible to include this data during independent inference for each spatial point.

A ‘mesh-based’ inference approach has been introduced in [9] which used synthetic data to demonstrate that inference of T_e , n_e and n_D can be performed without pre-inverting the camera data. By describing parameters at mesh vertices throughout the divertor, a comprehensive forward model can be directly constructed for the camera’s view. [9] notes that this mesh-based method allows for incorporating Langmuir probe and Thomson scattering data, both modelled in terms of T_e and n_e , into the inference. Furthermore, [9] recognised that T_e , n_e , and n_D parameters are not expected to vary significantly over short distances. Consequently, spatial smoothness can be integrated into the inference through the prior of equation (2.21). This approach helps constrain the uncertainties in the inferred parameters.

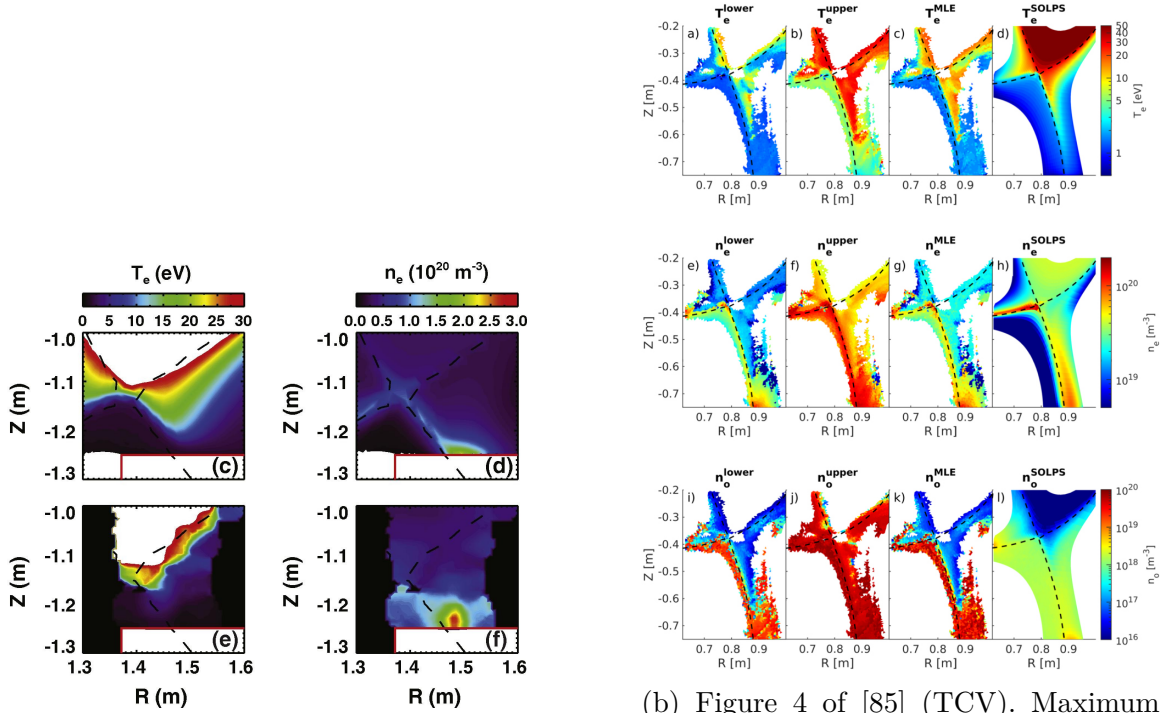
As has been reported [112], there is an additional contribution to Balmer emission from processes associated with molecular activated dissociation (MAD). Consequently, additional parameters are required to account for this molecular contribution when modelling Balmer-line emission. [47] parameterised the molecular contribution to the D_α and D_γ emissivity as a fraction of the total D_α and D_γ emissivity, respectively. Through assumptions in the collisional radiative models used for molecular contributions to Balmer emissivity—specifically the dominant molecular species, isotopic effects, and extent of vibrational excitation—the molecular density could also be recovered. The resultant 2D inferred distributions are presented in Figure 2.6c.

The need to account for molecular contributions to Balmer emission complicates tractable inference by exacerbating the divertor integrated data analysis's ill-posed nature. Large uncertainties in inferred parameters can lead to spurious profiles in mesh-based inference methods or substantial noise in pre-inversion inference methods. Although the results shown in 2.6b and 2.6c demonstrate effective characterisation of trends and absolute values, the resultant noise restricts the derivation of subsequent quantities from the inferred parameters. For example, conducted heat flux can be calculated using equation (2.18) which requires the electron temperature gradient. Gradient calculations are highly susceptible to noise, limiting the use of such techniques to derive important quantities in divertor physics. Smoothing techniques to handle noise reduce spatial resolution and therefore limit required in-depth analysis of the divertor region.

To reduce uncertainty in the inferred profiles, additional information is necessary. Alternative emission lines, such as the deuterium Fulcher band (approximately 560 to 640 nm) and helium emission lines, may provide this information. However, as noted in [47], Fulcher band emission measurements are susceptible to high and uncertain levels of Bremsstrahlung radiation. [8, 63] demonstrated that helium data can yield 2D electron temperature and electron density profiles, as shown in Figure 2.6d³. Chapter 4 investigates the impacts of including additional information in the inference.

Mesh-based Bayesian inference methods enable the incorporation of spatial information as a prior in equation (2.24). The authors of [106] introduced spatial information through Gaussian processes (see equation (2.25)), asserting that the underlying function is a sample from a Gaussian process, which provides a non-parametric approach for tomographic inversion. [59] further illustrated the application of Gaussian processes in Bayesian inference by treating the underlying pressure profiles (described by T_e and n_e parameters) as Gaussian processes and employing these parameters into the forward models of the included diagnostics. The resulting profiles varied smoothly, were consistent with the heterogeneous diagnostics, and included uncertainty information.

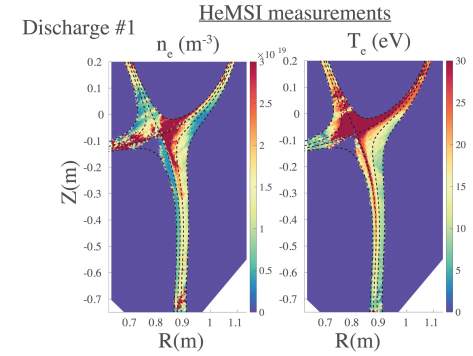
³[63] does not follow a Bayesian inference paradigm, however equation (12) of [63] is akin to a maximum likelihood estimation (MLE) which equates the MAP when no prior information is included.



(a) Figure 3(c-f) of [12] (D-IIID). T_e (e) and n_e (f) measurements through strike-point sweeping. Corresponding SOLPS interpretive model shown in (c) and (d).

(b) Figure 4 of [85] (TCV). Maximum likelihood estimate (MLE) (and \pm (upper/lower) 1σ confidence interval) of T_e , n_e , n_D (n_0) inference alongside inter-pretive SOLPS model.

(c) Figure 8 of [47] (JET). Inference of n_e (left column), n_D (n_{at}) (middle column) and molecular hydrogen density (n_{mol}) (right column) at 3 different outer strike point temperatures (rows).



(d) Figure 7 of [63] (TCV). n_e and T_e measurements using helium data.

Figure 2.6: Experimental measurements of two-dimensional electron temperature and electron density profiles as available in the literature. Please see original papers [12, 47, 63, 85], figures displayed here for visual comparison only.

2.5 Summary: Current Understanding of the Tokamak Divertor

The tokamak divertor is crucial for managing heat and particle fluxes to the plasma-facing components (PFCs). The detached divertor regime provides a mechanism to maintain these fluxes at levels tolerable for the PFCs. Understanding the detached divertor regime necessitates a comprehensive understanding of the divertor plasma. Achieving this requires high-resolution, two-dimensional profiles of electron temperature, electron density, neutral density, and molecular density.

The multitude of plasma species and processes in the divertor complicates modelling and diagnosis. However, integrated data analysis approaches have generated 2D profiles of T_e , n_e , n_D , and n_{D_2} that reveal trends and validate modelling approaches. This, alongside extensive spectroscopic analysis [114], has facilitated a solid understanding of the detached divertor state.

As outlined in Figure 2.7, substantial heat fluxes are expected to enter upstream. This heat flux, conducted through a relatively static plasma, is anticipated to peak sharply near the separatrix. It is expected to reduce significantly within the common flux region and be minimal in the private flux region. The heat flux interacts with intrinsic and/or seeded impurities and neutral hydrogen, which can be ionised and excited. This process removes power from the plasma and lowers its temperature. The lower temperature facilitates significant populations of neutral hydrogen atoms. These neutral hydrogen atoms are readily ionised by the hot plasma, providing an ion source and subsequent ion flux. Continued ionisation (and associated excitation) further removes power from the plasma, further lowering the temperature.

At temperatures around 1 eV, substantial molecular hydrogen densities can exist. These densities facilitate additional power removal and ion recombination through MAR. At temperatures below around 0.7 eV, atomic recombination processes further decrease power and ion densities in the plasma. Additionally, charge exchange between hot ions and neutral atomic and molecular hydrogen species reduces plasma momentum (pressure). Consequently, both the plasma's particle and heat fluxes are significantly reduced upon reaching the target.

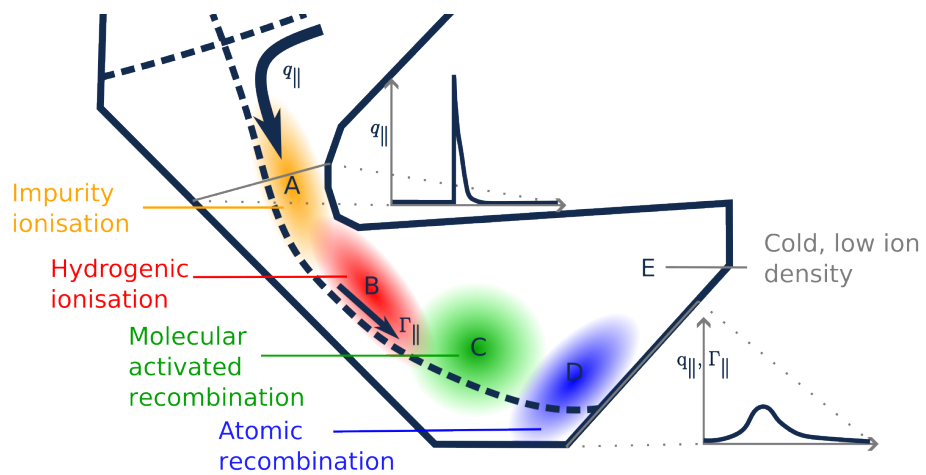


Figure 2.7: Qualitative summary of a detached divertor. Inspired by Figure 4 of [114]. Dashed line indicates the separatrix.

Chapter 3

Bayesian Inference using Balmer Emission Lines

For future power plants, passive diagnostics that require minimal space within the tokamak wall are advantageous. Filtered imaging cameras provide two-dimensional information on photon emission at specific wavelengths, making them a powerful, passive, and information-dense diagnostic tool.

The deuterium Balmer series ($n \rightarrow 2$ transitions) emits light in the visible spectrum. This series provides insights into the plasma state; a relative increase in emission from high n to low n Balmer lines indicates a cold, volumetrically recombining plasma [5, 109]. The Balmer series has proven extremely useful for divertor analysis [115] and integrated data analysis techniques [47, 49, 87]. However, various plasma conditions promote Balmer-line emission, including atomic recombination, molecular activated recombination, and ionisation (see Figure 2.1). These diverse conditions limit the information content of the Balmer line series and its ability to infer the plasma state with low uncertainty.

The extensive spatial coverage provided by filtered imaging of the Balmer series makes it a central focus of our work. In this chapter, we describe our model for Balmer line emission which includes molecular contributions. The model is a reparametrisation of the model developed for the BaSPMI method described in [112]. Both models are written as a function of key plasma state parameters: electron temperature, electron density, neutral atomic hydrogen density, and a metric related to molecular hydrogen density. These enable a comprehensive description of the divertor plasma as outlined in Section 2.3.1. We utilise this model to analyse the information contained in the Balmer lines. By generating synthetic datasets, we apply Bayesian inference under various plasma conditions to identify when accurate, low-uncertainty, and low-error inference of the plasma state can be achieved.

In this chapter, we conclude that the Balmer series alone cannot differentiate between conditions dominated by molecular activated recombination (MAR) and hydrogenic ionisation. However, additional information regarding whether the plasma state is above or below ~ 3 eV electron temperature allows for a distinction between MAR-dominant and ionisation-dominant conditions. This information enables accurate inference of electron temperature and density across a range of plasma conditions expected in a tokamak divertor.

3.1 Model for Molecular Contributions to Balmer Line Emission

The derivation of the model for molecular contributions to Balmer line emission follows the work of Dr. Verhaegh as used in [112]. It has been reparametrised in this study to accommodate inference when logarithms of parameters are considered. This is published in [36].

Hydrogenic Balmer line emission is typically modelled through two predominant atomic processes: electron-impact recombination (EIR) and electron-impact excitation (EIE). Recent work [49, 116] has shown that hydrogen molecules significantly influence both monitored hydrogen line emission and plasma dynamics through molecular activated recombination (MAR), molecular activated dissociation (MAD), and molecular activated ionisation (MAI). These processes create excited hydrogen atoms and must be included as additional contributors to Balmer line emission, equation (2.2).

Most Balmer line emission from molecular break-up in the divertor occurs when molecular ions (primarily D_2^+ , but also D_2^- and D_3^+) interact with the plasma [135], as outlined in Table 2.2. We model the Balmer line emission from plasma-molecular interactions in a manner analogous to equation (2.2). This requires the photon-emissivity coefficients¹ for each process (each a function of T_e and n_e), as well as the respective densities of the involved species. For clarity, we group the various processes by their driving molecule, resulting in: PEC_{D_2} , $\text{PEC}_{D_2^+}^{n \rightarrow 2}$, $\text{PEC}_{D_2^-}^{n \rightarrow 2}$, $\text{PEC}_{D_3^+}^{n \rightarrow 2}$. The respective PEC values for plasma-molecular interactions were obtained from the population coefficients retrieved from the YACORA-on-the-web [134, 135] collisional-radiative model. Assuming electron density and hydrogen ion density are roughly

¹Akin to those provided by ADAS for atomic processes, found through equation (2.1).

equal, $Z_{eff} \approx 1$, the molecular contribution to Balmer line emission is,

$$\varepsilon_{\text{mol.}}^{n \rightarrow 2} = n_e n_{D_2} \text{PEC}_{D_2}^{n \rightarrow 2} + n_e n_{D_2^+} \text{PEC}_{D_2^+}^{n \rightarrow 2} + n_e n_{D^-} \text{PEC}_{D^-}^{n \rightarrow 2} + n_e n_{D_3^+} \text{PEC}_{D_3^+}^{n \rightarrow 2}. \quad (3.1)$$

Under the assumption of no transport, the ratios of molecular ion densities to neutral molecular density, n_{D_2} , can be estimated by the ratio of creation to destruction rates of each molecular species [117, 120]. These rates were sourced from the default hydrogenic rates used by EIRENE [93] as functions of T_e and n_e [93]. Though these rates pertain to hydrogen, our analysis focuses on a deuterium plasma. Therefore, translation factors for hydrogen to deuterium, as adopted from [44], were applied (0.95 and 0.7 for $\frac{n_{D_2^+}}{n_{D_2}}$ and $\frac{n_{D^-}}{n_{D_2}}$ respectively), rather than employing an isotope mass-based rescaling, which increases errors in molecular charge-exchange rates at low temperatures [114].

With molecular density ratios expressed as functions of T_e and n_e , equation (3.1) can define an effective molecular photo-emissivity coefficient,

$$\text{PEC}_{D_2, \text{eff}}^{n \rightarrow 2}(T_e, n_e) = \text{PEC}_{D_2}^{n \rightarrow 2} + \frac{n_{D_2^+}}{n_{D_2}} \text{PEC}_{D_2^+}^{n \rightarrow 2} + \frac{n_{D^-}}{n_{D_2}} \text{PEC}_{D^-}^{n \rightarrow 2} + \frac{n_{D_3^+}}{n_{D_2}} \text{PEC}_{D_3^+}^{n \rightarrow 2}. \quad (3.2)$$

This encapsulates the molecular contribution to the Balmer-line emissivity, with each PEC term depending solely on n_e and T_e . This formulation eliminates the necessity of including additional parameters for each molecular process in equation (2.2).

The uncertainties in the relative ratios of molecular ion densities to D_2 in equation (3.2) are significant [114]. Molecular ion creation is driven by vibrationally excited molecules, which depend not only on local n_e and T_e , but also on molecular transport and plasma-wall interactions. However, significant uncertainties exist in the vibrational distribution model used by EIRENE [114, 117, 123]. This results in considerable uncertainties in the various creation rates of molecular ions, particularly in $D_2 + D^+ \rightarrow D_2^+ + D$, the dominant driver for D_2^+ under detached conditions [55, 114, 116, 117, 118, 123]. Consequently, there are large uncertainties in $\text{PEC}_{D_2, \text{eff}}^{n \rightarrow 2}$.

If only one emission process is dominant (typically the case at each spatial location in the plasma), the ratio between the emission coefficients of two Balmer lines, $\text{PEC}_{D_2, \text{eff}}^{m \rightarrow 2} / \text{PEC}_{D_2, \text{eff}}^{n \rightarrow 2}$, remains unaffected by uncertainties in the relative ratios of molecular ion densities. To determine $\varepsilon_{\text{mol.}}^{n \rightarrow 2}$, instead of multiplying $\text{PEC}_{D_2, \text{eff}}^{n \rightarrow 2}$ by n_e and the molecular density, we reformulate our emission model to use only the ratios of $\text{PEC}_{D_2, \text{eff}}^{n \rightarrow 2}$ between different Balmer lines. This reparameterisation employs $Q_{\text{mol.}}$, defined as the ratio of D_α 's molecular emissivity contributions, $\varepsilon_{\text{mol.}}^{3 \rightarrow 2}$, to atomic emissivity contributions, $\varepsilon_{\text{atm.}}^{3 \rightarrow 2}$.

$$Q_{D_\alpha}^{\text{mol.}} \equiv \frac{\varepsilon_{\text{mol.}}^{3 \rightarrow 2}}{\varepsilon_{\text{atm.}}^{3 \rightarrow 2}} = \frac{\varepsilon_{D_2, \text{eff}}^{3 \rightarrow 2}}{\varepsilon_{\text{rec}}^{3 \rightarrow 2} + \varepsilon_{\text{exc}}^{3 \rightarrow 2}}. \quad (3.3)$$

The emissivity of the Balmer lines $n \rightarrow 2$ is thus given by,

$$\varepsilon_{n \rightarrow 2} = \underbrace{n_e^2 \text{PEC}_{rec}^{n \rightarrow 2} + n_H n_e \text{PEC}_{exc}^{n \rightarrow 2}}_{\text{Atomic contribution, } \varepsilon_{\text{atm}}^{n \rightarrow 2}} + \underbrace{Q_{D_\alpha}^{\text{mol.}} \left(n_e^2 \text{PEC}_{rec}^{3 \rightarrow 2} + n_H n_e \text{PEC}_{exc}^{3 \rightarrow 2} \right)}_{\text{Molecular contribution, } \varepsilon_{\text{mol.}}^{n \rightarrow 2}} \frac{\text{PEC}_{D2,\text{eff}}^{n \rightarrow 2}}{\text{PEC}_{D2,\text{eff}}^{3 \rightarrow 2}}. \quad (3.4)$$

Consequently, no explicit molecular density is required. Furthermore, the effective molecular PEC is only represented as a ratio between lines $n \rightarrow 2$ and $3 \rightarrow 2$, minimising the impact of the aforementioned approximations. The molecular density can subsequently be recovered through,

$$n_{D_2} = Q_{D_\alpha}^{\text{mol.}} \left(n_e \text{PEC}_{rec}^{3 \rightarrow 2}(T_e, n_e) + n_H \text{PEC}_{exc}^{3 \rightarrow 2}(T_e, n_e) \right) \left(\text{PEC}_{D2,\text{eff}}^{3 \rightarrow 2}(T_e, n_e) \right)^{-1}. \quad (3.5)$$

However, this molecular density does incorporate uncertainties from the relative ratios of molecular ion densities in $\text{PEC}_{D2,\text{eff}}$.

3.1.1 Fulcher Band Model

Neutral hydrogen molecules excited in the plasma emit visible light between 595 and 615 nm. This Fulcher band emission provides supplementary information to Balmer line emission [116]. Spectroscopy systems measure Fulcher band brightness for multiple chords distributed in the poloidal plane, intersecting at various locations along the separatrix (see Figure 2.b. of [116]). Assuming limited changes in electron density, chords below 50% of the peak Fulcher brightness typically exhibited electron temperatures close to the separatrix, either below approximately 3 eV or above approximately 15 eV. In a tokamak divertor, the electron temperature is expected to decrease monotonically along the separatrix toward the target. Consequently, the region between the target and the point at which 50% of the Fulcher peak is measured corresponds to temperatures $\lesssim 3$ eV (Figure 11 of [116]).

Fulcher band emission can be modelled using the same four parameters as those required for the Balmer line emission model with molecular contributions (equation (3.4)). The emissivity, shown in Figure 3.1, is defined by

$$\varepsilon_{\text{Fulcher}} = n_{D_2} n_e \text{PEC}_{\text{Fulcher}}(T_e, n_e). \quad (3.6)$$

The photon emissivity coefficient for the Fulcher band, $\text{PEC}_{\text{Fulcher}}$, is provided by the AMJUEL collisional-radiative model (reaction 2.2.5fl [93]). n_{D_2} is parameterised by T_e , n_e , n_D , and $Q_{\text{mol.}}$ through equation (3.5).

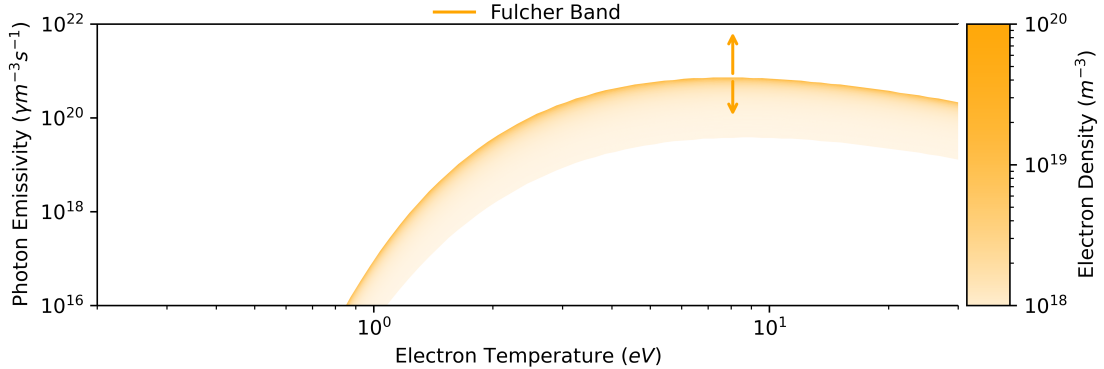


Figure 3.1: Photon emissivity of the Fulcher band. Emissivities at $T_e n_{D_2} = 10^{19} \text{ eV m}^{-3}$. Upward/downward arrows show impact on the peak due to a ten times increase/decrease in $T_e n_{D_2}$.

3.2 Balmer Line Ratios

In Figure 3.2, equation (3.4) is used to determine the emissivity of Balmer line photons at various temperatures and densities. Emission peaks correspond to the dominance of atomic recombination ($\lesssim 0.7 \text{ eV}$), molecular activated recombination ($\sim 1 \text{ eV}$) and ionisation ($\gtrsim 5 \text{ eV}$) processes (see Figure 2.1). The emissivity peak for molecular activated recombination scales approximately linearly with molecular density (solid arrow). The peak for ionisation scales approximately linearly with atomic neutral density (dashed arrow)². Furthermore, all emissivities scale with electron density (shading intensity). Consequently, various combinations of electron temperature, electron density, atomic neutral density, and molecular density can produce the same photon emissivity.

Despite exhibiting similar trends, the relative heights of the emissivity peaks shown in Figure 3.2 vary subtly for each Balmer line. This variation is due to the different processes outlined in Table 2.2 that dominate under different conditions, favouring the excitation of electrons to different energy levels. For example, Figure 3.2 shows the ratio of D_δ to D_α is far greater at lower electron temperatures. Additionally, at electron temperatures around 1 eV, an increase in molecular density enhances D_α emission more strongly than that of higher n Balmer lines.

The relative changes in emission between two Balmer lines provide substantial insight into the plasma state. By observing ‘line ratios’ between different Balmer lines, as shown in Figure 3.3, we can discern the strong dependencies of these ratios

²Due to the occurrence of MAI, this peak also has molecular density dependence.

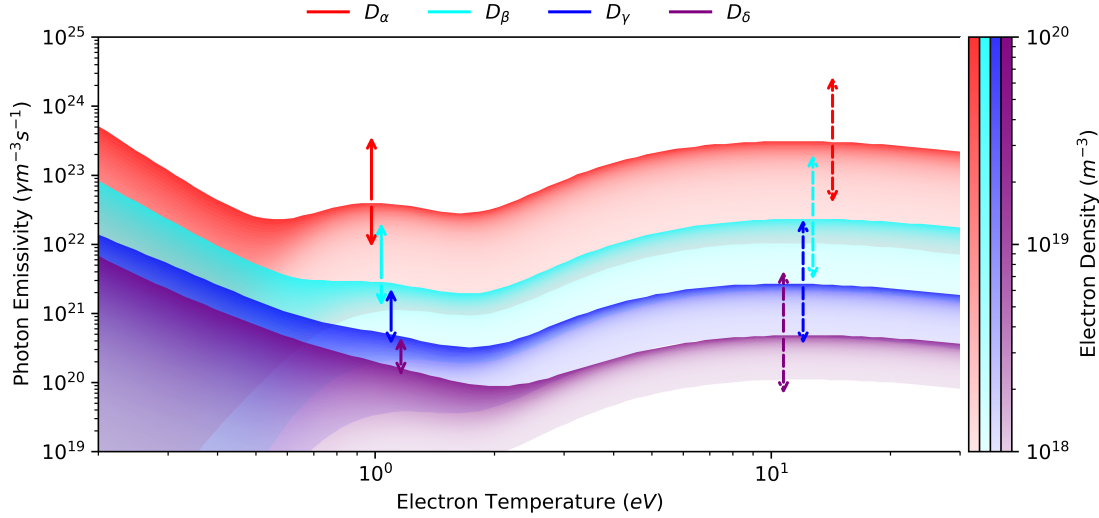


Figure 3.2: Photon emissivities of Balmer lines $3 \rightarrow 2$ (D_α) to $6 \rightarrow 2$ (D_δ) at $T_e n_D = 10^{20} \text{ eV m}^{-3}$ and $T_e n_{D_2} = 10^{19} \text{ eV m}^{-3}$. Upward/downward arrows show impact on the peaks due to a ten times increase/decrease in $T_e n_D$ (dashed arrow) and $T_e n_{D_2}$ (solid arrow).

on electron temperature and density. Due to uncertainties in the emission model, a significant difference in the line ratio is necessary to accurately infer the plasma state. In Figure 3.3a, the ratio of a low n (D_α) to a high n (D_δ) Balmer line differentiates between low electron temperatures $\lesssim 0.7 \text{ eV}$ (at $D_\alpha/D_\delta \lesssim 160$) and high electron temperatures (at $D_\alpha/D_\delta \gtrsim 160$). Above $\sim 5 \text{ eV}$, the ratio is relatively insensitive to electron temperature (meaning that it cannot be accurately inferred); however, there is a clear electron density dependence.

A key issue arises when the line ratio remains approximately constant despite significant variations in electron density and temperature. Consider the green band (contour 320 to 400) of Figure 3.3a. The line ratio is consistent for densities $\sim 10^{19} \text{ m}^{-3}$ at temperatures exceeding 2 eV , and for temperatures \sim of 2 eV at densities greater than $\sim 10^{19} \text{ m}^{-3}$. This similarity results from comparable energy level populations due to molecular activated recombination and electron-impact excitation processes. Consequently, inference using the line ratio cannot distinguish between lower temperatures with higher electron densities and higher temperatures with lower densities.

The Balmer line ratios are influenced by molecular density. Figure 3.3b presents higher molecular densities compared to Figure 3.3a. Focusing on the green band (contour 320 to 400), we observe that an increase in molecular density modifies the range of electron temperature and density combinations that produce the same line

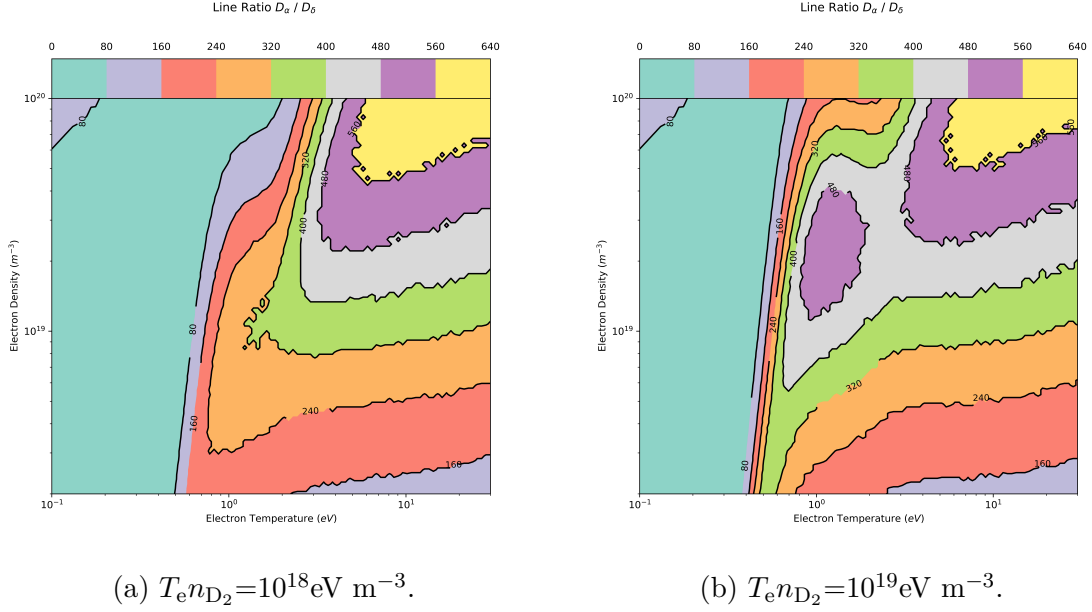


Figure 3.3: Balmer emission ratio of $3 \rightarrow 2$ (D_α) to $6 \rightarrow 2$ (D_δ) with different molecular densities ($T_e n_D = 10^{19} \text{ eV m}^{-3}$).

ratio. This further complicates the ability to distinguish between plasma states using Balmer line ratios.

There are similar line ratios across various plasma conditions. The line ratio between two Balmer lines is insufficient for accurately deducing the plasma state. Consequently, an integrated data analysis approach that considers more than just the line ratio between two Balmer lines is required.

3.3 Bayesian Inference with Balmer Lines

The parameters $\theta = \{T_e, n_e, n_D, Q_{\text{mol}}\}$ can model the emissivity of all Balmer lines using equation (3.4). Consequently, Bayesian inference with these free parameters can be performed with data for multiple Balmer lines. Assuming a Gaussian likelihood distribution and omitting prior arguments, equation (2.21) can be expressed as,

$$P(\theta|\mathcal{D}) \propto \prod_n \exp\left(-\frac{(\varepsilon_{n \rightarrow 2}(\theta) - \mathcal{D}_n)^2}{2\sigma_n^2}\right). \quad (3.7)$$

Here n is the n -th Balmer line, \mathcal{D}_n is the corresponding data, and σ_n is the uncertainty associated with the n -th Balmer line. Owing to uncertainties in the Balmer emission model, the uncertainty is assumed to be 10%, such that $\sigma_n = 0.1\mathcal{D}_n$. The resultant posterior distribution is four-dimensional and includes a maximum value (*maximum a posteriori*, MAP), representing the most plausible combination of parameters based on

the measured data. By marginalisation, the four-dimensional probability distribution can be condensed to a probability distribution over a single parameter of interest.

To assess the performance of Bayesian inference with multiple Balmer lines, three characteristic plasma states were examined. These states relate to regions D, C, and B in Figure 2.7, which represent atomic recombination, molecular activated recombination, and hydrogenic ionisation-dominant conditions, respectively³. For each state, the corresponding emission for each Balmer line was calculated using equation (3.1) evaluated at the ‘true’ parameters θ_{True} for the respective states. To account for measurement error, 10% Gaussian noise was applied to the emission, and three samples were drawn to demonstrate variability in the inference.

In Figure 3.4, Bayesian inference was conducted using data from the three regions of interest, with the marginalised temperature probability distributions presented. Inference utilised the ‘grid-based’ method (see Section 2.4.1.1). The upper row (a-c) considered only $n = 3$ and $n = 6$ (D_α and D_δ) data. The lower row (d-f) considered Balmer lines $n \in \{3, 4, 5, 6\}$ data.

Figure 3.4 highlight three critical figures of merit for an inference to be considered *good*:

1. Inference error: $|\theta_{\text{MAP}} - \theta_{\text{True}}|$.

The difference between the true parameter value and the parameter value corresponding to the maximum posterior probability (MAP). For non-Gaussian probability distributions, θ_{MAP} does not necessarily correspond to the mode of the marginalised probability distribution. The percentage error is taken to be $(|\theta_{\text{MAP}} - \theta_{\text{True}}| / \theta_{\text{True}}) \times 100\%$.

2. Inference uncertainty: $\theta_{\text{HDI}}^{\text{upr}} - \theta_{\text{HDI}}^{\text{lwr}}$.

The width of the region encompassing the 95% highest density interval (for the marginalised probability distribution). The percentage uncertainty is taken to be $\left(\left(\theta_{\text{HDI}}^{\text{upr}} - \theta_{\text{HDI}}^{\text{lwr}}\right) / \theta_{\text{True}}\right) \times 100\%$. For multimodal distributions, the total width (summed for each HDI region) is used.

3. Inference accuracy:

Whether the true parameter value is within the 95% highest density interval or not.

³Representative parameters for characteristic plasma states $\theta_X = \{T_e eV, n_e m^{-3}, n_D m^{-3}, Q_{mol}\}_X$:
 $\theta_B = \{10, 3 \times 10^{19}, 10^{18}, 0.25\}$; $\theta_C = \{1, 3 \times 10^{19}, 10^{19}, 100\}$; $\theta_D = \{0.3, 4 \times 10^{19}, 10^{19}, 0.05\}$

We define an *effective* inference as one that is accurate, with low inference error ($\lesssim 50\%$) and low uncertainty ($\lesssim 100\%$).

Figure 3.4(a-c) demonstrates that accurate inference can be achieved for the electron temperature under atomic recombination and molecular activated recombination conditions, but not under ionisation conditions. Extending the inference to four Balmer lines (Figure 3.4(d-f)) reduces the uncertainty. However, inference with four Balmer lines remains inaccurate for ionisation conditions.

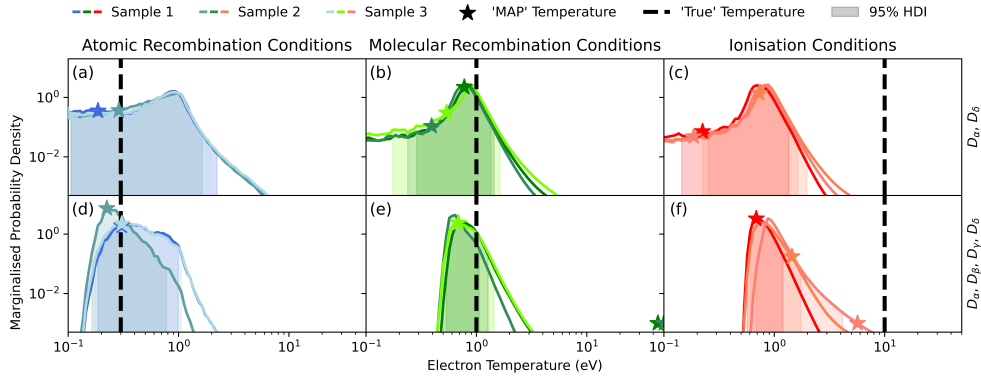


Figure 3.4: Inferred marginalised electron temperature probability distributions through Bayesian inference with different Balmer line combinations.

Upper row (a-c): D_α , D_δ

Lower row (d-f): D_α , D_β , D_γ , D_δ

Samples refer to multiple trials of the inference, each with different synthetic data owing to the expected 10% uncertainty on measurement data.

We observed in Figure 3.3 that both molecular recombination conditions ($T_e \sim 1-3$ eV, vertical bands) and ionisation conditions ($T_e \gtrsim 3$ eV, horizontal bands) gave similar line ratios. To find the marginalised electron temperature probability distribution, an integral is performed over the n_e , n_D and Q_{mol} axes. In Figure 3.3, the narrow vertical bands around 1 eV cover a significantly larger range of electron density than the horizontal bands with the equivalent line ratio. Consequently, when marginalising to the electron density axis, there is a spurious increase in the probability around 1eV (MAR associated emission) in comparison to the higher temperature found in ionisation conditions (aligned to electron-impact excitation, EIE). The large scale of the electron density parameter is the reason for the systematic offset between the true and inferred temperature in ionisation conditions.

3.3.1 Balmer Line Ability to Infer the Whole Plasma State

The electron temperature, electron density, neutral atomic hydrogen density, and molecular hydrogen density provide a comprehensive description of the tokamak divertor, as outlined in Section 2.3.1. In this section, we perform synthetic tests to evaluate whether these quantities, encoded in the parameters of the Balmer emission model in equation (3.4), can be accurately inferred from the Balmer lines, with low error and uncertainty.

To consolidate the diverse anticipated plasma states, five representative plasma conditions were selected using a self-organising map⁴. These conditions correspond to regions A, B, C, D, and E shown in Figure 2.7 and are defined by:

A High temperature, low hydrogenic ionisation conditions: $T_e > 4\text{eV}$, $R_{Hyd. Ion.} < 10^{23}\text{m}^{-3}$

B Ionisation dominant conditions: $T_e > 4\text{eV}$, $R_{Hyd. Ion.} > 10^{23}\text{m}^{-3}\text{s}^{-1}$

C Molecular activated recombination dominant conditions: $0.7\text{eV} < T_e < 3\text{eV}$, $\varepsilon_{D_\alpha} > 5 \times 10^{21}\gamma\text{m}^{-3}\text{s}^{-1}$

D Atomic recombination dominant conditions: $T_e < 0.7\text{eV}$, $\varepsilon_{D_\alpha} > 5 \times 10^{19}\text{m}^{-3}$

E Cold, low emission conditions: $T_e < 5\text{eV}$, $\varepsilon_{D_\alpha} < 5 \times 10^{21}\gamma\text{m}^{-3}\text{s}^{-1}$

Here $R_{Hyd. Ion.}$ is the rate of hydrogenic ionisation.

The analysis in Section 3.3 was repeated across many synthetic data sets, segmented into five key plasma conditions. Figure 3.5a illustrates a pictogram highlighting the inference capabilities within these regions for each parameter. These pictograms, where large green circles with a solid ring indicate a strong inference capability, are used to discuss the expected inference capability in each of the key plasma conditions.

3.3.1.1 Electron Temperature

The synthetic tests revealed that, when molecular contributions to the emission are considered, the Balmer line series can effectively infer electron temperature only at low temperatures. Under EIR-dominant conditions (region D of Figure 2.7), the electron temperature was effectively inferred, resulting in a median percentage error

⁴The inference performance varied considerably depending where in T_e -, n_{e^-} -, n_D - and Q_{mol} -space the data was generated. A self-organising map reduced the dimensionality of this data to its key components.

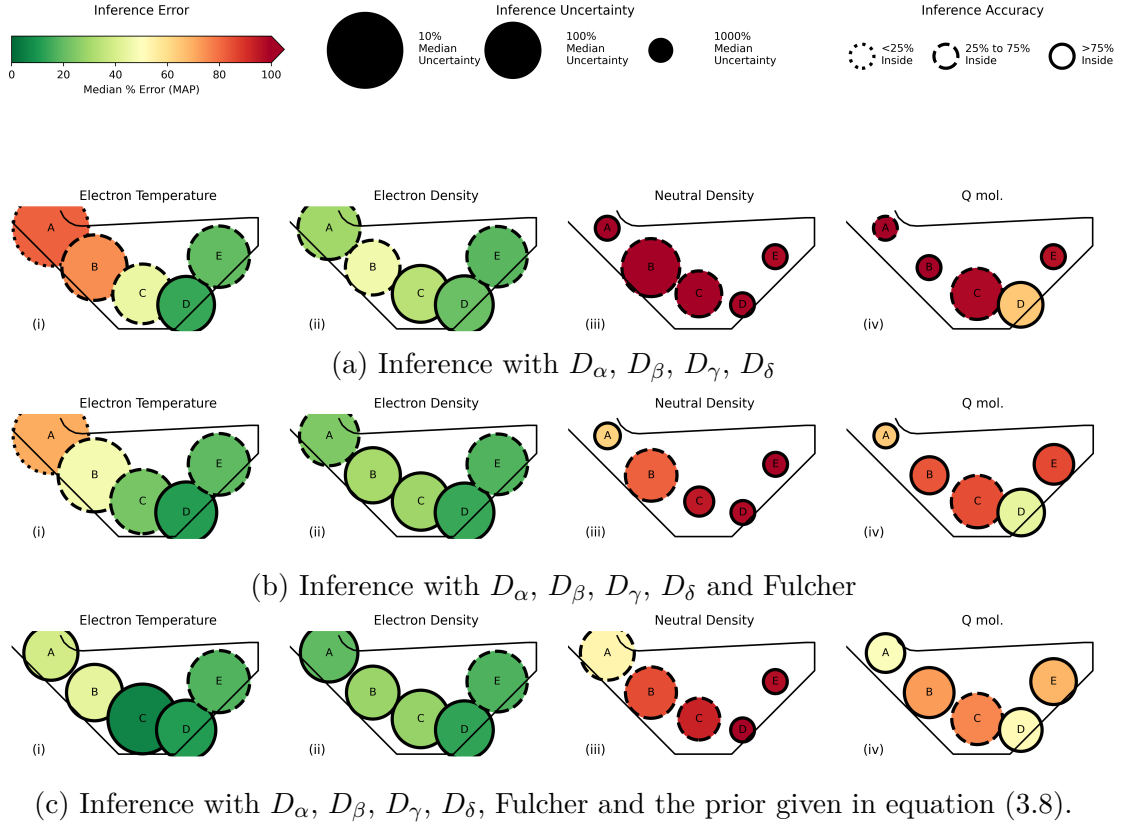


Figure 3.5: Qualitative pictogram of typical inference capability (error, accuracy and uncertainty) for key plasma conditions based on differing information included in the inference.

Large, green circles with solid rings indicate effective inference. Error, uncertainty and % inside correspond to inference error, inference uncertainty and inference accuracy respectively as defined in Section 3.3. Plasma conditions A, B, C, D and E roughly correspond to regions in Figure 2.7 and are defined in the prose. Specific quantities represented in the pictogram available on request.

of less than 25%. However, at the higher temperatures found in regions B and C, the similar trends in Balmer line emissions from MAR and EIE processes made the electron temperature too challenging to infer. This was in agreement with the results shown in Figure 3.4.

Region A of Figure 2.7 corresponds to high temperatures and low neutral deuterium densities. This results in low hydrogenic ionisation conditions and minimal Balmer line emission. However, the absence of Balmer line emission does not indicate a deficiency in inference capability; rather, the absence is itself valuable information.

Both region A (hot, low neutral density) and region E (cold, low electron density) conditions result in similar low levels of Balmer emission. As a result, the distinct con-

ditions cannot be differentiated. The inferences generally indicate low temperatures instead of high temperatures, with low uncertainty and, consequently, low accuracy for the electron temperature in region A. This is attributed to marginalisation over parameters across different scales, as with the results shown in Figure 3.4.

3.3.1.2 Electron Density

As shown in Figure 3.5a.ii., the electron density, n_e , was inferred with low error across all regions, with a median percentage error below 40%. However, accuracy remained below 75% across regions A, B, and E. This suggests that, for these regions, the true uncertainty in electron density is greater than the estimates obtained through Bayesian inference.

3.3.1.3 Atomic Neutral Density

Inference of atomic neutral density, using the Balmer lines showed high error and substantial uncertainty across all divertor regions (Figure 3.5a.iii.). n_D only significantly influences electron-ion excitation emission and weakly affects emission from plasma-molecule interactions. Consequently, in regions with low excitation emission, there is limited information on atomic neutral density. Even in regions B and C, where n_D affects the measured emission, the median percentage error still exceeded 100%. The proportion of emission from excitation and MAR cannot be determined by Balmer lines alone in these regions, resulting in the observed high n_D inferred error.

3.3.1.4 Q_{mol} .

The ratio of D_α emission from molecular to atomic processes, Q_{mol} , is crucial for determining the neutral molecular deuterium density via equation (3.3). Figure 3.5a.iv shows that Q_{mol} cannot be reliably inferred with reasonable error or uncertainty across all conditions. As a result, n_{D_2} also cannot be effectively determined. This significantly limits the comprehensive description of the divertor plasma state based solely on the Balmer lines.

3.4 Discussion

3.4.1 Inclusion of Additional Information

The results of the synthetic tests using the Balmer lines indicate that additional information is needed to enable comprehensive inference of the divertor plasma state.

Typically, EIR-dominant conditions can be distinguished from other conditions. However, the similar Balmer line emission signatures from both EIE and MAR processes render inference ill-posed under these conditions, making it ineffective.

Additional information can be provided through the Fulcher band, which can be modelled through the same parameters used to model the Balmer lines. After incorporating Fulcher band emission into equation (3.7), the analysis of Figure 3.5a was repeated. The results, shown in Figure 3.5b, indicate that including Fulcher data enhances inference effectiveness across all parameters and regions (except in region D, where atomic recombination dominates and its impact is minimal). This improvement is primarily due to its ability to differentiate between emission from MAR and electron-impact excitation processes. However, substantial uncertainties in the Fulcher band model, largely due to the explicit use of n_{D_2} rather than Q_{mol} , resulted in a 20% uncertainty being associated with Fulcher band data in equation (3.7). This limits the influence of the Fulcher band information.

To help separate Balmer emission from EIE and MAR processes, we consider the availability of information that explicitly specifies whether the electron temperature is above or below 3 eV. This information is incorporated into equation (3.7) by including the prior term of equation (2.21),

$$P(\theta) = \exp\left(-\frac{1}{2}\left(\frac{f(T_e)}{\sigma}\right)^2\right). \quad (3.8)$$

Here, $f(T_e) = \max\left(\frac{T_e}{3} - 1, 0\right)$ when the true temperature is known to be below 3 eV and $f(T_e) = \max\left(\frac{3}{T_e} - 1, 0\right)$ when the true temperature is known to be above 3 eV. The scale parameter σ controls the probability *penalty* applied to parameters that violate the prior (σ is here taken to be 0.01).

The analysis of Figure 3.5b was repeated with the prior of equation (3.8) included. The results, shown in Figure 3.5c, significantly improve inference accuracy of T_e and n_e . The definition of Q_{mol} (equation (3.3)) indicates that the substantial improvement in Q_{mol} inference suggests that the 3 eV prior information significantly aids in distinguishing between MAR and EIE Balmer emission. Although such an explicit prior is unavailable for inferences with experimental data, it highlights the importance of information regarding whether the electron temperature is above or below approximately 3 eV.

Helium line emission data has also been used to infer T_e and n_e in the divertor, and is considered to provide useful information at electron temperatures exceeding 10

eV[64]. However, including this data necessitates additional free parameters (helium densities) to model the Helium data.

3.4.2 Inferred Ion Source and Sink Rates

Figure 3.5a highlights the challenges of inferring individual plasma parameters T_e , n_e , n_D , and Q_{mol} from Balmer line emission. However, these parameters can be used to calculate the net ion source rate, $\Delta S = S_+ - S_-$, determined by considering the reaction rates of various processes (Figure 2.1) and their impact on the change in the number of ions (Table 2.2). The net ion source rate provides substantial insight into the divertor plasma state.

Inference of the net ion source rate proved more reliable than inferring the individual plasma parameters (T_e , n_e , n_D , and Q_{mol}). Inference was conducted using Balmer lines D_α , D_β , D_γ , D_δ , and Fulcher band information. Across different divertor plasma conditions (regions A-E), the median percentage error in the inferred net ion source rate, ΔS , was 25%. Consequently, when only Balmer line and Fulcher band information is available, the inference of the ion source rate is possible, even if the inference of n_D , Q_{mol} and T_e above 4 eV cannot be trusted.

Despite this, the region of dominant MAR, region C, exhibited an increased median percentage error in ΔS (65%). This was due to inference of many data sets in this region suggesting the conditions gave a net ion source when in reality there was a net ion sink (and vice versa). This issue arose again from the challenges in distinguishing between MAR-dominant (ion sink) and ionisation-dominant (ion source) conditions.

3.5 Summary

A model for Balmer line emission, with molecular contributions, has been developed as a function of the parameters $\theta = \{T_e, n_e, n_D, \text{ and } Q_{\text{mol}}\}$. The definition of Q_{mol} facilitated forward models of Balmer line emission with low uncertainty (approximately 10%). However, substantial uncertainties remain when deriving the neutral molecular deuterium density, n_{D_2} , from these parameters.

From the Balmer line emission model, the ratios between two Balmer lines vary significantly across different regions of the parameter space. This variation facilitated Bayesian inference using the Balmer lines to determine electron temperature and density in cold (<1 eV), EIR-dominant regions, achieving low inference error ($\lesssim 15\%$), low inference uncertainty, and high inference accuracy.

Similar ratios between two Balmer lines were observed across two parameter space regions. These corresponded to dominant-MAR (a narrow T_e region around 1 eV covering a wide range of n_e) and dominant-ionisation (a broad T_e region greater than 3 eV but a narrow range of n_e). Consequently, Bayesian inference with the Balmer lines struggled to differentiate between these regions. This resulted in low inference accuracy at electron temperatures above 1 eV.

Fulcher band emission can also be modelled using the parameters T_e , n_e , n_{D_2} , and Q_{mol} . However, the model's uncertainty is significant (assumed to be around 20 %). This arises from its direct dependence on n_{D_2} and the expectation that, at low Fulcher band emissivities, experimental data will be contaminated by a relatively high level of background emission. Including Fulcher band information improved the error and accuracy of the inference by facilitating the separation between dominant-MAR and dominant-ionisation conditions. Nevertheless, the uncertainties associated with the Fulcher model limited its overall impact. An alternative method to ascertain whether the temperature is $\gtrsim 3$ eV is found to be highly desirable.

3.5.1 Implications for Using Balmer Lines in Divertor Data Analysis

Alternative work [47, 85] has utilised Balmer lines to infer conditions in the divertor region, as illustrated in Figures 2.6b⁵ and 2.6c. We find that electron densities can be effectively inferred throughout the divertor using this method. However, in divertor plasmas with significant molecular densities, inferences of neutral density and electron temperature are challenged under conditions where $T_e \gtrsim 3$ eV. Despite this, the corresponding net ion source rate was typically well inferred, with a median absolute percentage error of 25%.

The availability of information indicating whether the temperature is $\gtrsim 3$ eV is limited in the divertor. [116] used the relative spatial position of the peak in the Fulcher band emission to provide this information. This distinction helped differentiate MAR from EIE emission. Such ‘spatially dependent’ information cannot be trivially inserted into the cell-based analysis performed in [47, 85] since their work treats each spatial location as an independent inference problem (as has been performed in this chapter). The mesh-based method introduced in [9] jointly considered all spatial locations in a single inference problem⁶. This allowed the inclusion of spa-

⁵Note: molecular contributions were not included in [85].

⁶Note: molecular contributions were not included in [9].

tially dependent information in the synthetic study. Whether this spatially dependent information is sufficient to distinguish MAR from EIE is investigated in Chapter 4.

Chapter 4

Synthetic Validation of Integrated Data Analysis in the Tokamak Divertor

Chapter 3 detailed how Balmer lines alone are insufficient to provide low uncertainty inferences of T_e , n_e and n_D across the entire divertor region. This was largely due to the inability to reliably distinguish between electron-impact excitation (EIE) and molecular activated recombination (MAR) processes. To improve the inference accuracy, it was found that additional information was required, such as whether the plasma state was above or below 3 eV.

Additional information can be provided to the Bayesian inference (see equation (2.24)) through additional diagnostics in the likelihood or through better use of the prior. As shown in [9], an ‘inference mesh’ can be used to specify vertices at which plasma quantities (e.g. T_e) are described as free parameters. This comprehensive description of the plasma state allows multiple diagnostics, observing different locations in the plasma, and ‘spatially-dependent’ information to be included in the inference. For example, localised temperature measurements provided by the Thomson scattering diagnostic (as well as n_D and Langmuir probe measurements) were included in the study of [9] alongside filtered camera imaging diagnostics of the Balmer lines. By demanding spatial smoothness of the two-dimensional T_e profile (spatially dependent information), the local T_e information provided by the Thomson scattering could be extended to a wider region.

In this chapter we extend the work of [9] by incorporating molecular contributions into the Balmer emission model and implementing them within the Divertor Multi-Instrument Bayesian Analysis System (D-MIBAS). The mesh-based framework exploits spatial information through a field-aligned inference mesh, enabling direct

modelling of camera data and contrasting with cell-based approaches that treat each location independently [47, 85]. Using MCMC sampling to characterise posterior distributions, we perform synthetic verification of D-MIBAS on two SOLPS-ITER simulations of the MAST-U Super-X divertor (attached and detached). This allows us to examine the consequences of neglecting molecular processes, to quantify the benefits of spatially informed inference, and to assess the value of additional diagnostics such as helium spectroscopy.

This chapter is comprised of the work of [36]:

D Greenhouse et al. *Two-dimensional inference of divertor plasma characteristics: advancements to a multi-instrument Bayesian analysis system*. 2025. DOI 10.1088/1361-6587/adab1b .

4.1 The D-MIBAS Method

D-MIBAS is designed to be flexible to multiple diagnostics and any divertor scenario. It requires a comprehensive description of the plasma state. This description is provided by a mesh in the poloidal cross-section such that, through linear interpolation between mesh vertices and the assumption of toroidal symmetry, the entire plasma state can be described. In this section, we layout the contributions to the posterior distribution (diagnostics and prior arguments) used in the synthetic test of this chapter. Our approach to characterising the posterior distribution is described in Section 4.1.6.

4.1.1 A Field-Aligned Inference Mesh

As with [9], we seek a mesh-based inference. Much of the physics in divertor plasmas is tied to magnetic flux surfaces and consequently is spatially-dependent. Bayesian inference permits *a priori* information to our inference in the form of the prior probability of equation (2.24). Consequently, to enable an efficient manner for including prior information on anisotropic plasma properties, we align our inference mesh to surfaces of constant poloidal magnetic flux.

The inference mesh used in this work is shown in Figure 4.1 where V vertices form a two-dimensional grid of normalised poloidal magnetic flux coordinates, ψ , and distances to the target (Tile 5) parallel to surfaces of constant poloidal magnetic flux. With knowledge of the magnetic pitch angle at each vertex, the magnetic connection length to the target (distance to the target parallel to a flux tube) at each mesh vertex can be trivially recovered.




Figure 4.1: The inference mesh and the diagnostic coverage in the MAST-U divertor. The magnetic geometry is obtained from SOLPS-ITER simulations of the MAST-U Super-X divertor.

4.1.2 The Posterior Distribution

The logarithm of the posterior distribution of equation (2.24), $\mathcal{L}_{\text{Posterior}}$, is the sum of the log-likelihood, $\mathcal{L}_{\text{Likelihood}}$, log-prior, $\mathcal{L}_{\text{Prior}}$ and a normalisation constant, c , which is neglected due to its lack of influence on either the MAP estimate or posterior distribution width (uncertainty). Both the log-likelihood and log-prior are themselves the sum of numerous log-probability distributions, respectively pertaining to different diagnostics, i , and different physics-informed priors, j . The total log-posterior probability is therefore,

$$\underbrace{\mathcal{L}(\theta|\mathcal{D})}_{\text{log-posterior}} = c + \underbrace{\sum_i \mathcal{L}_i(\mathcal{D}|\theta)}_{\text{log-likelihood}} + \underbrace{\sum_j \mathcal{L}_j(\theta)}_{\text{log-prior}}. \quad (4.1)$$

4.1.3 Likelihoods

For each diagnostic, a multivariate Gaussian likelihood distribution (uncorrelated between diagnostics) was assumed such that each diagnostic log-likelihood (minus an omitted normalisation constant) was given by

$$\mathcal{L}_i(\mathcal{D}|\theta) = -\frac{1}{2} (\mathbf{d}_i - \mathcal{M}_i(\theta))^T \Sigma_i^{-1} (\mathbf{d}_i - \mathcal{M}_i(\theta)). \quad (4.2)$$

\mathbf{d}_i specifies a vector of measurements for the i^{th} diagnostic. $\mathcal{M}_i(\theta)$ represents the forward model that predicts a diagnostic response from a plasma state (represented

by the model parameters, θ). Σ is the covariance matrix, which contains information on uncertainties.

For this study, the diagnostic set consisted of:

- divertor Thomson scattering (TS) measurements of T_e and n_e ;
- divertor Langmuir probe (LP) measurements of ion saturation current density, j_{sat} , at the target (tile 5);
- camera images ¹ of the brightness of
 - Balmer lines $\{3(\alpha), \dots, 7(\epsilon)\} \rightarrow 2$,
 - Helium lines $\{502\text{nm}, 668\text{nm}, 728\text{nm}\}$.

The model for channel j of each diagnostic, $\mathcal{M}_{i,j}(\theta)$ from equation (4.2), was given by

$$\mathcal{M}_{i,j}(\theta) = \sum_k G_{i,j,k} q_{i,k}. \quad (4.3)$$

Here, G_i represents the geometry matrix, which encodes the contribution of the relevant quantity, q , evaluated at mesh vertices k (see Figure 4.1), to the j 'th diagnostic channel.

Thomson scattering measurements of T_e and n_e , along with their uncertainties, are routinely available at MAST-U [41]. However, data points may be unavailable or not perfectly aligned to specific flux tubes, which we reflect in our chosen positions shown in Figure 4.1. Toroidal symmetry is assumed such that it is considered T_e and n_e Thomson scattering measurements are available at discrete poloidal locations. Therefore, for equation (4.3), $q = T_e$ and $q = n_e$ (respectively for electron temperature and electron density measurements) and G is the barycentric interpolation matrix.

Langmuir probe j_{sat} measurements are also routinely available at discrete locations along MAST-U's tile 5 (Figure 4.1). Assuming similar ion and electron temperatures in the target vicinity, the relevant quantity, q , of equation (4.3) is

$$j_{\text{sat.}} = \frac{en_e}{2} \sqrt{\frac{2eT_e}{m_i}} \quad (4.4)$$

where e is the elementary charge and m_i is the ion mass (taken exclusively as deuterium ions). As with the Thomson scattering data, toroidal symmetry is assumed such that G of equation (4.3) is the barycentric interpolation matrix.

¹Modelled after MAST-U's Multi-Wavelength-Imaging (MWI) diagnostic [29, 133] and TCV's MANTIS diagnostic [86]

Brightness images (calibrated both spatially and in absolute brightness) for wavelength windows associated with specific atomic transitions are also routinely available. For the brightness image corresponding to the atomic transition $m \rightarrow n$, $q = \varepsilon^{m \rightarrow n}$ in equation (4.3). As discussed in Section 4.1.3.1, the emissivity, $\varepsilon^{m \rightarrow n}$, is determined by the electron temperature, electron density, and the relevant neutral and ion species densities. Since brightness measurements are line-integrated along the camera lines of sight rather than local measurements in the poloidal plane, the geometry matrix, G , in equation (4.3) was constructed using a ray-tracing algorithm to determine the contribution of each mesh element to each camera pixel.

4.1.3.1 Emission Models

The emissivity of a species' atomic transition $m \rightarrow n$ (at a spatial location) can be modelled via:

$$\varepsilon^{m \rightarrow n} = \sum_i n_{i,0} n_{i,1} \text{PEC}_i^{m \rightarrow n}(T_e, n_e), \quad (4.5)$$

where i is a contributing process to the $m \rightarrow n$ emission. $n_{i,0}$ and $n_{i,1}$ are the number densities of the two species involved in the process. The T_e and n_e dependent photon-emissivity coefficient, PEC, of the process is generated from generalised collisional-radiative models (provided by ADAS [90] for atomic interactions).

To include molecular contributions to Balmer line emission, the molecular emission model developed in Section 3.1 was used. This described $\varepsilon^{m \rightarrow 2}$ as a function of T_e, n_e, n_D and Q_{mol} .

MAST-U's MWI and TCV's MANTIS cameras also routinely capture Helium-I emission lines which has been shown [63] to provide information on n_e and T_e . To avoid inaccuracies in Helium camera forward models, we explicitly include parameters for the densities of neutral and singly charged helium (n_{He^0} and $n_{\text{He}^{1+}}$ respectively) at each mesh vertex (at a cost of 2V additional free parameters) in our inference. This allowed both EIE and EIR contributions to be modelled via equation (4.5). Following the work of [8], 668, 728 and 502 nm singlet He I lines were chosen due to their insensitivity to transport of metastable states. Additional details and potential complexities are outlined in Appendix C.

4.1.4 Priors

We refer to each physics-based constraint considered in the inference as a prior. Each prior is imposed as a contribution, $\mathcal{L}_j(\theta)$, to the overall log-prior of equation (4.1)

with details of the probability function used provided in Appendix D.1. Each prior is evaluated at all mesh vertices relevant to that prior.

4.1.4.1 Spatially-Independent and -Dependent Priors

We consider ‘spatially-independent priors’ to be those that, when being evaluated at a certain mesh vertex, require no information about parameter values at a different mesh vertex. Consequently, these can be included in both the cell-based and mesh-based inference paradigms. They include priors that consider the interplay between parameters at a single mesh vertex such as an upper bound of the static electron pressure as well as bounds on parameter values. All spatially-independent priors used in this work are detailed in Table D.1 of Appendix D.2.

‘Spatially-dependent priors’, when being evaluated at a certain mesh vertex, require information of parameter values at other mesh vertices and so cannot be used in the cell-based inference paradigm. An example of such a prior is the expected monotonic reduction (toward the target) in the heat flux parallel to a surface of constant poloidal magnetic flux. All spatially-dependent priors used in this work are detailed in Table D.2 of Appendix D.2.

4.1.5 Inference Parameters

In the mesh-based inference, the model parameters,

$$\theta = \{F_1^{(1)}, \dots, F_1^{(V)}, \dots, F_N^{(1)}, \dots, F_N^{(V)}\}, \quad (4.6)$$

specify the value of each of N fields, F , at each of V mesh vertices. To compute forward models for all diagnostics, we require the parameter field set

$$F = \{T_e, n_e, n_H, Q_{\text{mol.}}, n_{\text{He}^0}, n_{\text{He}^{1+}}\}. \quad (4.7)$$

T_e , n_e and n_D are of direct interest for divertor physics. These, as well as $Q_{\text{mol.}}$, can be used to derive further quantities of interest to divertor physics including: static electron pressure; heat flux parallel to flux tubes; divertor ionisation sources; ion sinks (volumetric recombination due to electron-ion recombination and molecular activated recombination); hydrogenic power losses; and more [120]. n_{He^0} and $n_{\text{He}^{1+}}$ parameters can be identified as nuisance parameters which are necessary to evaluate forward models but not of direct interest in this work.

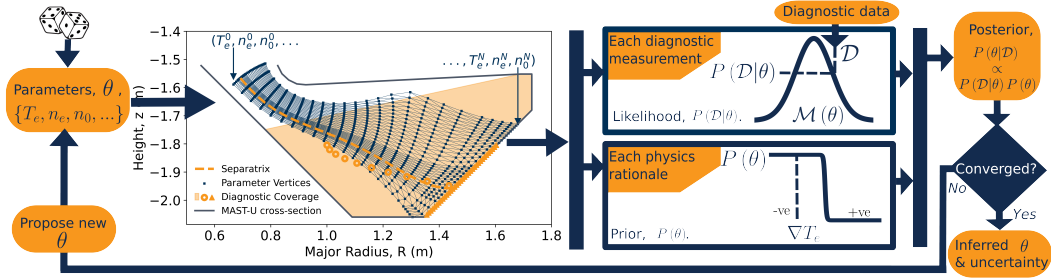


Figure 4.2: Activity diagram for characterising the posterior distribution in Bayesian inference.

4.1.6 Characterising the Posterior Distribution

The method for finding the mode of the posterior distribution and the associated uncertainty is outlined in Figure 4.2. Each iteration, an evaluation of the posterior function, has a computational cost. Consequently, the proposals for each new θ must be made in an efficient manner. The large number of inferred parameters ($NV > 5000$) requires Markov Chain Monte Carlo (MCMC) to characterise the posterior distribution and an efficient implementation of the MCMC algorithm is essential for a tractable accurate characterisation of the posterior distribution.

The complex, high-dimensional, highly correlated, multi-modal posterior distribution made MCMC challenging. To ease the difficulty, a generational optimisation algorithm, as described in the Appendix of [9], was used to find the *maximum a posteriori* (MAP). We found that the hybrid use of the first-order adam [50] and the second-order L-BFGS [68] optimisers performed best in the generational algorithm due to their strengths in different regions of the posterior distribution. Once in the MAP vicinity, Hamiltonian MCMC [80] was performed to collect samples describing the uncertainty.

Two factors were found to be essential for tractable use of Hamiltonian MCMC (HMC). Firstly, *kinks* in the probability distribution, where the probability altered dramatically within a short region of parameter space, required short steps to be taken within the HMC algorithm (to prevent divergence). Consequently, exorbitant times were required to adequately characterise the posterior distribution². To counteract this, the likelihood and prior probability functions had to be carefully constructed and transformed to remove kinks in the probability distribution. This is revisited in Section 5.3.2. Secondly, the ‘mass matrix’ used in the HMC algorithm had to be

²Each HMC iteration needed short ‘leap-frog’ time steps, but many leap-frog steps to avoid inefficient, random-walk like behaviour.

approximated. This was to contend with highly correlated parameters and parameters with substantially different uncertainty ³.

As discussed in Section 2.4.1.2, the mass matrix ideally approximates the inverse of the local covariance of the probability distribution. An approach to finding the covariance is to approximate the MAP vicinity as a multivariate Gaussian distribution. This MAP curvature approximation approach is detailed in Appendix A.2. This was found to perform well in the synthetic study in this chapter, but was adapted when contending with experimental data. This is revisited in Section 5.4.

The challenge, time, and complication of efficiently using HMC represent the key drawback of mesh-based IDA compared to cell-based IDA (which commonly does not require HMC). However, in this chapter, we seek to demonstrate that HMC can be successfully used and, with it, the benefits of using mesh-based IDA can be realised.

4.2 Synthetic Data Setup

The new integrated data analysis methodology developed in this work was applied to synthetic data. The data were generated from synthetic diagnostics applied to two SOLPS-ITER simulations: a ‘detached case’, shown in Figure 4.3aI, where the plasma was in a detached state in which PMI was expected to play a significant role; and an ‘attached case’, shown in Figure 4.3bII, where the plasma was in an attached state in which PMI was not expected to play a significant role. These simulations had an input power of 2.5 MW and were obtained from [31] using the ‘Super-X low α ’ geometry ⁴. A trace amount of helium was added to get realistic relative 2D spatial distributions of n_{He^0} and its ratio to $n_{\text{He}^{1+}}$, although the absolute numbers for n_{He^0} may not be characteristic of that obtained in the tokamak ⁵.

For each case, the SOLPS-ITER simulation was taken as the ‘true’ plasma state from which the relevant fields were extracted. To compute $Q_{\text{mol.}}$, $\text{PEC}_{\text{D}_2, \text{eff}}$ was used to post-process the D_2 density into the densities for the various molecular ions. It should be noted that this leads to significantly higher D_2^+/D_2 ratios than reported by SOLPS-ITER directly, which is known to under-predict the D_2^+/D_2 ratio due to inaccuracies in the molecular charge exchange rate (see [112, 114, 123]).

³The mass matrix specifies the mass of the ‘ball’ which undergoes Hamiltonian dynamics in the probability space. Adaptations to the mass alter the momentum with which the ball travels in each dimension.

⁴This geometry is observed more easily by the synthetic imaging diagnostic as the Super-X has no curvature in the separatrix near the target, see Figure 2 of [31].

⁵However, the absolute magnitude of n_{He^0} is not expected to impact our inference.




Figure 4.3: Comparison of ‘true’ fields used to generate synthetic data (a) and the inferred MAP estimate (b).

(a) Plasma fields for two synthetic cases from SOLPS-ITER.

(I) The detached case where the T_e (a.I.i.) drops below 5eV before it reaches the target.

(II) The attached case in which the n_e (a.II.ii.) peaks at the target.

(b) Maximum a posteriori (MAP) estimate inference of scenarios displayed in Figure 4.3(a) using posterior i. MAP was found using the generational algorithm outlined in Appendix A of [9].

To create the synthetic data sets, for each case, the forward models of equation (4.2) were used to predict diagnostic response and Gaussian errors akin to those

expected in experiments were added. As a comparison metric, for each field, the median absolute percentage error (MdAPE) between the ‘true’ plasma state and the inferred MAP estimate at each mesh vertex was found.

4.3 Results

We have applied the D-MIBAS framework to the two synthetic data sets outlined in Section 4.2 (with both including PMI in the measured Balmer-line brightness). Four different posterior distributions have been constructed under the framework to ascertain the influence of the advanced emission model, the use of spatially-dependent priors, and the inclusion of additional diagnostics. These posterior distributions are:

i Likelihood: TS, LP, helium and hydrogen Balmer multi-wavelength imaging diagnostics with a forward model that includes PMI contributions to Balmer emission.

Prior: spatially-dependent and spatially-independent priors.

This is the most complete version of the D-MIBAS.

ii As with i, but with a forward model that doesn’t include PMI contributions to Balmer emission ($Q_{\text{mol.}}=0$).

iii As with i, but without spatially-dependent priors.

iv As with i, but without helium multi-wavelength imaging.

4.3.1 Synthetic Diagnostic Test Results of the Full D-MIBAS analysis

As shown in Figure 4.3, the most complete version of the D-MIBAS framework (posterior i) was able to correctly infer the general spatial distribution of the electron temperature and electron density in both the detached case and the attached case. The atomic neutral density was less successfully inferred; however, the most prominent discrepancies occurred in regions of low Balmer and He I emission (deep into the common flux region, away from the separatrix). The inference achieved a MdAPE of 6%, 5% and 33% for T_e , n_e and n_D respectively for the detached case and 6%, 3% and 9% for the attached case.

The inferred uncertainties in the parameters (their 95% HDI) were found to be largely in agreement with the error between the true parameter values and the inferred

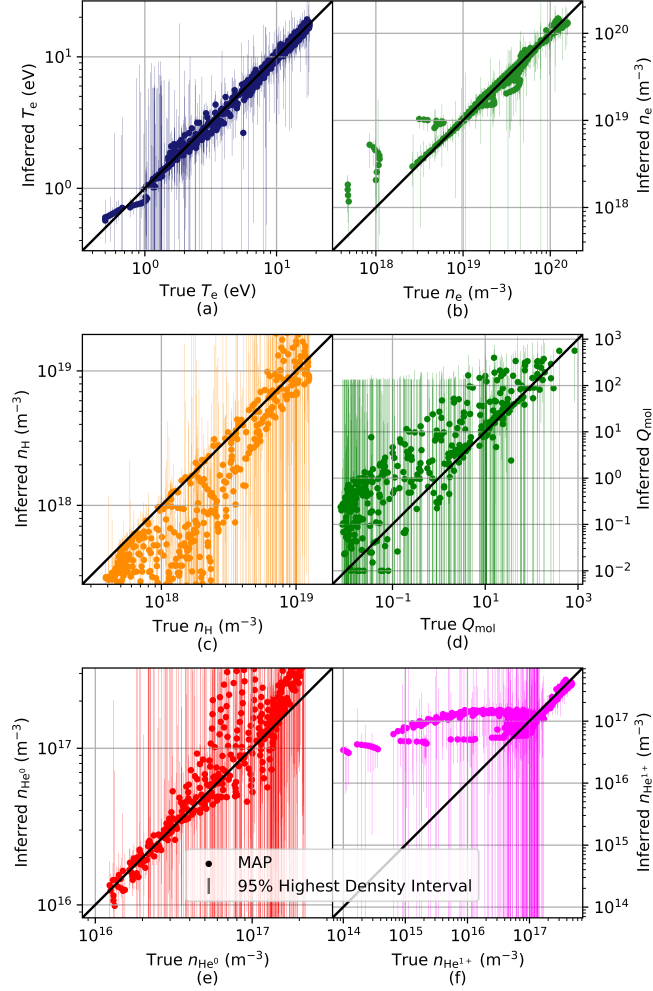


Figure 4.4: Comparison of true and inferred (MAP) parameter values using posterior i for the detached case. The 95% highest density interval (HDI) for the inferred parameters are displayed. The percentage of true data points lying within the 95% HDI are 95%, 88%, 63%, 91%, 95%, 63% for the T_e , n_e , n_D , Q_{mol} , n_{He^0} and $n_{\text{He}^{1+}}$ respectively.

MAP estimates. This is demonstrated in Figure 4.4. The uncertainty was found to be substantially greater for n_D and Q_{mol} parameters, (as well as the nuisance parameters n_{He^0} and $n_{\text{He}^{1+}}$) compared to T_e and n_e parameters. T_e and n_e parameters with large uncertainty were found to be in regions out of view of the synthetic camera diagnostic (upstream) and in regions of low Balmer and He I emission (deep into both the private flux and common flux regions).




Figure 4.5: The synthetic camera data for Balmer lines D_α ($3 \rightarrow 2$, red) and D_δ ($6 \rightarrow 2$, lilac) and their forward model predictions based on the MAP estimate of the four posterior distributions outlined in Section 4.3.

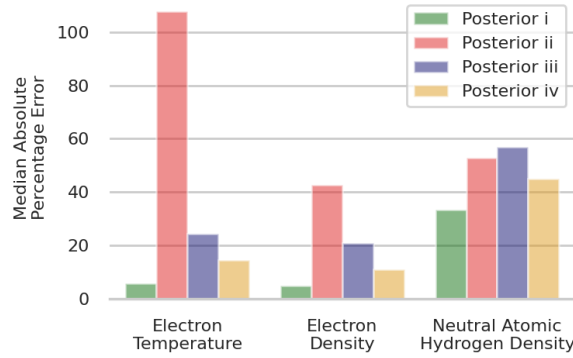


Figure 4.6: Performance of posteriors i-iv across fields of interest for the detached case.

4.3.2 Performance of the Different Posterior Distributions

Inference with each posterior distribution was able to accurately recreate the synthetic data, as highlighted in Figure 4.5. However, the inferred parameters (MAP estimate) of each posterior distribution were considerably different. The resultant MdAPE scores of inference with each of the posterior distributions outlined in Section 4.3 are shown in Figure 4.6. All parameters of interest were inferred most accurately by the most complete version of D-MIBAS (posterior i). We will now discuss each reduced analysis case (posteriors ii-iv) in more detail.

4.3.2.1 Importance of Including PMI in the Balmer Emission Model

Posterior ii only considers atomic contributions to Balmer line emission which results in a sustained overestimation of the electron temperature outside of the region where electron-ion recombination emission is dominant as can be seen in Figure 4.7a. The

cause of this was that the Balmer line emission is increased by the presence of plasma-molecule interaction. Because the Balmer line ratios from PMI and EIE are similar and significantly different from EIR, these brighter images could not be explained by EIR. Therefore, the model omitting PMI was compelled to attribute these bright images to EIE which, in turn, demanded higher electron temperatures, leading to the overestimate. This overestimation is an important demonstration that an incorrect model in Bayesian inference does not necessitate the true solution to fall within the found uncertainty interval if an invalid emission model is applied. This highlights the importance of using a model that accommodates PMI when inferring divertor plasmas with substantial molecular hydrogen densities.

4.3.2.2 Importance of Spatially-Dependent Priors

The spatially-dependent priors are permitted through jointly conducting inference with parameters at each mesh vertex as allowed by the mesh-based inference paradigm. Figure 4.7b demonstrates that the omission of these in posterior iii substantially increased uncertainty on the electron density (which was also observed for other inferred parameters) and the complete loss of information out of view of the synthetic camera diagnostic (as expected). The lack of additional information on trends in static electron pressure and smoothness leaves an ill-posed problem in which many different combinations of parameters can accurately recreate the multi-wavelength imaging data, thus causing large uncertainty in the inference. This highlights the benefit of jointly considering many parameters at different spatial positions in the inference as permitted by the mesh-based D-MIBAS.

4.3.2.3 Importance of Helium Emission

Posterior iv reflects situations where helium data are unavailable or unreliable. The helium singlet lines are known to carry information on both the electron temperature and electron density. The exclusion of helium data degraded the MdAPE for T_e and n_e to 14% and 11% respectively, as shown in Figure 4.6.

Along the separatrix, the inferred MAP estimate of posterior iv was in good agreement with the true plasma values; however, the inclusion of helium data in posterior i reduced the uncertainty in the inference of T_e and n_e (Figure 4.6). As shown in Figure 4.7c, the inclusion of helium data was found to significantly improve the n_D inference along the separatrix despite the n_D parameter not appearing in the helium forward models. This was because the reduced uncertainty in the inference of T_e and n_e managed to better constrain the process leading to Balmer-line emission

(EIR/PMI/EIE). By having a good understanding of how much EIE emission (which is n_D dependent) occurs, the n_D inference improved. This serves to demonstrate how the inclusion of additional diagnostics in the IDA can improve the accuracy of the inference of certain parameters even if they themselves are not involved in the forward model of the additional diagnostic.

4.4 Discussion

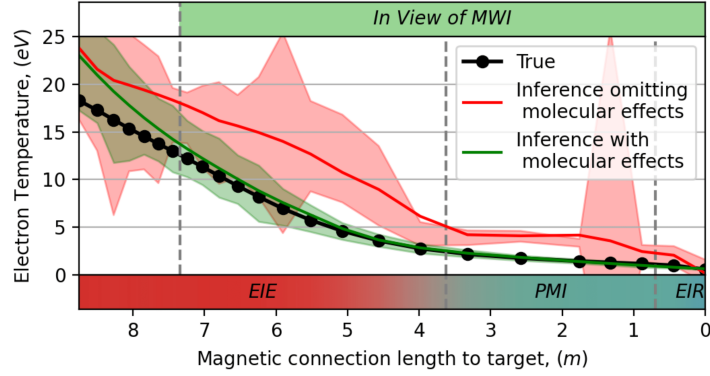

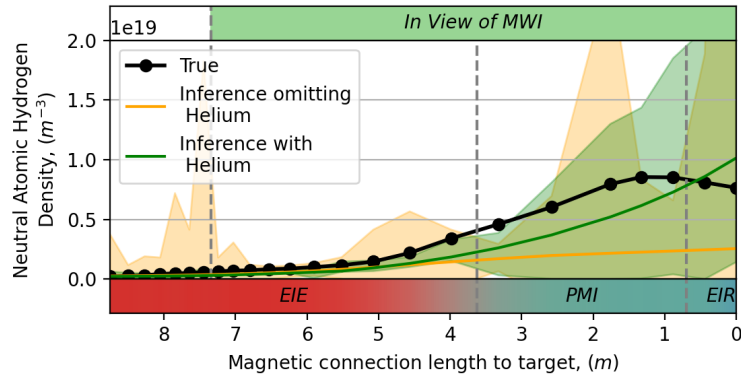
In this section, we will discuss the limitations of the current analysis and possible improvements for the future with an emphasis on the preparation of applying D-MIBAS to experimental data. We will go on to discuss possible applications of this analysis and what the implications are for such techniques.

4.4.1 Limitations of the Current D-MIBAS and Possible Improvements

4.4.1.1 Consideration for Experimental Data Application: Forward Model Inaccuracies

As observed in Section 4.3.2.1, inaccurate forward models can lead to MAP estimates with an inferred uncertainty range that is inconsistent with the ‘true’, known result despite achieving close agreement to the data (seen in Figure 4.5). There are uncertainties in the diagnostic descriptions that will cause inaccurate forward models. Such uncertainties are not observed when performing synthetic tests (since the forward models used for inference are the same as those used to generate synthetic data). Keeping the possibility of such uncertainties in mind is particularly important for experimental data application.

The diagnostic description of imaging and spectroscopic data includes multiple uncertainties. Multi-wavelength imaging employs bandpass filters to monitor a small, spectrally integrated part of the emission spectra. Plasma background emission and neighbouring spectral transitions can contaminate the monitored signal, resulting in overestimated camera brightnesses [87]. Reflections can displace the monitored emission which, if present, would require a more complex diagnostic description [49]. Similarly, photon opacity [70, 91], if present, would require a substantial modification to the diagnostic description. Additionally, there are uncertainties in the spatial and absolute calibrations of imaging and spectroscopy systems. The various photon emission coefficients used in the emission models can also have significant uncertainties of 10-20% [117].

(a) T_e comparison of posterior i (green) and posterior ii (red).(b) n_e comparison of posterior i (green) and posterior iii (blue).

(c) Neutral atom density comparison of posterior i (green) and posterior iv (gold).

Figure 4.7: Comparison of inferred parameters with the ‘true’ values (Figure 4.3Ia) for the detached case. The parameters follow a surface of constant poloidal flux, $\psi = 1$ (the separatrix). Solid lines indicate the inferred (MAP) parameters and the shaded region indicate the 95% HDI.

The additional uncertainties outlined above require that, when applied to experimental data, the forward model uncertainties (held in Σ of equation (4.2)) will exceed

those from the diagnostic measurements alone. In turn, the uncertainties on inferred parameters will increase from those observed in Figure 4.7. To limit the impact of the increase in forward model uncertainties, there are several methods that can be used:

- Physics and diagnostic studies to ascertain the most appropriate forward models to use, for example using an emission model that supports PMI (as introduced in Section 3.1).
- Introduction of additional model (nuisance) parameters describing uncertainties, for instance for the absolute calibration.
- Monte Carlo approaches can be used to perturb uncertain quantities within reasonable intervals (e.g., perturb spatial camera calibration uncertainties [85]) to create multiple posterior distributions that can be combined.
- Investigating the extent of forward model uncertainties via synthetic testing with intentional mismatches between the forward model used to generate synthetic data and the model used for inference (as performed in Section 4.3.2.1).

D-MIBAS provides a framework for comparing the consistency between the data obtained by the different diagnostics. This makes integrated data analysis particularly suitable for investigating errors in diagnostic descriptions (forward models). This can be strengthened by including additional diagnostics in the D-MIBAS inference (Section 4.4.1.2). Awareness of inadequate diagnostic descriptions can be used to improve instrument characterisation [85, 86] (e.g., improved calibrations) and to inform where more comprehensive models are required (e.g., modelling directly to the camera image data instead of pre-inverted data [9]).

4.4.1.2 Improving D-MIBAS Analysis Through Additional Diagnostics

The inclusion of additional diagnostics has multiple benefits. They can: reduce parameter inference uncertainty and improve inference accuracy; mitigate model uncertainties; expose deficiencies in diagnostic models; and act to validate diagnostic interpretation. However, care must be taken that the inclusion of an additional diagnostic does not result in too large an increase in the number of free parameters required. Examples of additional diagnostics that can be included in the current D-MIBAS setup without requiring additional free parameters include:

- Line-of-sight spectroscopy in the UV-Visible and VUV regime can monitor different hydrogen and helium emission lines which can be modelled using the current parameter set.
- Bolometry (including higher resolution imaging bolometry [25]) can provide additional information on the total radiation. This provides an upper-limit to the permitted hydrogenic radiation (of which can be modelled from the current parameter set).
- IR cameras provide target heat flux profiles. This provides an upper-limit on the heat flux calculated to be flowing from the core plasma which can be modelled from the current parameter set as described in Section 4.4.1.3.

4.4.1.3 Derived Quantities

Further derived quantities can be obtained from electron and neutral parameter information. The multi-dimensional posterior distribution provides correlated samples which can be used to compute derived quantities, including their uncertainty. The inferred parameter set can be used to obtain: static electron pressure [5]; heat and ion fluxes parallel to flux tubes [117]; estimates of ion sources and sinks (both due to atomic (EIE, EIR) as well as molecular processes (MAR, MAI)) [120] and (hydrogenic) power loss estimates [118, 120].

In Figure 4.8, the derived parallel heat flux and static electron pressure are shown at a roughly constant magnetic connection length to the target in the vicinity of the detached case's detachment front. We find that the low error in our inferred plasma state enables these derived quantities to be found with a reasonably low error. Their associated uncertainties are considerable because of the compounding of uncertainties from multiple parameters. However the results of this synthetic test suggest that these quantities, of substantial interest to divertor physics research, can be accessed through D-MIBAS.

Figure 4.9 displays the EIR, EIE and PMI emission fractions for $D\gamma$ for the detached case. In general, D-MIBAS is able to separate EIR, EIE and PMI emission contributions. However, there is a tendency to overestimate the PMI contribution and underestimate the EIE contribution. This is in agreement with our findings of Chapter 3 where the EIE and PMI emission sources could not be easily distinguished.

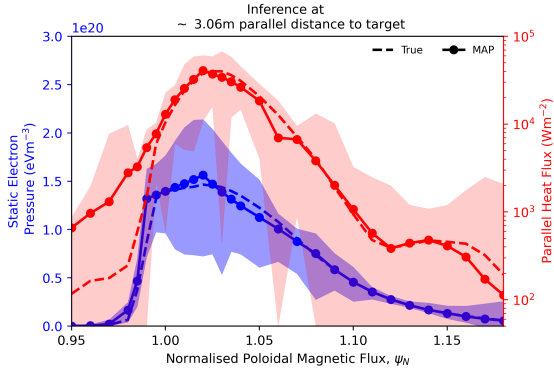


Figure 4.8: Comparison of true and inferred (using posterior i) static electron pressure and parallel heat flux values for the detached case. Evaluated for parameters perpendicular to surfaces of constant poloidal magnetic flux at a magnetic connection length of around 3 metres to the target




Figure 4.9: Comparison of true (a) and inferred (b) (using posterior i) dominant emission regions for the D_γ Balmer line for the detached case.

4.4.1.4 Performance of HMC

The mesh-based inference paradigm utilised by D-MIBAS requires Hamiltonian Markov Chain Monte Carlo (HMC) for characterisation of the posterior distribution. This technique, as outlined in Section 4.1.6, requires careful handling. It is possible for the HMC algorithm to not explore the posterior distribution adequately and therefore underestimate the inferred uncertainty. This is not the case for the cell-based inference paradigm which does not necessitate HMC.

Figure 4.10a shows the inference error between the inferred and true plasma state for the attached case using posterior i. Figure 4.10b displays the inferred uncertainty which we find aligns well with regions of increased error. This brings trust that the HMC algorithm is able to handle the complexity of our posterior distribution.

Testing found that the provision of a non-diagonal mass matrix (see Section 4.1.6) significantly aided the HMC algorithm. However, the MAP curvature approximation of the mass matrix (Appendix A.2) proved challenging to implement and was not globally representative of the full posterior distribution. Furthermore, it was found that prior arguments had to be implemented very carefully. Consider the prior related to the monotonicity of the conducted heat flux (equation (2.18)) along a flux surface. Small changes to the local electron temperature at a mesh vertex can cause large changes in the conducted heat flux. This can easily breach the prior argument so that any change in the local electron temperature has a very small probability. This

can cause the HMC algorithm to stick. More sophisticated methods for approaching such issues are outlined in Section 2.4.2.2.

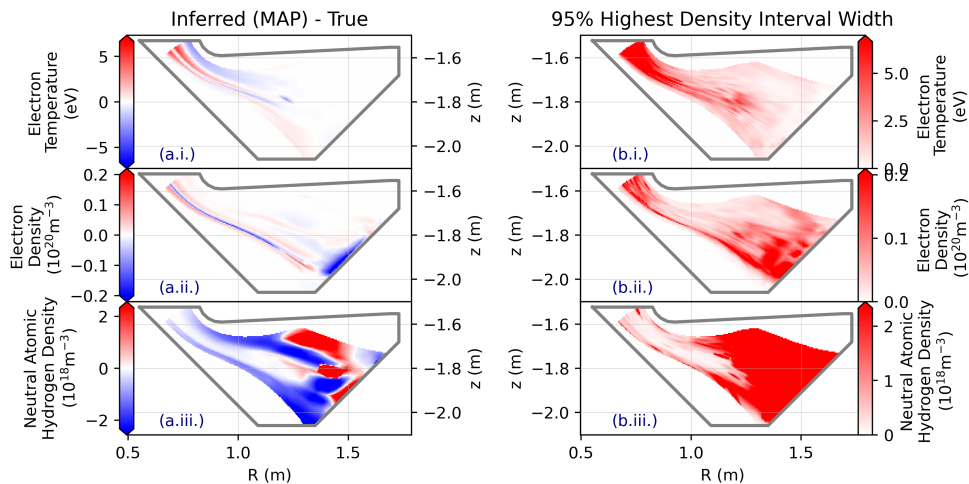


Figure 4.10: Visualisation of error (a) and uncertainty (b) for the detached case using posterior (i).

4.5 Summary and Conclusions

This study outlines a comprehensive approach for utilising multiple diagnostics to infer two-dimensional fields of plasma parameters, including electron temperature (T_e), electron density (n_e), and neutral density (n_D), across the divertor. The developed D-MIBAS incorporates the crucial role of plasma-molecule interactions (PMI) in the forward models of Balmer line emission. Through a mesh-based Bayesian inference paradigm, D-MIBAS utilises forward models of Langmuir probes, a divertor Thomson scattering system, and spectrally filtered camera images of Helium singlet and Hydrogen Balmer lines. Synthetic data sets, generated with realistic experimental errors from MAST-U Super-X SOLPS-ITER simulations, demonstrated D-MIBAS's ability to accurately infer electron temperature and electron density fields under both attached and detached divertor conditions.

The analysis found that PMI must be accounted for in the Balmer line forward models in order to achieve agreement between the true (those used to produce the synthetic data set) and inferred plasma parameters. Our findings of Chapter 3 highlighted that similarities between EIE and PMI emission prevented accurate inferences of all plasma fields unless additional information relating to the electron temperature being above approximately 3eV was provided. The analysis of this chapter found

that, even when combining multiple diagnostics, significantly diverse plasma parameter combinations could accurately replicate the data.

The study demonstrated the significant benefits of conducting inference on a mesh aligned to surfaces of constant poloidal magnetic flux. The mesh allowed various spatially-dependent priors to be included in the inference. These helped to isolate physically plausible parameter combinations (for example, implicitly giving insight into whether the local electron temperature is likely above or below $\sim 3\text{eV}$). This, in turn, substantially improved the accuracy and precision of inferred plasma parameters. The approach achieved a median absolute percentage error within 6% for electron temperature and 5% for electron density across the divertor for detached and attached plasma conditions.

A drawback of using the mesh based method is the requirement to use the complicated HMC method for characterising the posterior distribution. However, the synthetic testing of this chapter highlights that uncertainties in the plasma state found by this method can capture the error between the true and inferred plasma state. This opens up the possibility of obtaining accurate two-dimensional T_e and n_e inferences for MAST-U divertor experimental data.

Chapter 5

Divertor MIBAS as Applied to Experimental Data

In this chapter, we outline the methods taken to adapt D-MIBAS for use with experimental data. In the synthetic validation study of Chapter 4, we found that D-MIBAS’s ability to incorporate spatially-dependent information and multiple diagnostic information significantly improved inference precision. This was largely due to its ability to differentiate between the contributions of electron-ion excitation, plasma-molecule interaction, and electron-ion recombination to Balmer line emission. This differentiation enabled accurate and low-uncertainty inference of the plasma state in both attached and detached divertor conditions. However, comprehensive diagnostic models were found to be essential for accurate inference. In the synthetic study, the spatial and absolute calibrations of the various diagnostics were considered known. Uncertainty in these calibrations, as often encountered in experimental data, presents an additional significant challenge to achieving high-accuracy inferences of the plasma state with low uncertainty.

To address the complexities of the experimental data, three approaches were implemented. These approaches build upon the findings of the synthetic validation study. First, additional nuisance parameters were introduced to parameterise unknown quantities, such as absolute calibration values, ensuring comprehensive forward models (Section 5.2). Second, additional diagnostic information from the divertor region was incorporated into the inference (Section 5.2). Finally, prior information that utilised spatially dependent data was included in the inference (Section 5.3).

D-MIBAS is designed to be adaptable to any divertor. A key aim of the MAST Upgrade (MAST-U) tokamak is to investigate alternative divertor configurations [75]. Analysis of this divertor is complicated due to substantial molecular densities, which contribute significantly to Balmer line emissions through plasma-molecule interactions

[114, 116, 123]. With D-MIBAS able to distinguish between different contributions to Balmer line emission, our focus is initially directed toward the analysis of the MAST-U divertor.

5.1 Mega Amp Spherical Tokamak Upgrade

Mega Amp Spherical Tokamak Upgrade (MAST-U) is a spherical tokamak with a similar minor plasma radius ($a=0.65$ m) and major plasma radius ($R_0=0.85$ m) [38]. MAST-U's small aspect ratio, defined as the ratio of major to minor radius, leads to significant variations in the toroidal magnetic field across the plasma cross-section. Consequential drifts arising from the strong magnetic field gradients in a low aspect ratio tokamaks result in an asymmetry in the power balance, with the majority of the injected power directed to the outer divertor strike points [19].

In the context of exhaust, the variation in the toroidal magnetic field allows for significant total flux expansion between the magnetic X-point and the divertor chamber. This has been shown to facilitate the onset of detachment, allowing for low electron temperatures and substantial plasma-molecule interactions [116]. MAST-U features configurable poloidal magnetic fields, enabling various divertor configurations [75]. It also offers flexible core and divertor fuelling capabilities, along with impurity and helium seeding options [38]. Furthermore, MAST-U has extensive divertor diagnostic coverage, including Thomson scattering, Langmuir probes, spectroscopy, multi-wavelength imaging cameras, infrared cameras, a reciprocating probe system, bolometry, and magnetic diagnostics [38]. These capabilities make MAST-U an ideal test bed for our D-MIBAS system.

MAST-U's divertor region differs in several important respects from that envisaged for future fusion power plants. Conventional diagnostics such as Langmuir probes are not expected to survive the high heat and particle fluxes anticipated in reactor environments [46]. Thomson scattering systems require large optical access and substantial port space close to the plasma, and are therefore impractical for power plant style devices [20]. The interpretation of any spectroscopic diagnostics is expected to be more challenging for power plant style devices. MAST-U has carbon walls throughout the divertor, whereas future reactors will be tungsten-armoured, leading to increased surface reflectivity [34]. Moreover, MAST-U typically operates at relatively low divertor densities ($\lesssim 10^{21}$, m^{-3}), resulting in reduced photon opacity. In reactor regimes, enhanced photon opacity and wall reflections will complicate the analysis of line-emission spectra [70]. Therefore, D-MIBAS is not viable for future

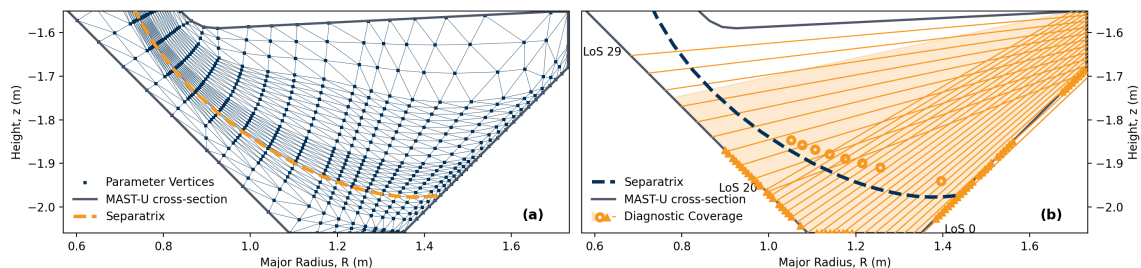


Figure 5.1: Inference mesh (a) and diagnostic coverage (b) for experimental use. Diagnostic coverage describes: spectroscopy - solid lines; multi-wavelength imaging - shaded region; Thomson scattering - hollow circles; and Langmuir probes - filled triangles.

power plant devices. Instead, testing D-MIBAS with a present-day machine enables: refinement of diagnostic forward models; a platform for studying the impact of removing or degrading certain diagnostic channels; and enhanced experimental comparison to high fidelity predictive models.

5.2 Comprehensive Diagnostic Forward Models

Accurate Bayesian inference relies on comprehensive forward models of diagnostic measurements. In this work, measurement errors are assumed to be Gaussian (see Appendix B), but experimental data often contain additional structure beyond simple noise. To account for such effects, we adopt two complementary strategies. First, when uncertainties can be parameterised, we introduce nuisance parameters into the forward models. These parameters vary within expected bounds during inference, and their marginalisation propagates the associated uncertainty into the inferred plasma state. Second, when uncertainties cannot be easily parameterised, we process the diagnostic data to identify regions poorly represented by the forward model. For these points, the model error is inflated, allowing us to retain a Gaussian likelihood form. While the discussion below focuses on MAST-U data, the same principles extend to other tokamak divertors.

5.2.1 Full Mesh Coverage

To comprehensively describe the divertor plasma, the inference mesh in Figure 4.1 was extended to fully occupy the divertor chamber. This is illustrated in Figure 5.1a.

Figure 5.1b shows the diagnostic coverage in the MAST-U divertor. Experimental models for these diagnostics are discussed in detail in this section.

5.2.2 Toroidal Asymmetries

The diagnostic forward models in this work assume a toroidally symmetric plasma. However, certain plasma phenomena, such as strike-point splitting (see Figure 5 of [133]) are not toroidally symmetric. To account for this, measured data are averaged over a 10-millisecond time window. Consequently, the inferred plasma state reflects the time-averaged plasma state rather than a snapshot at a specific moment.

Another source of toroidal asymmetry is the ‘ripple’ in MAST-U’s divertor tile 5 [75]. The tile varies in major radius throughout a toroidal rotation (see Figure 1 of [18]). With diagnostics positioned at different toroidal locations, an offset is expected when collapsing forward models for the various diagnostics into a single poloidal cross-section. This ripple in tile 5 is considered to affect the Langmuir probe (Section 5.2.3.2) forward model and forward models of camera and spectroscopy chords passing close to tile 5 (Sections 5.2.3.3 and 5.2.3.4).

5.2.3 Included Diagnostics

The log-likelihood, in D-MIBAS, is the sum of multiple log-probability distributions associated with different diagnostics

$$\mathcal{L}(\mathcal{D}|\theta) = \sum_i \mathcal{L}_i(\mathcal{D}|\theta). \quad (5.1)$$

We assume a Gaussian likelihood with no correlation for all diagnostics, such that

$$\mathcal{L}_i(\mathcal{D}|\theta) = -\frac{1}{2} \sum_j \left(\frac{\mathbf{d}_{i,j} - \mathcal{M}_{i,j}(\theta)}{\boldsymbol{\sigma}_{i,j}} \right)^2. \quad (5.2)$$

Here, $\mathbf{d}_{i,j}$, $\mathcal{M}_{i,j}$, and $\boldsymbol{\sigma}_{i,j}$ are the data, model, and uncertainty for the j ’th channel of the i ’th diagnostic respectively.

5.2.3.1 Thomson Scattering

Thomson scattering diagnostic systems examine the light scattered from a laser incident on plasma. The intensity of this scattered light is proportional to the electron density [127]. The Doppler effect shifts the wavelength of the scattered light due to the relative velocity of electrons in the plasma. Assuming a Maxwellian electron distribution, the isotropic thermal motion of electrons causes symmetric wavelength

broadening in the detected scattered light. The extent of this broadening relates to the electron temperature [127]. However, in detached plasmas, ionisation leads to ion and electron flows that create non-Maxwellian electron distributions. Additionally, the relatively low electron densities in the MAST-U divertor (commonly below 10^{20}m^{-3}) result in low levels of scattered light, leading to a poor signal-to-noise ratio (the noise arises from background radiation, further complicated by neutral and impurity species) [14]. These factors complicate the application of the Thomson scattering forward model to the raw scattered light data.

To address challenges in Thomson scattering diagnostics in the divertor, long acquisition times and increased scattering volumes are used alongside advanced fitting routines [15]. We utilise these validated analysis procedures [14, 15] at MAST-U, providing forward models for electron temperature and electron density at specific locations in the poloidal cross section. This includes the error identified in the MAST-U analysis procedure [15]. However, the low-frequency Thomson scattering laser system (30 Hz) limits data averaging across the preferred 10 millisecond window, which is necessary to account for toroidal asymmetries in the plasma (Section 5.2.2). Consequently, Thomson scattering data do not reflect the time-averaged plasma state but rather the nearest available time point. Thus, the error in our forward model must be larger than those reported in [15]. The uncertainty of the Thomson scattering electron density forward model was taken to be

$$\sigma_{n_e}^{TS} = \max\left(0.1n_e^{TS}, \varepsilon_{n_e}\right) \quad (5.3)$$

where ε_{n_e} is the Thomson scattering error reported from the analysis according to [15]. The same method was applied for the Thomson scattering electron temperature analysis.

5.2.3.2 Langmuir Probes

Langmuir probes routinely measure the ion saturation current, $j_{sat.}$, which corresponds to the local ion current density at the sheath edge. Under the assumptions outlined here, $j_{sat.}$ can be parameterised in terms of the electron temperature n_e and electron density T_e in the bulk plasma (i.e. at least several Debye lengths into the plasma, upstream of plasma-sheath effects). By accounting for the grazing angle between the probe's surface normal and the plasma's magnetic equilibrium, the ion current density along surfaces of constant magnetic flux at the target can be derived from the Langmuir probe measurements. Given the narrow extent of the sheath, the

ion current density at the target surface is assumed to be equal to that at the sheath edge. This is given by,

$$j_{sat.} = en_{se}c_s \quad (5.4)$$

where e is the elementary charge, c_s is the ion sound speed, and n_{se} is the ion density at the sheath edge. As described in equations (1.42) and (1.49-1.51) of [100], momentum conservation dictates that acceleration towards the sheath results in a decrease in ion density at the sheath edge compared to the upstream plasma. However, the exact location of this upstream position and the factor by which ion density increases depend on various assumptions outlined in Sections 1.8.2, 2.3, and 9.4 of [100].

The ion sound speed can be generalised as follows:

$$c_s = \sqrt{\frac{e(T_e + \gamma T_i)}{m_i}}, \quad (5.5)$$

where the electron and ion temperatures (T_e and T_i , respectively) are measured in eV, and m_i is the ion mass. γ describes the nature of the plasma flow towards the target. To address these unknowns, we introduce a nuisance parameter β , allowing the ion saturation current density measured by the Langmuir probe can be modelled as follows:

$$\mathcal{M}_{j_{sat.}}^{LP}(\theta) = \frac{en_e}{2} \sqrt{\frac{eT_e(1+\beta)}{m_i}}. \quad (5.6)$$

Here,

$$\beta = \gamma \frac{T_i}{T_e} \frac{4}{\alpha^2} + \frac{4}{\alpha^2} - 1, \quad (5.7)$$

where α is the ratio of electron density in the bulk plasma to ion density at the sheath edge, n_e/n_{se} . Common assumptions suggest that $\alpha = 2$ (the ion density halves from upstream to the sheath edge, quasi-neutrality exists upstream, and the plasma is hydrogen-dominated) and $\gamma=1$ (isothermal flow in the target vicinity). Under these assumptions, β describes the ratio of ion to electron temperature, T_i/T_e , and thus the inferred β and T_e provide a crude approximation of T_i in this work. To avoid overfitting, the same parameter β is used to describe all probe measurements.

The ripple of MAST-U's tile 5 complicates the description of the Langmuir probe model. To address this, synthetic studies were conducted to spatially adjust the radial position of the tile 5 probes based on expected displacements caused by the tile ripple. The studies indicated that corresponding uncertainties could reach 5% of the measured $j_{sat.}$. Consequently, the uncertainty of the forward model was taken to be

$$\sigma_{j_{sat.}}^{LP} = \max\left(0.05\overline{j_{sat.}}, s_{j_{sat.}}\right). \quad (5.8)$$

That is, the greater between 5% of the mean $j_{sat.}$ ($\overline{j_{sat.}}$) or the standard deviation of the measured $j_{sat.}$ over the 10-millisecond time window ($s_{j_{sat.}}$).

5.2.3.3 Multi Wavelength Imaging Camera

MAST-U's multi-wavelength imaging (MWI) camera provides a tangential view of the divertor plasma (Figure 1 of [133]). A mirror array directs light into multiple camera channels, each fitted with a band-pass filter that limits the incoming light to a specific wavelength range corresponding to a particular spectral line.

The brightness of each camera pixel corresponds to the emissivity of the selected spectral line, integrated along the line of sight (a viewing chord). An example chord is shown in Figure 5.2.a.ii. A geometry matrix, G , relates the camera's tangential view to the inference mesh in the poloidal plane. This matrix is computed by tracking the path of each chord through the inference mesh and determining the relative contribution of each mesh vertex to the total line-of-sight integral. The geometry matrix elements for the example chord are illustrated in Figure 5.2.b.

Since the camera focuses at a finite distance, each pixel collects photons from a cone rather than a single line (chord). This is modelled by splitting each chord into seven evenly spaced sub-chords, as shown in Figure 5.2.a.i. The resulting geometry matrix is derived by averaging the contributions from each mesh vertex across these sub-chords.

The forward model of the camera channel for spectral line i and pixel j is given by:

$$\mathcal{M}_{i,j} = c_i \sum_k G_{i,j,k} (\varepsilon_{i,k} + b_i) \quad (5.9)$$

where k represents the k^{th} mesh vertex. c_i is a nuisance parameter accounting for uncertainties in the absolute calibration of the camera. The nuisance parameter b_i represents background radiation from the continuum, photon opacity, stray light within the optical system, unaccounted vignetting effects, and interference effects. To avoid overfitting, a single background parameter was assigned to each camera channel i .

Spatial calibration issues were not easily addressed in the camera forward model. Any mismatch between the measured data and the forward model poses a risk of systematic offsets in the inference. Sharp features exist in the camera images, such as tile 3, indicated by the dashed arrows in Figures 5.2.c and d. A single pixel offset can lead the inferred results to place emission at anomalous locations. Additional challenges include the precise edge of the measured image, damage to the optical

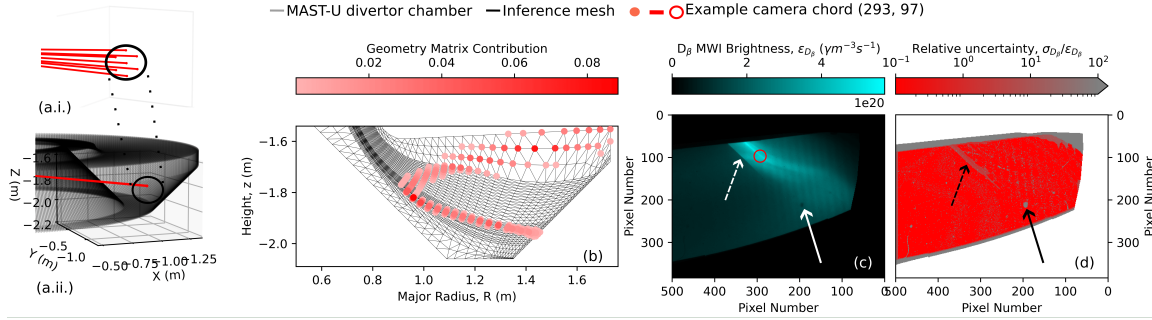


Figure 5.2: MWI D_β channel image processing.

(a.i.) Single camera viewing chord comprised of 7 evenly spaced individual chords
 (a.ii.) Camera viewing chord, for pixel (293, 97), through MAST-U lower divertor chamber.

(b) Camera chord passing through an inference mesh. Geometry matrix contribution for each mesh vertex for chord (293, 97).

(c) D_β MWI brightness image. Chord (293, 97) circled.

(d) Relative uncertainty (uncertainty / brightness) for D_β MWI brightness image.

system (solid arrow of Figures 5.2.c and d), and artefacts, such as visible vertical banding caused by interference (Figures 5.2.c). To mitigate these effects, a 5 by 5 grid was centred on each objective pixel, and the standard deviation of the pixel values within this grid was calculated.

The uncertainty of each pixel in the MWI forward models was taken to be

$$\sigma_{i,j}^{MWI} = \max\left(0.1\overline{B_{i,j}^{MWI}}, s_{B_{i,j}}\right). \quad (5.10)$$

That is, the greater between 10% of the mean pixel j brightness of camera i over the 10-millisecond time window ($\overline{B_{i,j}^{MWI}}$) and the standard deviation of neighbouring pixels (25 including the central pixel) $s_{B_{i,j}}$ over the same time window. The result of this uncertainty calculation is shown in Figure 5.2.d.

The MWI system includes a band-pass filter for the Balmer lines $\{3, 4, 5, 6\} \rightarrow 2$. To model these camera channels, the Balmer emission model from equation (3.1) is applied in equation (5.9). Although helium band-pass filters are available for the helium singlet lines discussed in the synthetic study of Chapter 4 (as well as a Balmer $7 \rightarrow 2$ filter), issues with their data necessitated resulted in their omission in the experimental analysis (Chapters 6 and 7). The MWI also features a band-pass filter for the Fulcher band, for which the emissivity can be modelled through equation (3.6). This data is included in the experimental analysis (Chapters 6 and 7), although to acknowledge the increased uncertainties in the Fulcher-band model, equation (5.10) is altered to $\sigma_{\text{Fulcher},j}^{MWI} = \max\left(0.2\overline{B_{\text{Fulcher},j}^{MWI}}, s_{B_{\text{Fulcher},j}}\right)$.

5.2.3.4 Spectroscopy

MAST-U features spectroscopy systems that offer multiple lines of sight in the poloidal plane (Figure 2b of [114]). These lines of sight are modelled similarly to individual camera pixels, as detailed in Section 5.2.3.3, but with a wider collection angle. With spectrometry data available across a broad wavelength range, background emission can be effectively removed. Therefore, the spectrometry data for the brightness of spectral line i and line of sight j was provided by,

$$\mathcal{M}_{i,j} = c_i \sum_k G_{i,j,k} \varepsilon_{i,k}, \quad (5.11)$$

where k represents the k^{th} mesh vertex. In this work, Balmer $\{3, 5, 6\} \rightarrow 2$ spectral lines were examined, with their data modelled using the Balmer line emission model described in equation (3.1) for the emissivity, ε of equation (5.11).

Due to the electric field in the plasma, each spectral line experiences Stark broadening. Consequently, equation (5.11) describes the brightness integrated over a wavelength window. Since the electron density influences the electric field, the extent of Stark broadening offers insights into the electron density. Thus, a forward model of each spectral line shape can be developed. However, such information on Stark broadening is currently omitted and will be addressed in future work.

5.3 Prior Probability Functions

The posterior distribution comprises multiple likelihood and prior functions. To ensure D-MIBAS achieves *tractable* inference (i.e. inference not dominated by large uncertainties, $\gtrsim 100\%$), despite the complexities of experimental data, a variety of prior functions beyond those used in the synthetic study of Chapter 4 were examined. The total log-prior probability, \mathcal{L} , in D-MIBAS is the sum of several individual priors (j).

$$\mathcal{L}(\theta) = \sum_j \mathcal{L}_j(x(\theta)). \quad (5.12)$$

The arguments x and the probability functions \mathcal{L} for these priors, as applied to experimental data in Chapters 6 and 7, are outlined in this section.

5.3.1 Calibrations

The impact of absolute calibration uncertainties is minimised in alternative work [48, 85] that considers the ratio of two emission lines. This minimisation occurs because the calibration error functions as a multiplicative factor for the measured brightness, consistent across each pixel in a camera and uniform across multiple camera channels, provided that the same instrument calibrates each filtered camera channel. Given this context, the calibration parameters in equations (5.9) and (5.11) are expected to be similar. To incorporate this information, the prior probability function was developed.

$$\mathcal{L}_{\text{cali.}}(\theta) = -\frac{1}{2}(\mathbf{c} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{c} - \boldsymbol{\mu}). \quad (5.13)$$

Here \mathbf{c} is a vector of all calibration values, and $\boldsymbol{\mu}$ is the expected value of these calibrations (typically one). A single standard deviation was used for all calibrations, σ , which describes the uncertainty of the calibration instrument. Since all calibrations were performed with the same calibration instrument, the correlation between different calibration values, ρ , was expected to approach 1. These values populated the covariance matrix of the camera calibrations through,

$$\Sigma_{ij} = \sigma^2 \begin{cases} 1, & i = j \\ \rho, & i \neq j \end{cases}. \quad (5.14)$$

For the experimental data in this work, $\rho = 0.99$ and $\sigma = 0.01$ was found to be appropriate.

5.3.2 Gaussian Process Priors

In the synthetic validation study of Chapter 4, spatially dependent information, encoded in the prior probability function, proved highly beneficial. The spatial smoothness of the parameter fields T_e , n_e and n_D was ensured by assigning low probabilities to fields with significantly nonzero Laplacian values, applied independently for dimensions parallel and perpendicular to surfaces of constant poloidal magnetic flux. The chosen probability function was designed to reflect the distribution of the Laplacian values (independently for each dimension) for each field derived from a range of SOLPS-ITER simulations. However, this distribution was found to be sensitive to the divertor geometry, prompting the need for a more generic method.

In our experimental work, we assume that the logarithm¹ of the collection of field values at all mesh vertices follows a multivariate Gaussian distribution. Consequently, each field in the plasma state θ is a sample from the Gaussian process that

¹The natural logarithm of the fields is taken since we expect the logarithm of the field to vary more smoothly than the field itself.

describes that plasma field. This follows from the argument of a Gaussian process (Section 2.4.2.2). For the field F_j (extracted from the total plasma state vector θ , see equation (4.6)), the overall log-prior probability, $\mathcal{L}(\theta)$, includes a term

$$\mathcal{L}_j(F_j) = -\frac{1}{2} (\ln(F_j) - \boldsymbol{\mu}_j)^T \boldsymbol{\Sigma}_j^{-1} (\ln(F_j) - \boldsymbol{\mu}_j). \quad (5.15)$$

$\boldsymbol{\mu}_j$ represents the mean function of the field, while $\boldsymbol{\Sigma}$ illustrates the correlation between different mesh vertices. $\boldsymbol{\Sigma}$ is populated using a kernel, as discussed in Section 2.4.2.2. Both $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ consist of hyperparameters that need to be deduced.

Importantly, this approach is not just restricted to fields that are explicitly parameterised through θ . Fields derived from θ , such as molecular density and parallel conducted heat flux, can also be subject to a Gaussian Process prior. In this thesis, we demonstrate this approach through just the static electron pressure ($n_e T_e$).

5.3.2.1 Hyperparameter Choice

The choice of hyperparameters, which specify the mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ functions, significantly impacts the posterior distribution. Although hyperparameters can be treated as free random variables within the inference process, doing so incurs substantial computational cost. To mitigate these costs, we pre-select hyperparameters using an empirical Bayesian approach (see Section 2.4.1.3).

Pre-selection of hyperparameters necessitates that resultant functions accommodate the wide range of expected divertor scenarios. To provide this flexibility, the Matérn and Rational Quadratic (RQ) kernels, supplemented by a white noise term, were considered. Both Matérn and RQ kernels are parameterised by a length scale l , amplitude s , and a shaping parameter κ . Example samples are shown in Figure 5.3. The covariance between two points in the field, \mathbf{X} and \mathbf{X}' (e.g. two mesh vertices) is given by:

$$\boldsymbol{\Sigma}(\mathbf{X}, \mathbf{X}') = K(\mathbf{X}, \mathbf{X}'; s, l, \kappa) + \sigma^2. \quad (5.16)$$

The inclusion of a white-noise term, σ , within the kernel plays a subtle yet important role. Firstly, it enhances numerical stability, which is particularly valuable for posterior characterisation via HMC sampling. Secondly, it increases model flexibility by enabling sharp, discontinuous variations in the underlying *latent* fields to be represented. This interpretation differs from the conventional GP formulation, where noise is assigned only to the data and the latent field is assumed to remain smooth.

The shaping parameter, κ , of the Matérn kernel is defined by ν , indicating the number of times, n , that the functions are differentiable ($\nu = n + 1/2$). The Matérn

kernel is given by

$$K_{\text{Matérn}}(\mathbf{X}, \mathbf{X}'; s, \boldsymbol{\ell}, \nu) = s^2 \cdot \frac{2^{1-\nu}}{\Gamma(\nu)} \left(r(\mathbf{X}, \mathbf{X}'; \boldsymbol{\ell}) \sqrt{2\nu} \right)^\nu K_\nu \left(r(\mathbf{X}, \mathbf{X}'; \boldsymbol{\ell}) \sqrt{2\nu} \right) \quad (5.17)$$

where Γ is the gamma function and K_ν is the modified Bessel function of the second kind. r represents the scaled distance between points in the field \mathbf{X} and \mathbf{X}' as given by:

$$r(\mathbf{X}, \mathbf{X}'; \boldsymbol{\ell}) = \sqrt{\sum_{i=1}^d \left(\frac{x_i - x'_i}{\ell_i} \right)^2}. \quad (5.18)$$

Here ℓ can be defined per dimension i (representing an anisotropic kernel) or uniformly across all dimensions (representing an isotropic kernel).

The shaping parameter, κ , of the Rational Quadratic (RQ) kernel is defined by α and provides relative weighting to shorter and longer length scales. The RQ kernel is given by

$$K_{\text{RQ}}(\mathbf{X}, \mathbf{X}'; s, \boldsymbol{\ell}, \alpha) = s^2 \left(1 + \frac{r(\mathbf{X}, \mathbf{X}'; \boldsymbol{\ell})^2}{2\alpha} \right)^{-\alpha}. \quad (5.19)$$

As $\alpha \rightarrow \infty$, the RQ kernel converges to the extremely smooth squared exponential (SE) kernel, which corresponds to the Matérn kernel in the limit $\nu \rightarrow \infty$.

To select a single set of hyperparameters for the Gaussian Process priors, SOLPS-ITER simulations across various divertor geometries and conditions were collated. Hyperparameters $(\hat{\mu}_j, \hat{s}_j, \hat{\boldsymbol{\ell}}_j, \hat{\kappa}_j, \hat{\sigma}_j^2)$ for each field of interest, j , (e.g. electron temperature) were *learned* in an empirical Bayesian manner by selecting those hyperparameters that maximised the summed log marginal likelihood when fitting to T different SOLPS-ITER data sets. Each dataset contained vertex coordinates $\mathbf{X}_j^{(t)}$ and *physical* field values $\exp(\mathbf{y}_j^{(t)})$ (where $\mathbf{y}_j^{(t)}$ are the *latent* field values being modelled by the GP). The summed log marginal likelihood was calculated through

$$(\hat{\mu}_j, \hat{s}_j, \hat{\boldsymbol{\ell}}_j, \hat{\kappa}_j, \hat{\sigma}_j^2) = \arg \max_{\mu_j, s_j, \boldsymbol{\ell}_j, \kappa_j, \sigma_j^2} \sum_{t=1}^T \log p(\mathbf{y}_j^{(t)} \mid \mathbf{X}_j^{(t)}, \mu_j, s_j, \boldsymbol{\ell}_j, \kappa_j, \sigma_j^2), \quad (5.20)$$

where

$$\log p(\mathbf{y}_j^{(t)} \mid \mathbf{X}_j^{(t)}, \mu_j, s_j, \boldsymbol{\ell}_j, \kappa_j) = \quad (5.21)$$

$$- \frac{1}{2} (\mathbf{y}_j^{(t)} - \mu_j)^\top \left(K_{s_j, \boldsymbol{\ell}_j, \kappa_j}(\mathbf{X}_j^{(t)}, \mathbf{X}_j^{(t)}) + \sigma_j^2 I \right)^{-1} (\mathbf{y}_j^{(t)} - \mu_j) \quad (5.22)$$

$$- \frac{1}{2} \log \det \left(K_{s_j, \boldsymbol{\ell}_j, \kappa_j}(\mathbf{X}_j^{(t)}, \mathbf{X}_j^{(t)}) + \sigma_j^2 I \right) - \frac{nt}{2} \log(2\pi). \quad (5.23)$$

The joint fit defined in equation (5.20) yields hyperparameters that generalise across the various SOLPS-ITER data sets. The resulting values for all fields are

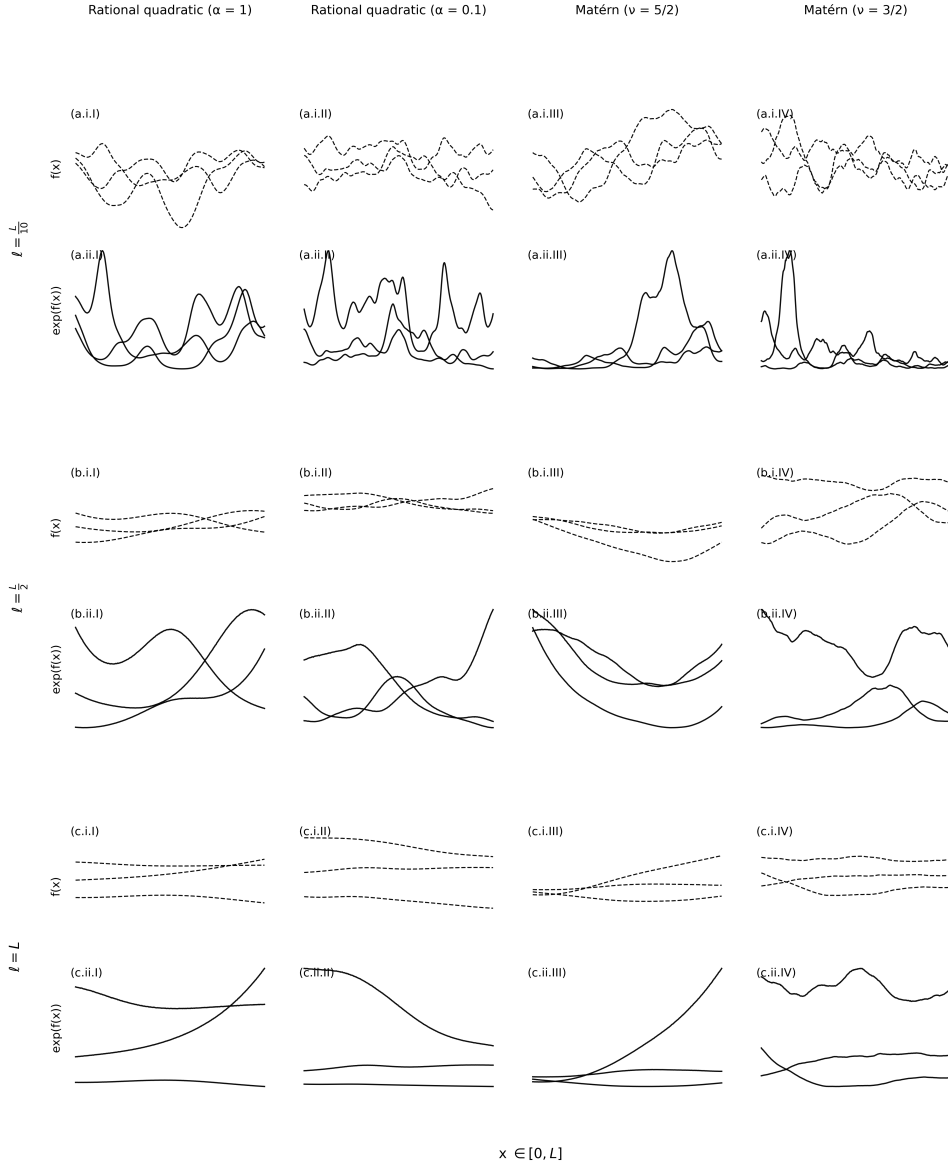


Figure 5.3: Sample functions from different Gaussian Process priors using the Matérn and Rational Quadratic kernels given in equations (5.17) and 5.19 with arbitrary scale, varying lengthscale (ℓ), and varying shaping parameters (α and ν). Dashed curves show *latent* GP samples $f(x)$ and solid curves show the corresponding *physical* fields $\exp(f(x))$. Input defined over a domain from 0 to L with lengthscales, ℓ , given relative to L .

summarised in Table 5.1. To capture the broad range of divertor conditions, the optimisation found relatively short length scales, low smoothness parameters, and appreciable white-noise terms. While this flexibility allows the Gaussian process prior to accommodate diverse plasma states, it reduces its ability to act as a strong

Table 5.1: Empirical Bayes selection of hyperparameters for various plasma fields in the MAST-U divertor. Where two values are provided for the lengthscale, ℓ , an anisotropic kernel is used with the first element corresponding to the direction orthogonal to surfaces of constant poloidal

Quantity, j	Transformed Input field, $f_j(\mathbf{X})$	Input coordinates, \mathbf{X}	κ	μ	s	ℓ	σ
T_e	$\ln(T_e)$	(ψ, s_{\parallel})	$\nu = \frac{5}{2}$	0	1.5	(0.03, 1)	0.05
n_e	$\ln(n_e)$	(ψ, s_{\parallel})	$\nu = \frac{5}{2}$	41.5	1.5	(0.03, 1.5)	0.03
n_D	$\ln(n_D)$	(R, Z)	$\alpha = 0.1$	40.5	1.5	0.08	0.03
$Q_{\text{mol.}}$	$\ln(Q_{\text{mol.}})$	(ψ, s_{\parallel})	$\alpha = 0.1$	0	3.	(0.02, 0.6)	0.03
p_e	$\ln(T_e n_e)$	(ψ, s_{\parallel})	$\nu = \frac{5}{2}$	41.5	2	(0.02, 1)	0.05

constraint. A more effective strategy would be to tailor hyperparameters to specific divertor scenarios, thereby reducing uncertainty in the inferred fields. Ultimately, the optimal approach is to treat hyperparameters as free *nuisance* parameters inferred jointly with the plasma state in a fully Bayesian manner. However, this comes at significant computational cost and is therefore not pursued in this work.

5.3.3 Monotonicity Priors

In Chapter 4, we introduced spatially dependent prior assumptions reflecting the expected monotonic decrease of electron temperature, static electron pressure, and conducted heat flux toward the target along surfaces of constant poloidal magnetic flux. In experimental scenarios, uncertainties in the magnetic equilibrium can compromise the reliability of these assumptions. Consequently, the monotonic behaviour is not strictly enforced but encouraged probabilistically through the use of the Cauchy distribution. Compared to the Gaussian distribution, the Cauchy distribution imposes a weaker penalty on large deviations, allowing for greater flexibility in accommodating plasma states that locally violate prior expectations.

Each monotonicity prior j contributes to the total prior probability equation (5.12) and is applied at each mesh vertex k . The log probability for each monotonicity prior is given by

$$\mathcal{L}_j(x(\theta)) = - \sum_k \ln \left(\pi \gamma \left(1 + \frac{x_j(\theta)|_k^2}{\gamma^2} \right) \right). \quad (5.24)$$

The scale of the Cauchy probability distribution, γ , was determined by perturbing the magnetic equilibrium across various SOLPS-ITER simulations and observing expected violations of the expected monotonicity. The results are shown in Table 5.2.

Table 5.2: Monotonicity priors. $\partial_{s_{\parallel}}$ denotes the derivative, found via finite difference across neighbouring mesh vertices, parallel to the surface of constant poloidal magnetic flux. ‘symlog’ denotes a symmetric log transform, i.e. $f_{\text{symlog}}(x) = \frac{x \log_b(b+|x|)}{b+|x|}$ where b is the logarithm base (in this case, Euler’s number, e).

Prior (j)	Prior argument, $x(\theta) _k$	γ
Electron Temperature Monotonicity	$\max\left(0, \partial_{s_{\parallel}} \ln(T_e)\right) _k$	0.05
Electron Pressure Monotonicity	$\max\left(0, \partial_{s_{\parallel}} \ln(n_e T_e)\right) _k$	0.05
Conducted Heat Flux Monotonicity	$\max\left(0, \partial_{s_{\parallel}} \text{symlog}\left(-\kappa T_e^{5/2} \partial_{s_{\parallel}} T_e\right)\right) _k$	0.1

5.3.4 Upstream priors

The spatially dependent priors extend the domain over which information is known about the plasma. Upstream of the divertor chamber, close to the magnetic X-point, various additional physics arguments can be made. By extending the inference mesh towards the magnetic X-point, these arguments can influence downstream plasma conditions through the spatially dependent priors.

Low neutral populations between the mid-plane and magnetic X-point limit cross-field transport. Consequently, orthogonal to the magnetic flux surfaces and slightly downstream of the magnetic X-point, the static electron pressure and conducted heat flux are expected to peak at the separatrix. Prior arguments were curated to capture this expected single peak at the separatrix for the static electron pressure and parallel conducted heat flux fields as shown in Table 5.3.

MAST-U produces little fusion power. Consequently, the input power to the scrape-off layer (SOL) approximately equals the external heating power. Upstream, near the magnetic X-point, plasma power in the SOL is predominantly transported by conduction. This is given by

$$P_{\text{SOL}}(\theta) \approx P_{\text{SOL cond.}}(\theta) = \int_0^{2\pi} \int_{\perp} -R\kappa T_e^{5/2} \partial_{s_{\parallel}} T_e ds_{\perp} d\phi, \quad (5.25)$$

where ϕ is the toroidal angle and s_{\perp} is the poloidal distance travelled orthogonal to the magnetic flux surfaces. This allows a prior argument to be placed on the inferred power in the plasma carried by conduction as shown in Table 5.3. To help accommodate inconsistencies with the magnetic equilibrium used within the prior arguments of Table 5.3, the Cauchy probability distribution of equation (5.24) was used.

Table 5.3: Upstream priors. $H(\theta) = 1$ for mesh vertices within 30 cm of the magnetic X-point, and 0 otherwise. $\text{sgn}(\theta) = 1$ for $\psi(\theta) < 1$ (mesh vertices in the private-flux region) and -1 otherwise. ∂_{s_\perp} and ∂_{s_\parallel} denote the spatial derivative perpendicular and parallel to surfaces of constant poloidal magnetic flux respectively. These derivatives were found via finite difference across neighbouring mesh vertices. P_{max} is the external heating power supplied to the plasma.

Prior (j)	Prior argument, $x(\theta) _k$	γ
Static electron pressure up-stream peak	$\max(0, H(\theta) \cdot \text{sgn}(\theta) \cdot \partial_{s_\perp}(n_e T_e) _k)$	1
Conducted heat flux up-stream peak	$\max(0, H(\theta) \cdot \text{sgn}(\theta) \cdot \partial_{s_\perp}(-\kappa T_e^{5/2} \partial_{s_\parallel} T_e) _k)$	1
Conducted plasma power	$\max(0, \frac{P_{\text{SOL cond.}}(\theta)}{P_{max}} - 1)$	0.8×10^6

Table 5.4: Parameter bounds used in the inference.

Parameter	Bounds, (lower, upper)	Units	Type
T_e	(0.05, 150)	eV	Field
n_e	(10^{17} , 10^{21})	m^{-3}	Field
n_D	(10^{17} , 10^{20})	m^{-3}	Field
$Q_{\text{mol.}}$	(10^{-5} , 10^5)	-	Field
c_i^{mwi} $i \in \{D_\alpha, D_\beta, D_\gamma, D_\delta, \text{Fulcher}\}$	(0.9, 1.1)	-	Single
c_i^{dms} $i \in \{D_\alpha, D_\gamma, D_\delta\}$	(0.9, 1.1)	-	Single
$b_b^{\text{cam}}(D_\alpha)$	(10^{16} , 10^{20})	γm^{-3}	Single
$b_b^{\text{cam}}(D_\beta)$	(10^{16} , 10^{19})	γm^{-3}	Single
$b_b^{\text{cam}}(D_\gamma, D_\delta, \text{Fulcher})$	(10^{16} , 10^{18})	γm^{-3}	Single
β	(1.0, 20.0)	-	Single

5.3.5 Parameter Bounds

To aid the hamiltonian MCMC sampling used in this work, soft priors were used to provide parameter bounds. These lower (l) and upper (u) bound values are provided in Table 5.4 and enforced via the Gaussian probability distribution with $\sigma = 10^{-3}$:

$$\mathcal{L}_j(\theta) = -\frac{1}{2} \sum_i \left(\frac{\max(0, \frac{\theta_i}{u_i} - 1, \frac{l_i}{\theta_i} - 1)}{\sigma} \right)^2. \quad (5.26)$$

5.4 Covariance Matrix Approximation

In Chapter 4, it was found that an approximation of the covariance matrix was extremely beneficial for Hamiltonian MCMC (HMC) sampling. This was due to the




Figure 5.4: Covariance matrix approximation based on the covariance kernels of equation (5.16). Parameter number represents position in parameter vector θ (see equation (4.6)), with numbers assigned sequentially across the rows and columns of the inference mesh. Labels A to D described in the prose.

covariance matrix providing information on the considerable correlations between parameters. However, the MAP curvature approximation of the mass matrix used was not globally representative of the full posterior distribution. An alternative method for approximating the mass matrix is the *a priori* construction described in Appendix A.3. This predominantly uses the kernels, equation (5.16), of the Gaussian process priors to make a crude approximation of the covariance matrix.

An example of such a covariance matrix approximation is shown in Figure 5.4. Covariance between parameters of the same type, such as T_e , are given by the kernels of equation (5.16). These highlight significant correlations among parameters across the inference mesh; greater correlations are observed between neighbouring parameter vertices (label B of Figure 5.4) and correlation trends exist along rows and columns (label C of Figure 5.4). Approximating the impact of the likelihood on the posterior distribution’s covariance is more complex. A simplified method involves reducing the variance of parameters close to diagnostic data points (e.g., label D of Figure 5.4 represents T_e parameters close to a Thomson scattering data point). Additionally, SOLPS-ITER cases were used to find correlations between different field parameters, notably the strong negative correlation between T_e and n_D . This led to negative covariance assignments for these combinations (label A of Figure 5.4). Overall, these methods provided a basic but effective approximation of the covariance matrix, which sufficiently informed the HMC algorithm with an initial guess of expected correlations.

5.5 Computational Cost

The maximum a posteriori (MAP) estimate was first identified using the evolutionary gradient ascent (EGA) algorithm. A population of 32 candidate solutions was

optimised for 200 steps per generation, after which the population was refined by retaining the highest posterior probability third of candidates, ‘breeding’ another third from high-performing candidates, and replenishing the remaining third through random sampling of the input domain. This process was repeated for 50 generations, balancing global exploration with local optimisation [9]. Candidate optimisation used either the Adam or L-BFGS algorithm.

Following MAP estimation, Hamiltonian Markov Chain Monte Carlo (HMC) sampling was performed from the MAP location. Each chain used 500 burn-in samples followed by 2000 posterior samples. The HMC step size was adaptively adjusted to maintain an acceptance rate of approximately 60%. Four chains were run in parallel, with samples collated and thinned to 100 representative posterior sample states.

Posterior evaluations consisted primarily of matrix multiplications within computationally inexpensive forward models and priors. On an Apple M2 chip with 16 GB of memory, each posterior and gradient evaluation required approximately 0.01 seconds. Analytic first derivatives were derived for both L-BFGS and HMC, minimising the overhead of gradient computation.

Typical D-MIBAS problems involved approximately 5,000 input parameters. At this scale, dense covariance or mass matrix operations, including storage, factorisations, and temporary work arrays, required approximately 0.5–1 GB of memory per optimisation candidate or MCMC chain.

Overall, the computational cost was therefore dominated by the total number of optimisation and sampling steps required for convergence, rather than by individual posterior evaluations. The use of EGA prior to HMC provided physically plausible initial parameter estimates and therefore reduced the convergence time of the HMC algorithm. Both EGA populations and HMC chains were naturally parallelisable, allowing for significant reductions in wall-clock time on suitable multi-core or distributed computing hardware.

5.6 Summary

D-MIBAS performs integrated data analysis by constructing a posterior probability distribution of plausible plasma states, θ , given diagnostic measurements, \mathcal{D} . The likelihood and prior probability distributions that comprise the posterior distribution are themselves comprised of numerous probability distributions, respectively pertaining to different diagnostics, i , and different physics-informed priors, j . Omitting the

normalisation constant, the total log-posterior probability is therefore,

$$\underbrace{\mathcal{L}(\theta|\mathcal{D})}_{\text{log-posterior}} = \underbrace{\sum_i \mathcal{L}_i(\mathcal{D}|\theta)}_{\text{log-likelihood}} + \underbrace{\sum_j \mathcal{L}_j(\theta)}_{\text{log-prior}}. \quad (5.27)$$

This chapter outlines the complexities of integrated data analysis at MAST-U. These complexities primarily arise from challenges in the absolute and spatial calibration of diagnostics. D-MIBAS addresses these challenges by incorporating nuisance parameters into diagnostic models and increasing the uncertainty of these models. However, this approach allows for a broader range of plasma states to describe the diagnostic data, leading to a significant increase in the uncertainty of the inferred plasma states.

The use of an inference mesh allows for spatially-dependent information to be incorporated as prior arguments. In this chapter, we outline several additional prior functions based on physical principles that utilise this spatially-dependent information. These have been developed by analysing multiple SOLPS-ITER simulations of MAST-U and managed to ensure compatibility with the Hamiltonian Markov Chain Monte Carlo sampling algorithm. These additional prior arguments limit the number of plausible plasma states. Thus, their application enables D-MIBAS to be used with experimental data, facilitating low-uncertainty inferences.

Chapter 6

D-MIBAS Inference of a MAST-U Super-X Divertor

In this chapter, we apply D-MIBAS to infer the plasma state throughout a MAST-U discharge operated in the Super-X divertor configuration. This represents the first deployment of D-MIBAS on experimental data, thus this chapter places strong emphasis on validation, verification, and uncertainty quantification (VVUQ). These VVUQ activities utilise the comprehensive description of the divertor plasma state provided by D-MIBAS to: evaluate cross-diagnostic consistency within the inference; compare results against alternative analysis methods; and benchmark directly against interpretive SOLPS-ITER modelling of the MAST-U Super-X divertor. Additionally, we relate the inferred plasma dynamics to the expected behaviour of a detached divertor, providing a qualitative assessment of the physical plausibility of the results.

6.1 Discharge #46860

This study examines the effects of upstream electron density changes on the divertor plasma state. This was investigated by applying D-MIBAS to MAST-U discharge #46860. The discharge involved a core density fuelling ramp and operated in a low confinement mode (L-mode) with a Super-X divertor configuration. It featured 1.6 MW of neutral-beam injection heating and a plasma current of 0.75 MA. A stable divertor configuration was held during the increase in Greenwald fraction from 27% (450 ms) to 48% (790 ms) before the plasma disrupted. Intrinsic carbon, from the wall, and trace helium puffs for diagnostic purposes were present, but no additional impurities were seeded.

Experimental analysis of MAST-U L-mode plasmas in the Super-X configuration has widely indicated a detached plasma state throughout the discharge [39, 122].

The Super-X configuration generates a long connection length from the midplane to the target (19 m) and high total flux expansion (2.3) [122]. In the L-mode plasma, the core density rise is closely correlated with additional particles diffusing into the scrape-off layer (SOL) plasma increasing the SOL plasma density. This increasing SOL density, combined with the geometric effects of long connection length and flux expansion, promotes increased rates of plasma-neutral interactions. This removes power, momentum, and particles from the plasma before they reach the divertor targets, thereby facilitating and sustaining detachment [54, 66, 99].

Detachment can be diagnostically observed as a roll-over in the target ion flux, typically measured using Langmuir probes [100]. This roll-over arises because the ionisation source moves upstream, reducing the replenishment of ions near the target and enabling volumetric recombination. This is observed spectroscopically through changes in Balmer line emission (the relative intensity of high- n to low- n transitions) resulting in a change to the plasma hue [66, 116]. The downstream end of the ionisation region is referred to as the detachment front [122]. This front marks the position along the divertor leg where the plasma begins to cool rapidly and loses momentum and particles through volumetric processes (within the plasma itself, rather than at the target).

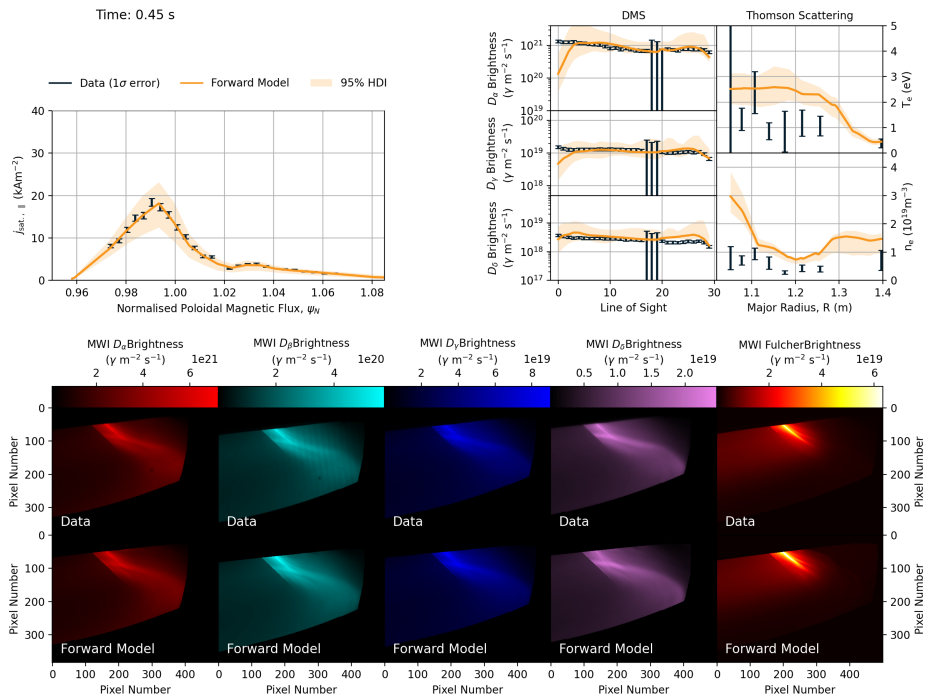
6.1.1 Verification, Validation and Uncertainty Quantification

Due to the novelty of the D-MIBAS technique, this section addresses verification, validation, and uncertainty quantification (VVUQ) to establish the credibility of the results.

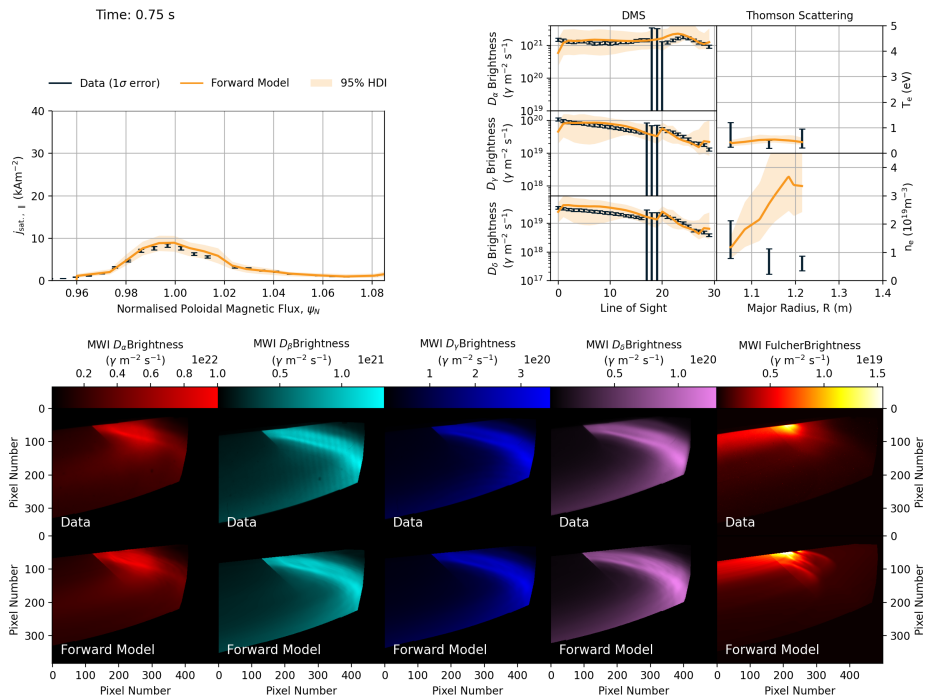
6.1.1.1 Diagnostic Predictions

The plasma state was inferred at time intervals throughout discharge #46860. The inferred plasma states represent the most probable states given the diagnostic measurements, our model assumptions, and the included physics-based priors. Verification of the inference is considered by assessing how well the forward-modelled diagnostics, derived from these inferred states, match the actual diagnostic measurements. Figures 6.1a and 6.1b illustrate this verification, showing the agreement between predicted and measured diagnostic data at the beginning and end of the core density ramp, respectively.

The diagnostic comparisons shown in Figures 6.1a and 6.1b are representative of the inference at all investigated time points and demonstrate reasonable consistency



(a) At 450 ms, a Greenwald fraction of 27%.



(b) At 750 ms, a Greenwald fraction of 45%.

Figure 6.1: Comparison between diagnostic data and forward model at the MAP inference.

across the diagnostics. Notable discrepancies include the Thomson scattering electron density measurements frequently report values below the inferred densities, particu-

larly at low electron temperatures. It may be that this discrepancy is expected, as the Thomson scattering signal weakens at low temperatures and becomes increasingly affected by stray light and instrumental noise. Therefore, it is plausible that the actual measurement uncertainty of the Thomson scattering diagnostic is likely greater than the reported values. However, analysis across more discharges are required to properly investigate this difference.

The most significant discrepancy between diagnostic measurement and inference appears in lines of sight 0–3 (those closest to the target) for the divertor monitoring spectroscopy (DMS). This discrepancy can be partially attributed to alignment issues related to ripple structures in tokamak tile 5; however, it is plausible that the physics-based priors are too restrictive in the target vicinity. The electron density may experience sharper transitions in the target vicinity than suggested by physics-based priors based on SOLPS-ITER modelling. Similarly, it may be that the physics-based priors for electron temperature (i.e. monotonic gradient of the temperature and the conducted heat flux) and pressure are too restrictive in the deeply detached divertor conditions close to the target. To address this D-MIBAS would need to be constructed with separate priors in the bulk plasma, and in the target vicinity. A limitation of this inference approach is that the posterior results reflect a balance of many competing contributions. To establish whether specific features are physically realistic or simply artefacts of the inference, analysis across a broader set of plasma states is a necessary focus for future work.

6.1.1.2 Experimental Comparison to Alternative Methods

Several alternative analysis methods provide complementary insights into the divertor plasma. However, none explicitly infer the plasma state directly in terms of electron temperature, electron density, and neutral density simultaneously, as accomplished by D-MIBAS. Figure 6.2 presents validation exercises comparing quantities derived from the inferred plasma state with results obtained using independent analytical methods.

Figure 6.2.a compares the inferred emissivity profiles obtained from (i) D-MIBAS and (ii) a SART inversion [133] for the $4 \rightarrow 2$ Balmer line (D_β). The observed agreement in both spatial distribution and absolute values validates that the integrated approach captures the emissivity structure. This consistency is noteworthy; unlike the SART inversion, which models the measured brightness in terms of localised emissivity, D-MIBAS models D_β Balmer line emissivity via T_e , n_e , n_D , and Q_{mol} while balancing fits to multiple diagnostics and prior arguments. The agreement was found

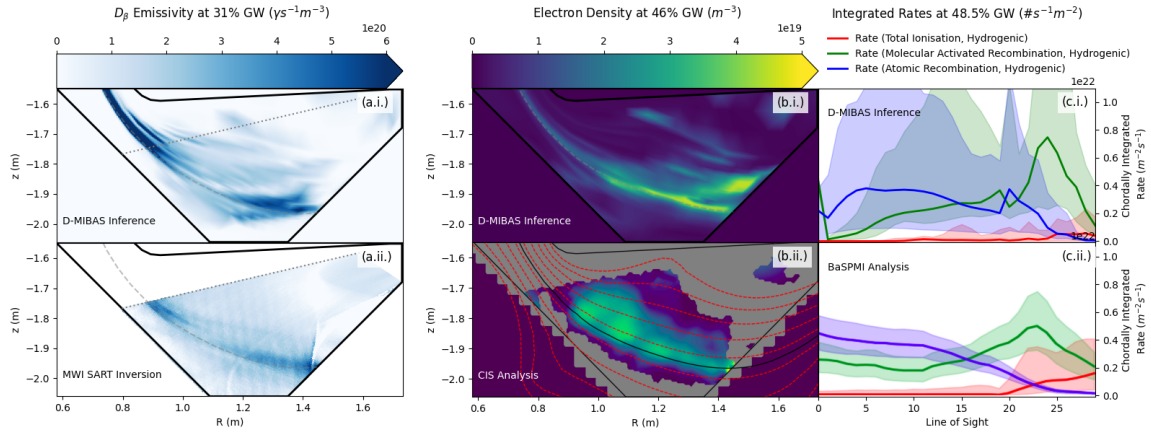


Figure 6.2: Comparison of alternative analysis procedures (bottom row) to results recreated from the D-MIBAS inferred plasma state (top row). Bottom-left: A SART inversion from techniques given in [133] Bottom-centre: Coherence-imaging spectroscopy, recreated from Figure 16 of [72]. Bottom-right: Spectroscopic analysis through the BaSPMI method [112], recreated from Figure 12 of [89]. Note: shaded region corresponds to 95% HDI in D-MIBAS results (top-right), but 67% equal-tailed HDI in BaSPMI Analysis (bottom-right).

to be consistent across different times and camera channels. Notably, the D_β camera channel experiences interference-induced banding patterns, which complicate the SART inversion, as illustrated by the spurious banding structure in Figure 6.2.a.ii. The absence of this artifact in the D-MIBAS inference highlights the robustness of the integrated data analysis framework which prevents over-fitting to any single diagnostic.

Figure 6.2.b provides a comparison of the inferred electron density from D-MIBAS with results from coherence imaging spectroscopy (CIS), which estimates density fields from spectral line broadening [71, 72]. Both approaches yield similar absolute values along the separatrix (roughly $3 \times 10^{19} \text{m}^{-3}$). However, the CIS results suggest a larger electron density at the strike point than D-MIBAS with the CIS result outside of the 95% highest density interval (inferred uncertainty) of the D-MIBAS results. Despite this, given the considerable uncertainties inherent in both techniques (furthermore, the CIS results are from a repeat discharge of discharge #46860 at the same Greenwald fraction), the emphasis here is on reproducing the correct order of magnitude and broad spatial trends. Within this context, the loose agreement between the two methods provides limited validation of the D-MIBAS inferred plasma state.

Figure 6.2.c validates our inferred plasma state through comparison with the BaSPMI analysis method, which employs Balmer line spectroscopy to estimate the rates of hydrogenic ionisation, molecular activated recombination, and atomic re-

combination [112]. Our inferred state reproduces similar trends and magnitudes across multiple lines of sight and time points, thereby strengthening confidence in D-MIBAS's ability to separate the Balmer line emissivity into distinct physical processes.

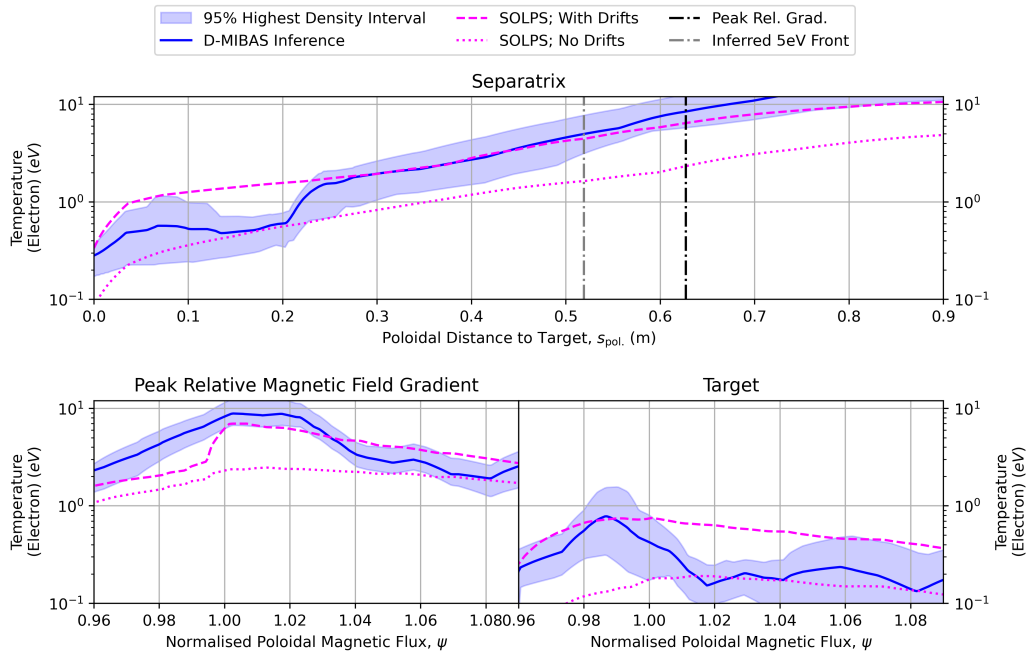
6.1.1.3 Modelling Comparison

SOLPS-ITER has been used for interpretive analysis of discharge #46860, as presented in [122] with data made available in [113]. Unknown input parameters contribute to the uncertainty of the modelling prediction. These include atomic and molecular rate coefficients, cross-field transport effects, and the treatment of plasma drifts [76, 114, 124, 130]. Plasma drifts arise because the plasma experiences forces due to the shape and inhomogeneity of the magnetic field and the presence of electric fields. They drive a systematic drift of ions and electrons, redistributing particles and heat between the inner and outer divertor legs. Specifically, drifts reduce the width of the upstream heat flux and increase the proportion of power going to the outer leg. Although comprehensive uncertainty quantification for SOLPS-ITER is not available, modelling uncertainty is alluded to in this section by presenting SOLPS-ITER predictions with and without the inclusion of drifts.


Comparison to SOLPS-ITER models of discharge #46860 are shown in Figure 6.3 for electron temperature and electron density fields. The comparison is made along the separatrix, cross-field (orthogonal to surfaces of constant poloidal magnetic flux) in the region of peak relative magnetic field gradient (generally close to the 5 eV front), and cross-field, along the target. We observe similar qualitative behaviour and magnitudes between the inferred fields and the SOLPS-ITER predictions of discharge #46860 as shown in Figure 6.3. However, notable differences exist.

One key difference between interpretive SOLPS-ITER modelling and D-MIBAS inference is that both the electron temperature and electron density drop more sharply along the separatrix, close to the target in the SOLPS-ITER result. It is challenging to determine which is *true*. Discrepancies in the DMS lines of sight noted earlier (Figures 6.1a and 6.1b) suggest this may well be an issue with the D-MIBAS results. However, the CIS results of Figure 6.2 suggest that electron density may *increase* at the strike point differing to both the SOLPS-ITER and D-MIBAS results. This highlights the limitations of fine-detailed comparisons between SOLPS-ITER and D-MIBAS at this stage.

Additional analysis across more discharges are required to better contextualise the D-MIBAS results and develop a more robust understanding of the discrepancies



(a) Electron Temperature



(b) Electron Density

Figure 6.3: D-MIBAS inferred plasma state at 600 ms (37% Greenwald fraction) compared to SOLPS-ITER interpretive modelling of discharge #46860 .

between SOLPS-ITER and experimental results inferred through D-MIBAS. This would benefit from the inclusion of additional diagnostics, such as CIS, to D-MIBAS permitting a reduced dependency on priors and a further reduction in inferred un-

certainties. Doing so would enable unprecedented insight and experimental guidance for the various unknown input parameters in SOLPS-ITER simulations.

6.1.1.4 Inference Uncertainty

The percentage uncertainties of the inferred fields, shown in Figure 6.4, are considerably greater than those found in synthetic validation of D-MIBAS in Chapter 4 (see Figure 4.10). The median percentage uncertainties, across all analysed times and mesh vertices, are 60%, 55%, 65%, and 65% for T_e , n_e , n_D , and Q_{mol} fields, respectively. This increased uncertainty reflects the additional complexity of experimental data. The increase is considered to be primarily due to the unknown absolute calibration value of the spectroscopic data. Additionally, the Gaussian process priors enforcing spatial smoothness were made highly flexible to avoid biasing the inferred results. Both of these factors substantially increase the range of plasma states that could plausibly have produced the measured data.

It is non-trivial to determine whether the reported uncertainties (derived from posterior samples obtained via Hamiltonian Markov Chain Monte Carlo) fully represent all sources of error. Reported uncertainties are reduced for parameters that significantly influence the modelled diagnostic data (i.e. those close to Thomson Scattering points and the separatrix, where Balmer line emission is strongest) as expected. In comparison to the synthetic results of Chapter 4, experimental n_D and Q_{mol} uncertainties do not increase as much as T_e and n_e parameters. This is considered to be due to the Gaussian process priors used in the experimental work, which impose strong spatial correlations and significantly impact the posterior uncertainty. This illustrates the dual role of the prior: it regularises the inference in poorly constrained regions, but can also lead to overconfident posteriors if its assumptions are not entirely valid. These considerations highlight the challenges of achieving comprehensive UQ within the D-MIBAS framework. While great care has been taken in the formulation of diagnostic forward models and the construction of prior arguments, in the absence of ground truth or a fully validated diagnostic benchmark, the resulting uncertainty quantification must be regarded as approximate.

6.2 Results: Inferred Plasma State

The inferred parameters are presented in Figure 6.5(a-c.i.-iii.) for three time points during the density ramp, corresponding to 27%, 36%, and 45% of the Greenwald fraction. Utilising ADAS [90] and Yacora-on-the-web [134], we derive the total hydrogenic




Figure 6.4: Percentage uncertainty of inferred parameters at 600 ms (37% Greenwald fraction). 95% HDI used to quantify the uncertainty.

ionisation (Ion.), molecular activated recombination (MAR), and atomic recombination rates (EIR), as shown in Figure 6.5(a-c.iv). At the examined time points, the hydrogenic ionisation region was *detached* from the target. This indicates that the divertor was in a detached state throughout the investigated time points. The 5 eV electron temperature front is found near the downstream end of the hydrogenic ionisation region. In this chapter, it is therefore adopted as a quantitative marker of the detachment front.

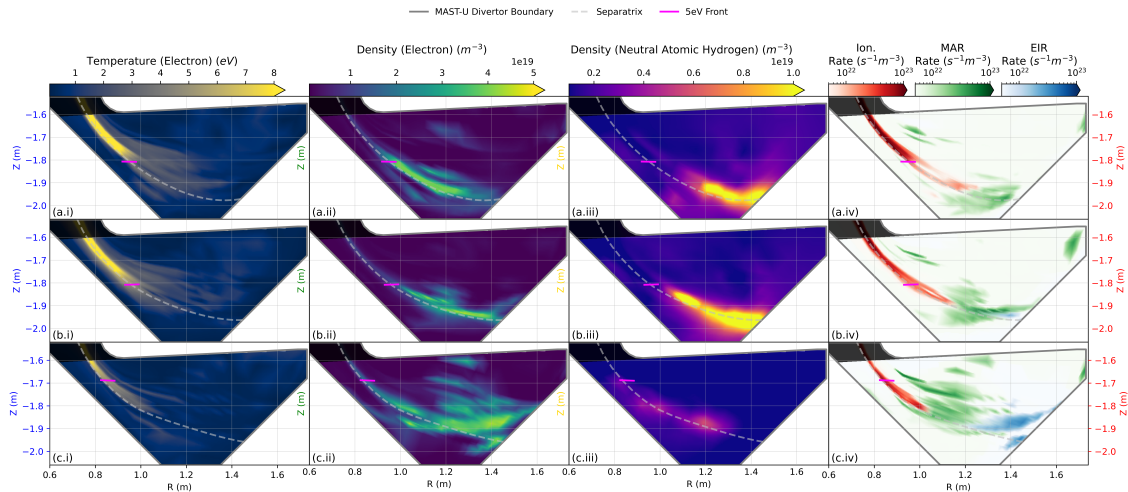


Figure 6.5: Inferred plasma states throughout discharge #46860.

- a) 27% Greenwald fraction.
- b) 37% Greenwald fraction.
- c) 45% Greenwald fraction.

Magenta dash represents the 5 eV electron temperature front. Shaded region represents lack of diagnostic coverage.

The 5 eV electron temperature front recedes gradually during the plasma density ramp, indicating a smooth progression of detachment. Figure 6.5(a-c.i) shows that

this front moves radially inward from a major radius of approximately 1 m to 0.8 m. This behaviour aligns with bolometry observations, which show the radiating plasma region gradually shifting upstream in the outer divertor leg of MAST-U plasmas [25, 27]. The steady movement of this front, even under substantial changes in upstream electron density, suggests a stable and manageable detachment process that is highly advantageous for divertor control [52].

The region of high neutral atomic hydrogen density shifts during the core density ramp, consistent with expectations for a detached divertor. As shown in Figures 6.5(a–c.iii), the atomic neutral density build-up appears downstream of the ionisation front, and its movement tracks the ionising region. This can be explained by recombination processes within the plasma, which provides a source of neutral particles far from the divertor target. Such behaviour supports the interpretation of a moving ionisation region that is fuelled by recombination-produced neutrals [92].

The electron density profile in the divertor broadens across flux surfaces during the density ramp, rather than significantly increasing in peak value along the separatrix. Figures 6.5(a–c.ii) show this broadening away from the separatrix, while the extent of hydrogenic ionisation remains relatively unchanged. This broadening may be driven by stronger plasma–neutral interactions in the divertor chamber. However, limited diagnostic coverage raises the possibility that the broadening could instead result from a widening of the upstream SOL density profile, projecting into the divertor as an apparent broadening.

The dominance of different atomic and molecular processes in the divertor plasma is closely linked to the electron temperature, and the processes shift upstream intrinsically linked to the movement of the 5 eV front. Between 5 and 3 eV, electron-impact ionisation remains dominant; however, as the temperature continues to decrease downstream, molecular activated recombination (MAR) rates surpass ionisation. MAR remains dominant as the temperature continues to decrease until approximately 0.8 eV by which point electron–ion recombination (EIR) becomes the dominant process. These transitions are consistent with observations from JET [47]. During the density ramp, the sequential upstream shift of the 5 eV front therefore carries with it the regions of ionisation, MAR, and EIR dominance, illustrating the evolving particle balance in the detached MAST-U divertor [114, 133].

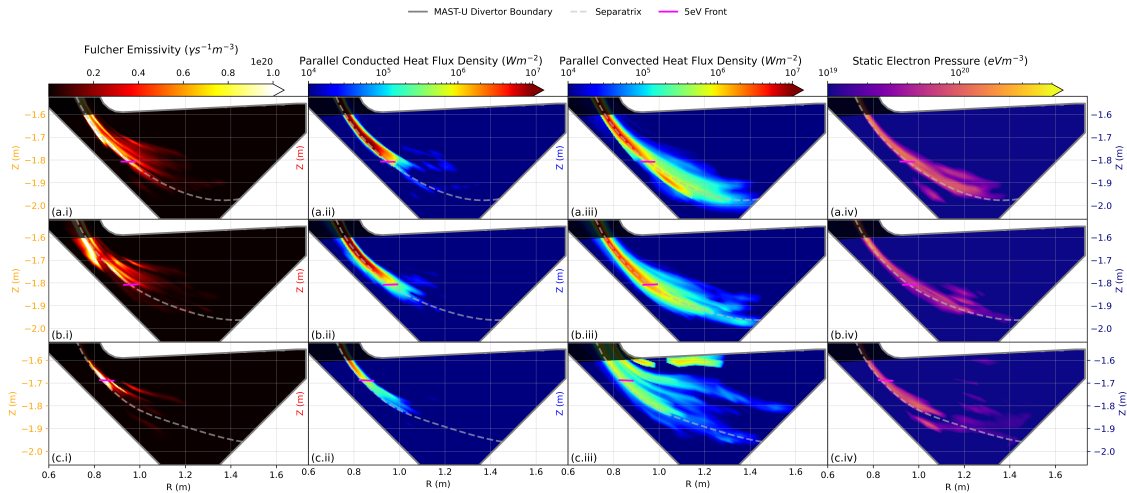


Figure 6.6: Derived fields from the inferred plasma states throughout discharge #46860.

a) 27% Greenwald fraction, 450 ms.

b) 37% Greenwald fraction, 600 ms.

c) 45% Greenwald fraction, 750 ms.

Magenta line represents the 5 eV electron temperature front. Shaded region represents lack of diagnostic coverage.

6.3 Discussion

6.3.1 Progression of the Detached Divertor State

The divertor plasma is observed to be in a detached state at all analysed time points in Figure 6.5. As the 5 eV front moved upstream, a larger region of dominant electron–ion recombination processes (EIR) developed. The cold, heavily recombining conditions at the target imply that substantial power, particles, and momentum are removed before reaching the divertor plates. Such conditions are consistent with a *deeply* detached divertor state [54, 126]. This section focuses on understanding how this detached plasma evolves during the upstream core density ramp.

Figure 6.6 presents profiles of Fulcher band emissivity, parallel conducted heat flux, parallel convected heat flux, and static electron pressure. For visual ease, Figure 6.7 presents the inferred plasma state along the separatrix at the start (Figure 6.7.a, 450 ms, 27% Greenwald fraction) and end (Figure 6.7.b, 750 ms, 45% Greenwald fraction) of the density ramp. Additionally, changes at numerous time points throughout the density ramp are shown in Figure 6.8.

The Fulcher band provides an experimentally accessible proxy for the hydrogenic ionisation region and corresponding 5 eV electron temperature front, making it a

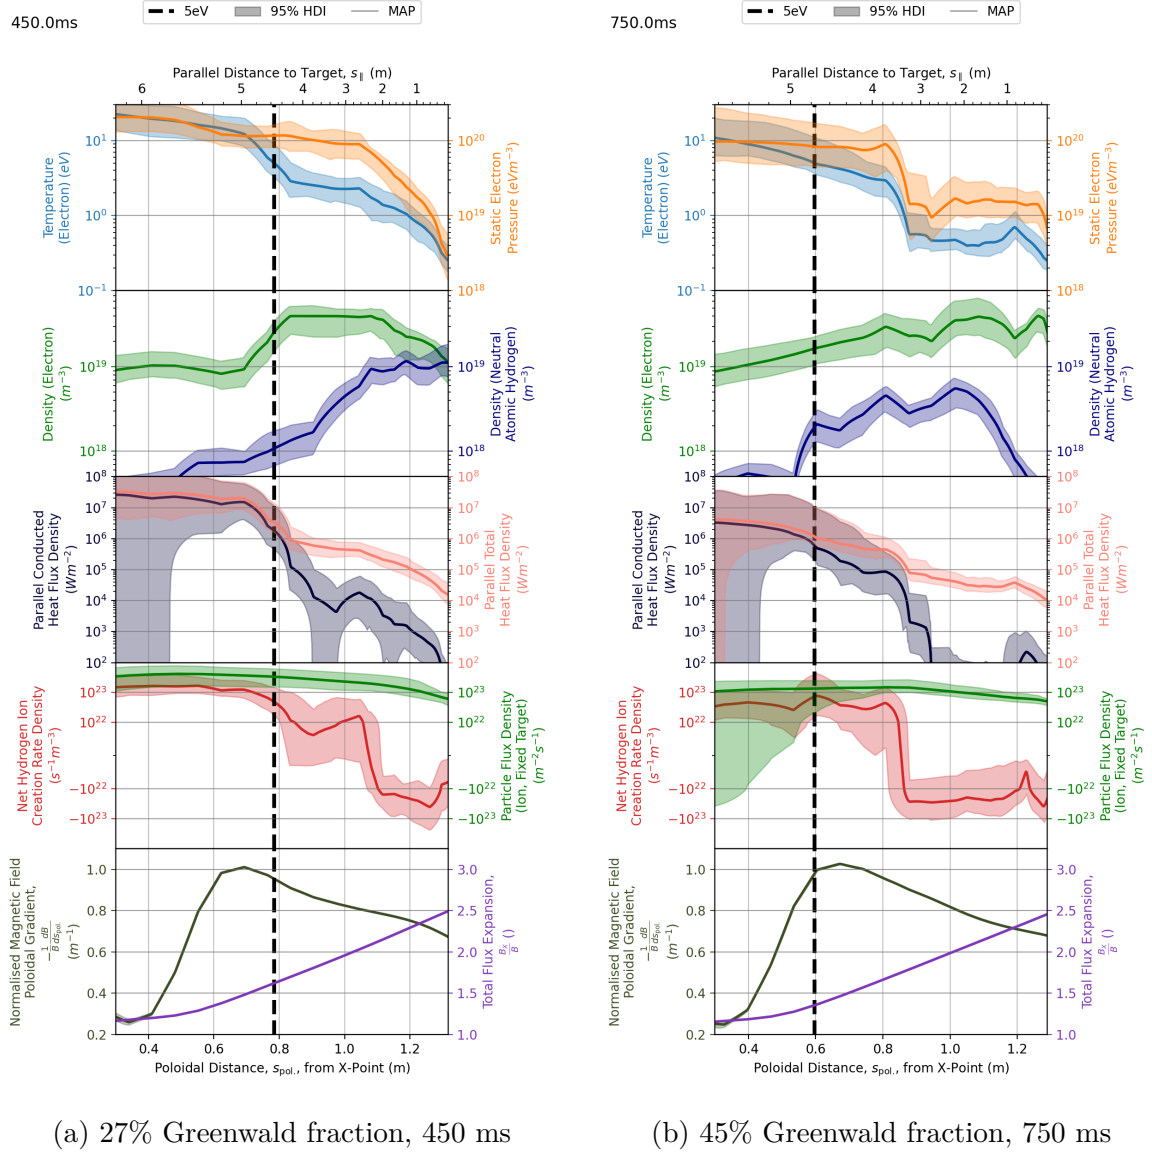


Figure 6.7: Inferred and derived fields following the separatrix ($\psi=1$). Shaded region indicates 95% highest density interval. Particle flux density calculated using equation 2.19 and the assumption of negligible cross-field transport.

powerful tool for tracking detachment. The 5 eV front marks a pronounced change in plasma behaviour and has long been recognised as an indicator of the detached divertor state [66, 100]. During the density ramp, the region of strong Fulcher band emission moves upstream towards the X-point alongside the region of dominant hydrogenic ionisation (Figures 6.5 and 6.6) and variations in other plasma parameters, as shown in Figures 6.6 and 6.7. Fulcher band emission, as demonstrated here and in other studies [57, 82, 83, 120], is particularly valuable because it peaks in emissivity just upstream of the 5 eV front and then decreases rapidly as T_e declines toward 3

eV. As such, it provides a clear, experimentally measurable signal [52, 82, 83] that coincides with the transition region where key detachment processes occur.

The detachment process strongly modifies the particle balance. The net hydrogen ion creation rate density ($S_{\text{ion}} - S_{\text{rec}}$ shown in Figure 6.7) peaks just upstream of the 5 eV front, while the transition to net recombination occurs further downstream at 2–3 eV. The ion flux density reaches its maximum within ionising plasma rather than at the transition to a net-recombining plasma. This is primarily due to the large flux expansion of the Super-X divertor, which spreads ionisation-driven fluxes over a broader cross-sectional area [100, 111]. Additionally, the high neutral densities enhance cross-field diffusion [100], further distributing the ion flux. As the plasma transitions to a net-recombining state, the ion flux decreases, and once electron temperatures fall sufficiently low, two- and three-body electron–ion recombination leads to a sharp reduction in the flux delivered to the target.

Power and momentum losses occur sequentially along the separatrix, defining the characteristic structure of the detached divertor. Although not observed in our diagnostic set, previous studies indicate that impurity ionisation provides the dominant power sink responsible for the initial reduction in conducted heat flux [26, 92, 133]. Figure 6.6 demonstrates that the parallel conducted heat flux decreases sharply across the 5 eV front, coinciding with the region of dominant hydrogenic ionisation. This highlights the close coupling between heat flux and particle balance: ionisation removes energy from the plasma, thereby reducing the conducted heat flux, until the plasma becomes starved of the power required to sustain further ionisation. Ion flows driven from upstream ionisation sources allow the convected heat flux to persist further downstream than the conducted component. Similarly, the static electron pressure remains elevated until electron temperatures drop to around 3 eV. In this low-temperature region, the high neutral population facilitates volumetric sinks (charge exchange and recombination processes) that dissipate both power and momentum [54]. This sequential reduction of conducted heat flux, convected heat flux, and static pressure is consistent with the expected detached divertor structure [100].

The upstream movement of the 5 eV front is moderated by the magnetic geometry, providing experimental evidence for detachment location sensitivity (DLS). As shown in Figure 6.8(d), the front shifts progressively away from the target as upstream density increases, consistent with the Lengyel model for impurity radiation [62] in which, to achieve the same drop in conducted heat flux (relative to local magnetic field) a shorter magnetic connection length is required for a greater electron density. However, the 5 eV front is observed to move slowly through the region of

maximum absolute poloidal gradient of the normalised magnetic field, $\left| \frac{1}{B} \frac{dB}{ds_{\text{pol}}} \right|$. This is in agreement with the poloidal detachment location sensitivity (DLS) study of [17] (see equation 13). In [17], it is argued that a region with high relative magnetic field gradient has a reduced detachment sensitivity. Therefore, increasing the upstream density only minimally moves the region of majority conducted heat flux reduction which, in our study, is closely related to the the 5 eV front. These results indicate that magnetic geometry plays a stabilising role in detachment, linking experimental observations to theoretical models of DLS [17, 67].

6.3.2 Changing Target Conditions

The evolution of target conditions during the density ramp, as inferred by D-MIBAS, is summarised in Figure 6.8. Although the plasma remained detached throughout the analysed period, pronounced changes in the target state are observed. These changes are found to be dominated by the onset of electron-ion recombination

Consistent with other findings (Figure 5.a of [122]), electron-ion recombination processes removed less power than ionisation and MAR processes. However, the onset of EIR led to the greatest change in target conditions. During the core density ramp, the power and ion current delivered to the target from the plasma (considering only charged species, neglecting neutral components and radiation) decreased slowly before decreasing more rapidly at around 600 ms. This coincides with the onset of electron-ion recombination (EIR). The power delivered to the lower outer target from the plasma is given by:

$$P_{\text{Target}} = \int_S (q_{\text{cond.}} + q_{\text{conv.}}) \sin \alpha \, dA, \quad (6.1)$$

where A is the elemental divertor tile area, S is the entire lower outer leg divertor, and α is the grazing angle between the magnetic field lines and the divertor tile surface plane. Similarly, the ion current is given by:

$$I_{\text{Target}} = \int_S j_{\text{sat.}} \sin \alpha \, dA. \quad (6.2)$$

Despite remaining in a detached state throughout the discharge, the strike point (the region where the separatrix intersects with the target) electron density was observed to roll-over during the core density ramp. The strike point electron temperature was observed to be below 1 eV throughout. Increased rates of radiative and three-body electron-ion recombination processes were observed to develop at around 600 ms. This additional ion sink coincides with the roll-over in strike point electron density.




Figure 6.8: D-MIBAS inferred plasma state during the core density ramp discharge #46860.

The radial profile of the target conditions were also observed to change throughout the core density ramp. The integral decay length, λ_{int} [24, 69], normalises the cross-field integral of the heat flux to its peak value, providing an effective measure of

its width. Despite detachment suggesting a flattening of the target heat flux profile (increasing λ_{int}), λ_{int} is initially observed to reduce. This was considered to be due to the core density fuelling ramp increasing the divertor electron density. Without strong recombination rates, this can increase target electron density and convected heat flux close to the target. However, once a strong recombination sink (EIR) is established, the strike point electron density rolls over and λ_{int} begins to rise again. The net result is a flattening of target profiles, indicative of a deeply detached divertor.

6.4 Summary and Conclusions

This chapter has presented the first application of D-MIBAS to experimental data, focusing on a MAST-U Super-X discharge with a core density fuelling ramp. By combining multiple diagnostic measurements with physics-based prior constraints, D-MIBAS enabled the reconstruction of key divertor plasma parameters and their evolution during detachment.

The inferred plasma states provide some qualitative alignment with established understanding of detachment physics. Substantial ionisation rates remove power from the plasma and lead to a sharp reduction in conducted heat flux across the 5 eV electron temperature front. In the detached states observed, this front resides well away from the target, leading to reduced plasma temperatures and the accumulation of neutral populations in the divertor leg. Although target recycling diminishes with reduced surface fluxes, neutral densities are maintained by molecular-activated recombination (MAR), which supplies neutrals at higher plasma temperatures than those required for purely atomic recombination.

Power and particle delivery to the target from the plasma, as well as the target temperature, are all observed to decrease during the core density ramp. The convected heat flux is observed to extend beyond the conducted heat flux, closer to the target. This is considered to be primarily due to substantial parallel flows which maintain elevated electron densities near the target. However, as additional power is removed from the plasma, the target temperatures continue to cool. Combined with the raised target electron density, these conditions promote two- and three-body electron-ion recombination, further dissipating power, momentum, and particle fluxes.

VVUQ analysis has shown the D-MIBAS inference framework to have some consistency amongst the included diagnostics, interpretive modelling, alternative analysis methods, and the current understanding of detachment. For example, it indicates

experimental support for the predicted location of the detachment front from the Detachment Location Sensitivity (DLS) model [17, 67]. However, inconsistencies exist, and the current level of uncertainty, typically $\sim 60\%$ median percentage uncertainty, limits the confidence with which new physics conclusions can be drawn. To address this, three avenues for future work are proposed:

- Extending the analysis to multiple discharges to test the robustness of priors and establish whether the inferred behaviour is general or discharge-specific.
- Inclusion of additional diagnostics such as coherence imaging spectroscopy and spectroscopic measurements between the magnetic X-point and divertor chamber.
- Improved characterisation of existing diagnostics such as reducing calibration uncertainties in spectroscopic data and learning a fixed β in the Langmuir probe model (see equation (5.7)).

Chapter 7

D-MIBAS Inference in Alternative Divertor Configurations

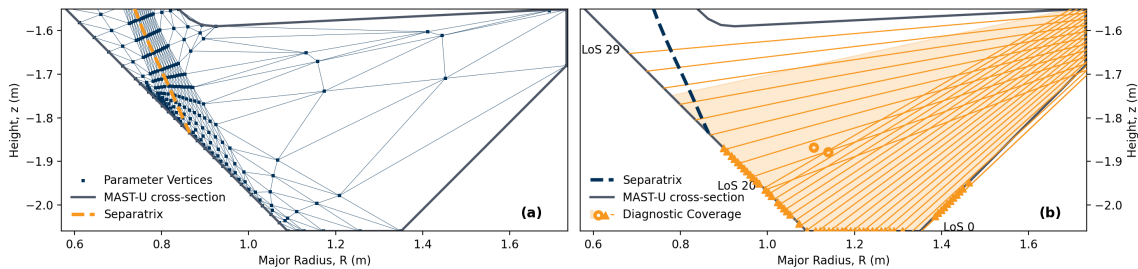
In this chapter, we explore applying D-MIBAS to alternative divertor configurations: the conventional divertor, elongated divertor, and Super-X divertor. The D-MIBAS framework is designed with a modular approach for incorporating diagnostics and prior arguments into the overall posterior distribution, making it non-machine-specific and non-scenario-specific. It requires knowledge of the magnetic equilibrium to construct the field-aligned inference mesh and forward models for the chosen diagnostics. However, the priors developed in this thesis use hyperparameters learned from multiple SOLPS-ITER simulations, which are fixed for simplicity, as summarised in Chapter 5 and used in Chapter 6. This chapter explores D-MIBAS’s performance in inferring alternative divertor configurations. In doing so, it provides insight into whether this version of D-MIBAS, with fixed hyperparameters, is sufficiently generalisable or whether D-MIBAS necessitates an alternative treatment of the hyperparameters.

7.1 Discharge #46895

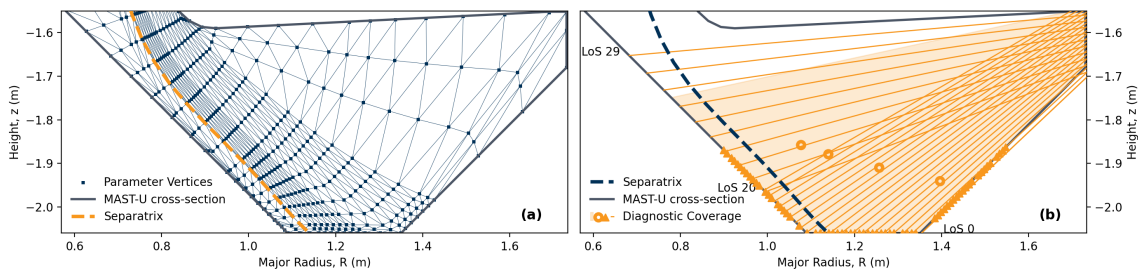
MAST-U discharge #46895 was performed as a controlled geometry sweep of the outer-leg separatrix. It has been extensively studied in [122] which contains interpretive SOLPS-ITER modelling. During the discharge, the divertor configuration transitioned from a conventional divertor (CD) at approximately 400 ms, to an elongated divertor (ED) at approximately 600 ms, to a Super-X divertor (SXD) at approximately 800 ms. Throughout this sweep, the core density and power were held constant at approximately 30% of the Greenwald limit ($n_{e,\text{sep}} \approx 0.8 \times 10^{18} \text{ m}^{-3}$) and $P_{\text{SOL}} \approx 1.0 \text{ MW}$, maintained by neutral beam heating and active deuterium fuelling. This corresponds to upstream conditions akin to those at approximately 480 ms in

discharge #46860 studied in Chapter 6. As the strike point was swept radially outward, the major radius of the strike point increased from 0.8 to 1.4 m, the magnetic connection length from the upstream midplane to the target increased from 13 to 19 m, and the total flux expansion increased from $F_R = 1.2$ to 2.2.

The diagnostic coverage for the conventional divertor configuration is extremely limited as shown in Figure 7.1a. The separatrix is far from the Thomson scattering point, only intersects with five spectroscopy chords, is minimally viewed by the multi-wavelength imaging camera, and no functioning Langmuir probe data was available in the strike point vicinity. Consequently, this configuration relies heavily on the physics-based priors developed throughout this thesis.



(a) Conventional divertor (400 ms).



(b) Elongated divertor (600 ms).

Figure 7.1: Inference mesh (a) and diagnostic coverage (b) for experimental use. Diagnostic coverage describes: spectroscopy - solid lines; multi-wavelength imaging - shaded region; Thomson scattering - hollow circles; and Langmuir probes - filled triangles.

7.1.1 Interpretive Modelling

D-MIBAS provides the opportunity for direct comparison between experiment and modelling. Interpretive SOLPS-ITER modelling of the divertor configurations found in discharge #46895 has been performed and presented in [122]. Data has been made

available at [113]. The interpretive SOLPS-ITER modelling results are repeated here, in Figures 7.3 and 7.6, for direct comparison to the fields inferred by D-MIBAS.

7.2 Results

The results of D-MIBAS are presented in Figure 7.2. The results suggest that ED and SXD configurations were partially detached. Dominant hydrogenic ionisation rate is observed to give-way to dominant MAR rates in advance of the target. The target electron temperature is below 5 eV at the strike point for all divertor configurations. This is in agreement with the interpretive SOLPS-ITER modelling. The same is true for the CD configuration, however, the 5 eV front and transition from dominant hydrogenic ionisation to dominant MAR rate is observed to occur within two cm of the target.

Consistency is found between the detachment location sensitivity (DLS) model, and the location of the detachment front, as summarised by the 5 eV front, across all three divertor configurations. As the strike point is swept radially outward, the 5 eV front is observed to move from a major radius of 0.8 m to 1 m. This corresponded to the region of greatest absolute poloidal gradient of the magnetic field, $\left| \frac{1}{B} \frac{dB}{ds_{\text{pol}}} \right|$. This follows the same observation found in Chapter 6. A similar position for the 5 eV front is also found in the SOLPS-ITER simulations.

Inconsistencies are observed in the inferred electron density profiles, most plausibly due to limited diagnostic coverage in the conventional (CD) and elongated (ED) divertor configurations. In these cases, D-MIBAS infers significantly higher electron densities than both the Super-X divertor (SXD) configuration and the corresponding SOLPS-ITER predictions. While SOLPS-ITER simulations consistently predict a peak density of $3\text{-}4 \times 10^{19} \text{m}^{-3}$, located a few centimetres from the strike point, D-MIBAS infers densities approaching $7 \times 10^{19} \text{m}^{-3}$ along the separatrix in the CD and ED cases, peaking near the tail end of the hydrogenic ionisation region. Although such values are physically plausible, the inference of the CD and ED cases was found to have large inferred spectroscopy calibration (c_i) uncertainties. The limited availability of alternative diagnostic information (such as Thomson scattering points near the high-emissivity region) reduced the ability to constrain the calibration robustly thus permitting high uncertainties in inferred electron densities (exceeding 50%). To establish whether the increase CD and ED electron density is significant, either improved diagnostic coverage in the CD and ED configurations, or fixed and well-characterised spectroscopy calibration values, is required.

Establishing consistent hyperparameter priors across different divertor configurations is shown in this chapter to be plausible, but not always desirable. This challenge is clearest in the case of the atomic neutral density. Lengthscales must be short enough to capture the sharp drop within the hydrogenic ionisation region, yet long enough to represent the broader neutral distribution. The rational quadratic kernel partly accommodates this by allowing smoothness across multiple scales (see Section 5.3.2). However, as seen when comparing the inferred results in Figure 7.2 with the interpretive modelling predictions in Figure 7.3, the inferred structure differs markedly downstream of the ionisation region. Improving neutral density inference will therefore require either additional diagnostic constraints or refined hyperparameter strategies. The latter could involve assigning different expected lengthscales between divertor configurations, or adopting a non-stationary approach in which lengthscales vary with spatial location (e.g. downstream of the ionisation region). Such higher-fidelity priors, however, come at the cost of greater model complexity and reduced transparency in the interpretation of results.

Crucially, the spatially dependent priors used in this work enabled D-MIBAS to generally distinguish hydrogenic ionisation regions from MAR-dominated regions across all three divertor configurations. Comparison of the D-MIBAS results (Figure 7.2) with the modelling predictions (Figure 7.3) shows the expected structure: a narrow region dominated by ionisation, followed by a broader MAR region extending between the target and the ionisation front. However, the MAR rates inferred by D-MIBAS are considerably lower than those predicted by SOLPS-ITER, both in peak magnitude and spatial extent. The higher MAR rates predicted by SOLPS-ITER lead to increased power, particle, and momentum losses, which in turn enabled dominant EIR rates and a broadened electron density profile at the target. Such behaviour was not observed in the D-MIBAS inference, which instead remained consistent with expectations for comparable Greenwald fractions (30%) in discharge #46860 (Chapter 6).

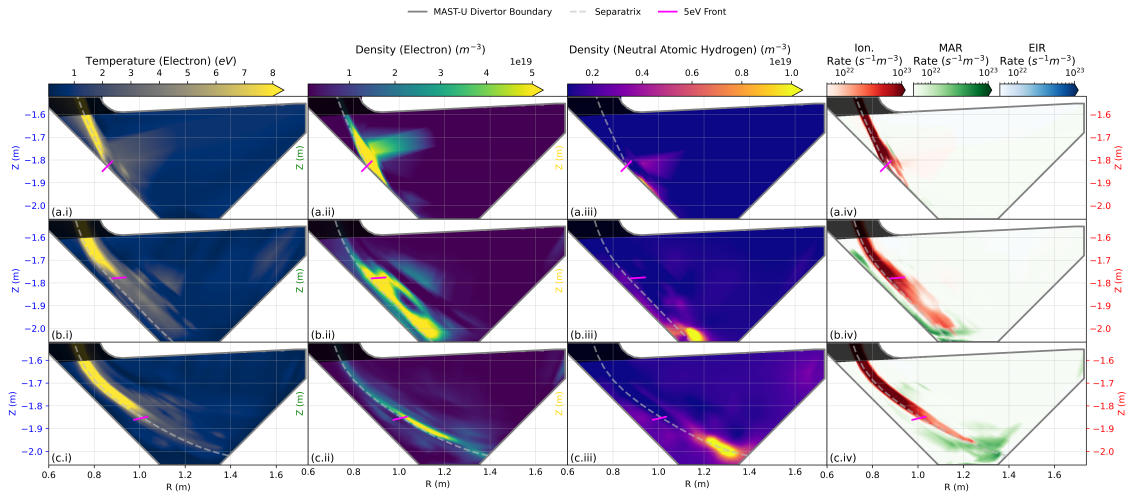


Figure 7.2: D-MIBAS inferred fields for discharge #46895. Top row (a): Conventional divertor (CD), 400 ms. Middle row (b): Elongated divertor (ED), 600 ms. Bottom row (c): Super-X divertor (SXD), 800 ms.

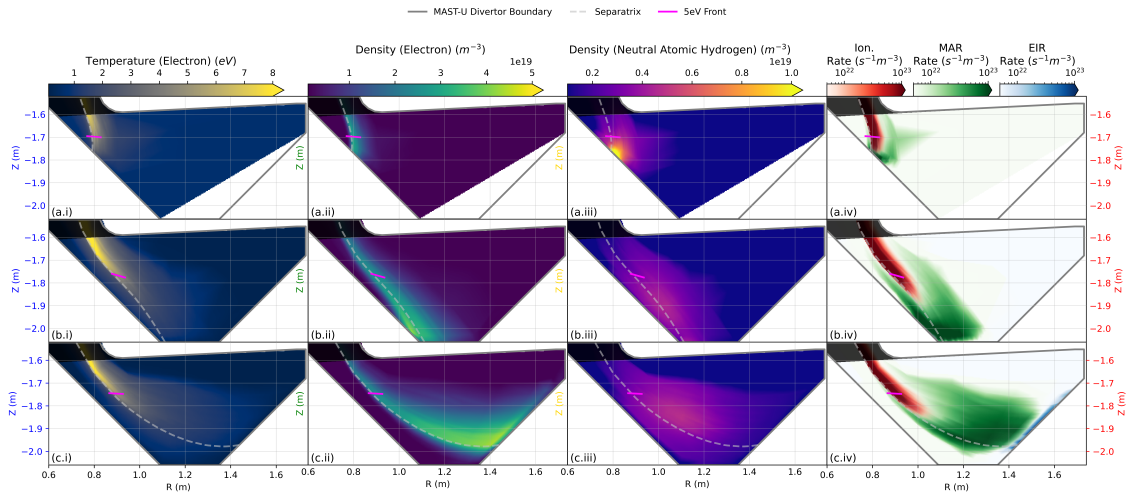


Figure 7.3: Interpretive SOLPS-ITER modelling of CD, ED, and SXD MAST-U divertor configurations presented as part of [122] and available at [113]. Top row (a): Conventional divertor (CD). Middle row (b): Elongated divertor (ED). Bottom row (c): Super-X divertor (SXD).

7.3 Discussion

7.3.1 Diagnostic Agreement

Inference in the conventional (CD) and elongated (ED) divertor configurations is considerably more challenging than in the Super-X divertor (SXD), primarily due to

limited diagnostic coverage. Figures 7.4a and 7.4b show the diagnostic agreement for the CD and ED cases, respectively.

The camera data tends to dominate the combined likelihood. This arises because each camera pixel contributes an individual data point, creating an abundance of constraints compared to other diagnostics. However, the calibration values for the spectroscopic data are unknown, meaning that complementary diagnostics are essential for constraining and learning these calibrations. When such complementary diagnostics are scarce, the inference may align with the camera data but not with the other diagnostics.

Evidence of this imbalance is seen in Figures 7.4a and 7.4b. In these cases, plasma states that would be in closer agreement with non-camera diagnostics are either not explored by the HMC sampling algorithm or are penalised as too inconsistent with the camera model. This signals underlying disagreement between diagnostic models. A particular issue is found in the Balmer brightness model at low emissivities. Disagreement is clear in DMS lines of sight 0–5, which, as shown in Figures 7.1a and 7.1b, do not intersect regions of substantial emission. This is plausibly due to missing background terms in the DMS model (see equation (5.11), as well as vignetting-related calibration challenges in the camera.

This highlights the inherent complexity of relying solely on camera data for integrated data analysis. While progress can be made by learning and reducing current uncertainties (e.g. calibration values and rates in collisional-radiative models), baseline issues such as camera degradation during operation remain a plausible source of error. Consequently, the most robust approach combines camera data with complementary diagnostics and physics-based information, enabling cross-calibration and consistency during operation. Such integration of heterogeneous diagnostics with physics-informed constraints is precisely what D-MIBAS is designed to achieve.

7.3.2 Derived Fields

The Fulcher band emissivity, parallel conducted heat flux, parallel convected heat flux, and static electron pressure corresponding to the results shown in Figure 7.2 are presented in Figure 7.5. As highlighted in Chapter 6, the Fulcher band provides a convenient experimental observation, while the parallel conducted heat flux, parallel convected heat flux, and static electron pressure provide insight into the removal of power, particles, and momentum from the plasma.

The derived fields allow a clear view into the impact and benefits of the various divertor configurations. The observations support the extensive analysis of [122]. In each of the CD, ED, and SXD divertor configurations, strong ionisation and the associated power losses lead to a suppression of the parallel conducted heat flux before the plasma reaches the target and corresponding reduced electron temperatures. The parallel convected heat flux extends further than the conducted component, reflecting enhanced particle fluxes. However, in the ED and particularly in the SXD, these fluxes subsequently diminish as molecular-activated recombination increases, resulting in a reduced ion flux to the target. The increased space and divertor baffling in the ED and SXD configurations accommodates extended regions of high neutral densities. This contributes to a decrease in the static electron pressure through momentum withdrawal due to charge exchange and elastic collisions.

The derived fields provide a mechanism not only for validating interpretive modelling, but also for identifying which model assumptions and parameters may require refinement to better reproduce experimental observations. As discussed in Section 2.3.2, SOLPS-ITER itself contains various assumptions and uncertainties. The D-MIBAS results (Figure 7.5) show qualitative similarities with the interpretive SOLPS-ITER modelling (Figure 7.6), though clear differences remain in the scrape-off layer width and peak magnitudes. In particular, the experimentally inferred scrape-off layer appears narrower than predicted by SOLPS-ITER, with correspondingly elevated peak heat fluxes and reduced peak Fulcher band emissivities. If the D-MIBAS results are assumed to be accurate, these discrepancies can guide modelling refinements, including adjustments to molecular reaction rates (affecting Fulcher band emission), cross-field transport and drift physics (affecting scrape-off layer width), and scrape-off layer power balance and impurity ionisation (affecting peak heat fluxes). However, this also highlights a fundamental challenge in validation: both D-MIBAS and SOLPS-ITER contain substantial uncertainties and modelling assumptions, and neither can be assumed to represent the true plasma state exactly. Consequently, robust quantitative validation requires analysis across substantially

larger discharge databases, comparison against independent diagnostic techniques and alternative modelling approaches, and continued development of divertor physics understanding to identify systematic differences between inference and simulation.

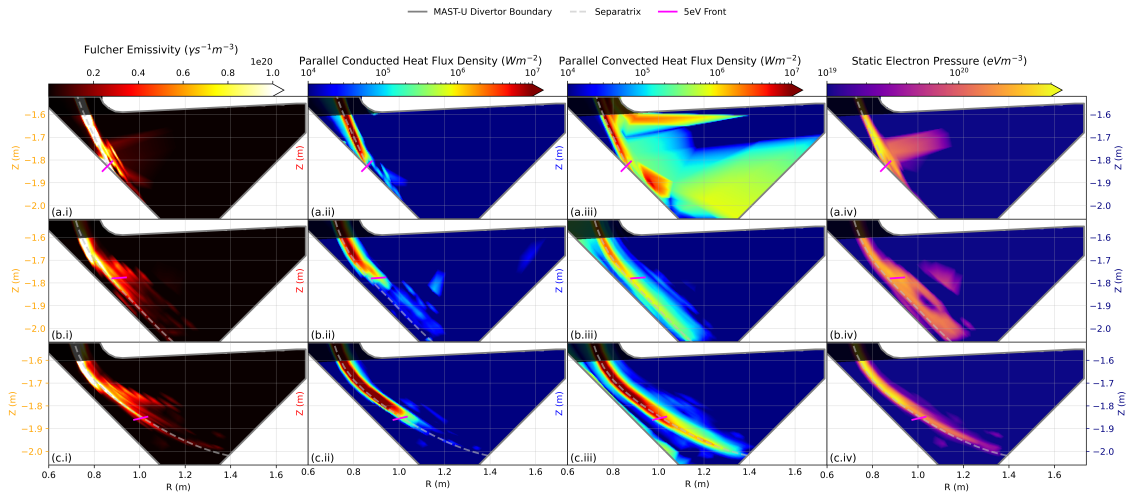


Figure 7.5: Derived fields from the inferred plasma states throughout discharge #46895.

Top row (a): Conventional divertor (CD).
 Middle row (b): Elongated divertor (ED).
 Bottom row (c): Super-X divertor (SXD).

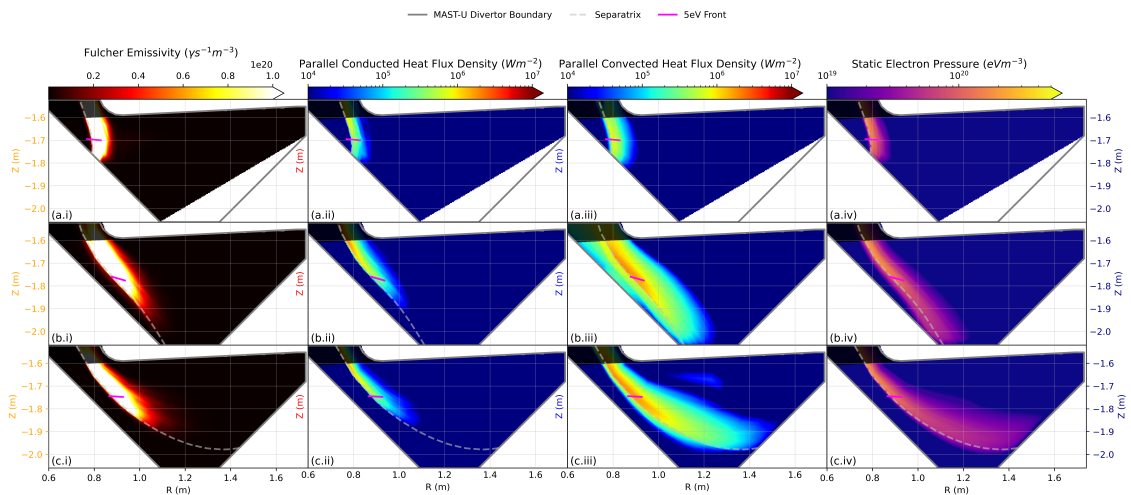


Figure 7.6: Interpretive SOLPS-ITER modelling of CD, ED, and SXD MAST-U divertor configurations presented as part of [122] and available at [113].

Top row (a): Conventional divertor (CD).
 Middle row (b): Elongated divertor (ED).
 Bottom row (c): Super-X divertor (SXD).

7.4 Summary

This chapter has applied D-MIBAS to a discharge that accessed three distinct divertor geometries (conventional, elongated, and Super-X) while maintaining the same fixed hyperparameters inferred from earlier analysis (Chapter 5). The results demonstrate that the framework is sufficiently flexible to capture expected detachment behaviour across all three cases, with consistent identification of ionisation- and MAR-dominated regions. There was loose agreement between the results and interpretive SOLPS-ITER modelling and divertor-physics-based expectations.

The differences observed between SOLPS-ITER predictions and D-MIBAS inferences guide where additional effort is required. In the conventional and elongated divertors, electron densities were inferred to be higher than modelling predictions. However, large uncertainties persisted due to the limited diagnostic coverage which was unable to constrain the inferred calibration values. The inferred neutral density highlighted the limitations of fixed hyperparameters. The selected kernel was designed to capture multi-scale structures; however, it is extremely challenging to provide sufficient flexibility in regions of sharp transition (i.e. near the detachment front) while maintaining broader, smooth structures elsewhere. These findings underscore the dual importance of multiple, complimentary diagnostics and carefully chosen priors in the D-MIBAS framework.

As with the results of Chapter 6, the substantial uncertainties and limited number of discharges analysed prevent detailed quantitative comparisons and conclusions. Repeat analysis in future work will support uncertainty reduction (e.g. by learning and fixing camera calibrations rather than allowing them to be free parameters in each inference) and will ascertain where clear systematic disagreements occur.

Chapter 8

Conclusion and Outlook

In this thesis, a novel Bayesian framework for integrated data analysis of tokamak divertor plasmas has been developed and deployed to infer two-dimensional plasma states in terms of electron temperature, electron density, and neutral atomic hydrogen density. This work contributes to the broader challenge of controlling heat and particle exhaust in tokamaks by providing a generic methodology for investigating this complex and crucial region of a tokamak. It offers deep experimental insight, enables direct comparison between interpretive models and experimental results, and acts as a platform for informing the diagnostic suite and optimising heat and particle exhaust of future fusion power plants.

The central contribution of this work is the development of a mesh-based inference framework, the Divertor Multi-Instrument Bayesian Analysis System (D-MIBAS), which unifies multiple diverse diagnostics and physics-informed arguments into a single coherent inference problem. In contrast to existing cell-based inference approaches, which preprocess diagnostic data and perform independent inferences at discrete spatial locations, D-MIBAS models the plasma state over a continuous two-dimensional spatial mesh. This spatial coupling is what facilitates the incorporation of diagnostics probing different plasma regions and supports the use of physics-based prior arguments, such as the expected monotonicity of electron temperature, static electron pressure, and conducted heat flux profiles along magnetic field lines. These advancements enable D-MIBAS to interpret ill-posed individual diagnostic signals, resulting in accurate, low-uncertainty inference.

Key achievements include:

- The development of a mesh-based inference framework capable of handling thousands of interrelated parameters (represent multiple physical quantities, fields, distributed across the mesh) using Hamiltonian Markov Chain Monte Carlo.

- The inclusion of molecular emission processes to the Balmer line emission model in a concise, tractable manner.
- The probabilistic incorporation of physics-based priors that guide the inference without over-constraining it, thereby maintaining robustness against local inconsistencies such as mesh misalignment or unmodelled physics.
- The data-driven *learning* of hyperparameters within these physics-based priors, achieved through a physics-informed probabilistic machine learning framework.
- Demonstration of robust performance across Super-X and elongated divertor configurations using experimental MAST-U data.
- Demonstration of performance in the conventional divertor configuration with limited experimental diagnostic coverage.
- Validation of inferred results against both independent diagnostic measurements and interpretive SOLPS-ITER simulations.

Remaining challenges and future work include:

- Extending validation efforts to high-power H-mode scenarios, where additional physics (e.g. small local perturbations to the plasma state due to filamentary transport) may require new modelling assumptions such as the consideration of three-dimensional effects [51].
- Extending the domain of D-MIBAS to observe plasma characteristics upstream of the divertor.
- Inclusion of additional diagnostics that could provide additional physics quantities. These additional diagnostics can be directed at different regions of the plasma such as an X-point camera, or make better use of currently included diagnostics such as investigating Stark broadening profiles in the spectroscopy data [73] and coherence imaging spectroscopy [21, 71, 72] channels in filtered imaging data. This will reduce assumptions made when deriving power and ion flux quantities from the inferred parameters.
- Applying D-MIBAS to promising alternative divertor configurations such as the snowflake or X-point target [16].
- Applying D-MIBAS to alternative tokamaks to test its generalisability across machine geometries and diagnostic suites.

- The learning of hyperparameters within the inference itself to enable more appropriate choices to be made for the varying divertor states.

8.1 On Furthering Exhaust Understanding

The divertor is one of the most complex plasma regions in a tokamak. It hosts a wide variety of species including neutral atoms, neutral molecules, hydrogen ions, and impurity ions at a large range of different charge states. These can all undergo various processes and impact the divertor dynamics. The divertor is inherently two- or even three-dimensional in nature, with charged species exhibiting anisotropic transport along magnetic field lines, while neutral species are largely isotropic. This complexity prevents simplified, one-dimensional geometric descriptions from fully capturing the physics of the divertor.

This work addresses these challenges in several novel ways:

- D-MIBAS can provide a comprehensive description of the divertor plasma state to support the interpretation of detachment dynamics. It reconstructs two-dimensional plasma fields by inferring key quantities, from which additional quantities (such as heat and particle fluxes) can be derived under reasonable assumptions. Currently inaccessible quantities include impurity and neutral temperatures, densities, and velocities. However, these could be incorporated if suitable diagnostics become available, or inferred indirectly through physics-based arguments that link them to the reconstructed parameters, albeit with increased uncertainty.
- D-MIBAS incorporates expected physics to constrain inferred uncertainties. Its mesh-based approach provides spatial links that promote physical consistency, for example by encoding isotropic correlations for neutral species and anisotropic correlations for charged species.
- D-MIBAS integrates multiple diagnostics that provide complementary information, thereby constraining inferred uncertainties. Its mesh-based approach enables the joint modelling of diagnostics that cover different regions of the plasma. While individual diagnostics may be insufficient on their own, together they constrain each other and provide physically interpretable results.

- D-MIBAS’s probabilistic structure provides inherent uncertainty quantification. This enables explicit assessment of consistency between diagnostics as well as between experiment and modelling results.
- Applied to MAST-U data, D-MIBAS revealed power and momentum loss patterns aligned with established detachment theory. It provided quantitative insight into the spatial structure of the detachment process. It also provided experimental relationships between experimentally observable quantities, such as Fulcher band emission, and key plasma quantities such as the 5 eV electron temperature front.

8.2 On Advising Future Integrated Data Analysis Efforts

Measurement of Balmer line brightness remains a valuable diagnostic for divertor plasmas due to its broad coverage and non-invasive nature. This work confirms that high- n to low- n Balmer line ratios can reliably identify cold (<1 eV electron temperature), recombining regions from other plasma states. However, outside this regime, Balmer line signals alone are insufficient to uniquely determine plasma states. This is found to be due to the similar signatures from different atomic and molecular processes generating an ill-posed problem. Thus, when significant densities of molecules are present, an advance integrated data analysis approach is required to interpret Balmer line information.

D-MIBAS overcomes this limitation by embedding Balmer line data within a broader inference problem. This is facilitated by a mesh-based approach that supports the inclusions of: multiple diagnostics covering different regions; and physics-based prior arguments (e.g. the monotonicity of T_e and smoothness of n_D) that implicitly resolve the ill-posed nature of the problem.

Key methodological advances include:

- Gaussian process priors: Used to encode spatial correlations and introduce physics-based assumptions on the smoothness of plasma fields. A *white-noise* term was added directly to the field to maintain numerical stability and to prevent the priors from becoming over-constraining. These priors also provide an approximation to the *mass-matrix* required for efficient Hamiltonian Monte Carlo sampling, which is crucial for scalability to large parameter sets.

- Cauchy prior probability distributions: Used for physics-based assumptions that cannot be guaranteed with certainty (e.g. due to mesh misalignment or incomplete physics). The heavy-tailed Cauchy form is tolerant to outliers, ensuring that D-MIBAS can incorporate physics-based priors without strongly biasing the inference when confronted with contradictory information.
- Hyperparameter selection: A data-driven, empirical Bayesian approach was taken to hyperparameter selection. Multiple SOLPS-ITER simulations across a diverse range of divertor scenarios were used to *learn* values for the hyperparameters. The resulting hyperparameters were generalisable across a multiple divertor scenarios meaning that they could be fixed within the inference. In future implementations, these hyperparameters could be considered as nuisance parameters and inferred alongside the plasma state itself.

The modular architecture of D-MIBAS makes it well-suited for future integrated data analysis efforts. A major challenge lies in the very large number of parameters (typically multiple thousands) required to model multiple diagnostics. This large *parameter space* must be explored to characterise the posterior distribution and quantify uncertainties. Synthetic verification has confirmed that Hamiltonian Markov Chain Monte Carlo (HMC) can accurately handle parameter sets of this scale, provided that inter-parameter correlations are captured through HMC's *mass matrix*. This robust uncertainty quantification enables systematic testing of diagnostic and model assumptions.

8.3 On D-MIBAS's Role in Fusion Power Plants

As the fusion community transitions towards the design and realisation of power-plant tokamaks, the demand for reliable, interpretable, and uncertainty-aware diagnostic systems becomes critical. Divertor plasmas in power plant conditions will operate under constraints that challenge both measurement and modelling. Power plant devices will encounter significant reductions in diagnostic coverage due to harsh operating environments (e.g., neutron flux, limited port access, and intense thermal loads). D-MIBAS is not expected to be directly used in fusion power plants.

Instead, D-MIBAS presents a step toward fusion power plants by supporting the validation of plasma-edge modelling and enhances our understanding of exhaust. By reconstructing the divertor plasma state from experimental data with quantified uncertainties, D-MIBAS facilitates direct comparisons with interpretive models such as

SOLPS-ITER. This provides a rigorous approach to simulation validation with today's well diagnosed research tokamaks and is essential for reducing uncertainty when extrapolating divertor performance to future power plants. In doing so, it contributes to the development of resilient exhaust systems, bringing the prospect of abundant energy from fusion power plants closer to reality.

Bibliography

- [1] International Atomic Energy Agency. *Evaluated Nuclear Data File (ENDF)*. <https://www-nds.iaea.org/exfor/endl.htm>. Accessed: 2024-08-30. 2024.
- [2] M Agostini et al. “Neutral density estimation in the ASDEX upgrade divertor from deuterium emissivity measurements during detachment and shoulder formation”. In: *Plasma Physics and Controlled Fusion* 61.11 (Sept. 2019), p. 115001. DOI: 10.1088/1361-6587/ab4122.
- [3] S.L. Allen et al. “First measurements of electron temperature and density with divertor Thomson scattering in radiative divertor discharges on DIII-D”. In: *Journal of Nuclear Materials* 241-243 (1997), pp. 595–601. ISSN: 0022-3115. DOI: [https://doi.org/10.1016/S0022-3115\(97\)80106-7](https://doi.org/10.1016/S0022-3115(97)80106-7).
- [4] M-H. Aumeunier et al. “Development of inverse methods for infrared thermography in fusion devices”. In: *Nuclear Materials and Energy* 33 (2022), p. 101231. ISSN: 2352-1791. DOI: <https://doi.org/10.1016/j.nme.2022.101231>.
- [5] et. al. B. Lipschultz. “Divertor Physics Research on Alcator C-Mod”. In: *Fusion Science and Technology* 51.3 (2007), pp. 369–389. DOI: 10.13182/FST07-A1428.
- [6] RC Bergmayr, D Wunderlich, and U Fantz. “Molecular data needs for advanced collisional-radiative modelling for hydrogen plasmas”. In: *The European Physical Journal D* 77.7 (2023), p. 136.
- [7] M. Bernert et al. “The X-Point radiating regime at ASDEX Upgrade and TCV”. In: *Nuclear Materials and Energy* 34 (2023), p. 101376. ISSN: 2352-1791. DOI: <https://doi.org/10.1016/j.nme.2023.101376>.
- [8] Thomas van den Biggelaar. “Determination of a 2D electron density and electron temperature profile in the MAST-U divertor using spectral line ratios analysis”. MSc Thesis. Eindhoven University of Technology, 2022.
- [9] C Bowman et al. “Development and simulation of multi-diagnostic Bayesian analysis for 2D inference of divertor plasma characteristics”. In: *Plasma Physics and Controlled Fusion* 62.4 (Feb. 2020), p. 045014. DOI: 10.1088/1361-6587/ab759b.
- [10] Chris Bowman. “Applications of Bayesian Probability Theory in Fusion Data Analysis”. Doctor of Philosophy. University of York, 2016.
- [11] J Bryant et al. “Impact of Yacora evaluated molecular effective rate coefficients on detached SOLPS-ITER simulations”. In: *Nuclear Fusion* 65.3 (2025), p. 036025.
- [12] J.M. Canik et al. “Modeling of detachment experiments at DIII-D”. In: *Journal of Nuclear Materials* 463 (2015). PLASMA-SURFACE INTERACTIONS 21, pp. 569–572. ISSN: 0022-3115. DOI: <https://doi.org/10.1016/j.jnucmat.2014.11.077>.
- [13] M Cavedon et al. “Experimental investigation of L-and H-mode detachment via the divertor Thomson scattering at ASDEX Upgrade”. In: *Nuclear Fusion* 62.6 (2022), p. 066027.

- [14] J. G. Clark, M. D. Bowden, and R. Scannell. “Low temperature Thomson scattering on MAST-U”. In: *Review of Scientific Instruments* 92.4 (Apr. 2021), p. 043545. ISSN: 0034-6748. DOI: 10.1063/5.0043813.
- [15] J. G. Clark et al. “First divertor Thomson scattering measurements on MAST-U”. In: *Review of Scientific Instruments* 93.10 (2022), p. 103534. DOI: 10.1063/5.0101635.
- [16] C Cowley et al. “Novel SOLPS-ITER simulations of X-point target and snowflake divertors”. In: *Plasma Physics and Controlled Fusion* 65.3 (Feb. 2023), p. 035011. DOI: 10.1088/1361-6587/acb4ba.
- [17] C. Cowley et al. “Optimizing detachment control using the magnetic configuration of divertors”. In: *Nuclear Fusion* 62.8 (July 2022), p. 086046. DOI: 10.1088/1741-4326/ac7a4c.
- [18] Y. Damizia et al. *First Ion Temperature Measurements in the MAST-U Divertor via Retarding Field Energy Analyzer*. 2024.
- [19] G De Temmerman et al. “Thermographic study of heat load asymmetries during MAST L-mode discharges”. In: *Plasma Physics and Controlled Fusion* 52.9 (2010), p. 095005.
- [20] AJH Donné, AE Costley, and AW Morris. “Diagnostics for plasma control on DEMO: challenges of implementation”. In: *Nuclear Fusion* 52.7 (2012), p. 074015.
- [21] RS Doyle et al. “Development and calibration of a multi-delay coherence imaging diagnostic on the MAST-U tokamak”. In: *Review of Scientific Instruments* 95.5 (2024).
- [22] Benjamin Daniel Dudson and Jarrod Leddy. “Hermes: global plasma edge fluid turbulence simulations”. In: *Plasma Physics and Controlled Fusion* 59.5 (2017), p. 054010.
- [23] T. Eich et al. “Inter-ELM Power Decay Length for JET and ASDEX Upgrade: Measurement and Comparison with Heuristic Drift-Based Model”. In: *Phys. Rev. Lett.* 107 (21 Nov. 2011), p. 215001. DOI: 10.1103/PhysRevLett.107.215001.
- [24] T. Eich et al. “Scaling of the tokamak near the scrape-off layer H-mode power width and implications for ITER”. In: *Nuclear Fusion* 53.9 (Aug. 2013), p. 093031. DOI: 10.1088/0029-5515/53/9/093031.
- [25] Fabio Federici et al. “Design and implementation of a prototype infrared video bolometer (IRVB) in MAST Upgrade”. In: *Review of Scientific Instruments* 94.3 (2023).
- [26] Fabio Federici et al. “Design and implementation of a prototype infrared video bolometer (IRVB) in MAST Upgrade”. In: *Review of Scientific Instruments* 94.3 (2023).
- [27] Fabio Federici et al. “Evolution of radiation profiles in a strongly baffled divertor on MAST Upgrade”. In: *Nuclear Materials and Energy* 43 (2025), p. 101940. ISSN: 2352-1791. DOI: <https://doi.org/10.1016/j.nme.2025.101940>.
- [28] X. Feng et al. “Development of an 11-channel multi wavelength imaging diagnostic for divertor plasmas in MAST Upgrade”. In: *Review of Scientific Instruments* 92.6 (June 2021), p. 063510.
- [29] Y. Feng et al. “Monte-Carlo fluid approaches to detached plasmas in non-axisymmetric divertor configurations”. In: *Plasma Physics and Controlled Fusion* 59.3 (2017), p. 034006. ISSN: 0741-3335.
- [30] A Fil et al. “Separating the roles of magnetic topology and neutral trapping in modifying the detachment threshold for TCv”. In: *Plasma Physics and Controlled Fusion* 62.3 (Jan. 2020), p. 035008. DOI: 10.1088/1361-6587/ab69bb.
- [31] A. Fil et al. “Comparison between MAST-U conventional and Super-X configurations through SOLPS-ITER modelling”. In: *Nuclear Fusion* 62.9 (Aug. 2022), p. 096026. DOI: 10.1088/1741-4326/ac81d8.
- [32] R Fischer, A Dinklage, and E Pasch. “Bayesian modelling of fusion diagnostics”. In: *Plasma Physics and Controlled Fusion* 45.7 (May 2003), p. 1095. DOI: 10.1088/0741-3335/45/7/304.

-
- [33] R Fischer et al. “Thomson scattering analysis with the Bayesian probability theory”. In: *Plasma Physics and Controlled Fusion* 44.8 (July 2002), p. 1501. DOI: 10.1088/0741-3335/44/8/306.
- [34] J. Gaspar et al. “Emissivity measurement of tungsten plasma facing components of the WEST tokamak”. In: *Fusion Engineering and Design* 149 (2019), p. 111328. ISSN: 0920-3796. DOI: <https://doi.org/10.1016/j.fusengdes.2019.111328>.
- [35] W.J. Goedheer. *The influence of diffusion and of reabsorption of radiation on the particle and energy balance of an infinitely long quasi-cylindrical discharge in hydrogen gas*. 1978.
- [36] Daniel Greenhouse et al. *Two-dimensional inference of divertor plasma characteristics: advancements to a multi-instrument Bayesian analysis system*. 2024.
- [37] M. Greenwald et al. “A new look at density limits in tokamaks”. In: *Nuclear Fusion* 28.12 (Dec. 1988), p. 2199. DOI: 10.1088/0029-5515/28/12/009.
- [38] J.R. Harrison et al. “Overview of new MAST physics in anticipation of first results from MAST Upgrade”. In: *Nuclear Fusion* 59.11 (June 2019), p. 112011. DOI: 10.1088/1741-4326/ab121c.
- [39] James R Harrison et al. “Benefits of the Super-X divertor configuration for scenario integration on MAST Upgrade”. In: *Plasma Physics and Controlled Fusion* 66.6 (2024), p. 065019.
- [40] W. K. Hastings. “Monte Carlo sampling methods using Markov chains and their applications”. In: *Biometrika* 57.1 (Apr. 1970), pp. 97–109. ISSN: 0006-3444. DOI: 10.1093/biomet/57.1.97.
- [41] J Hawke et al. “Outline of optical design and viewing geometry for divertor Thomson scattering on MAST upgrade”. In: *Journal of Instrumentation* 8.11 (Nov. 2013), p. C11010. DOI: 10.1088/1748-0221/8/11/C11010.
- [42] S. S. Henderson et al. “Determination of volumetric plasma parameters from spectroscopic N II and N III line ratio measurements in the ASDEX Upgrade divertor”. In: *Nuclear Fusion* 58.1 (2018), p. 016047. ISSN: 0029-5515.
- [43] Nicholas J. Higham. “Computing a nearest symmetric positive semidefinite matrix”. In: *Linear Algebra and its Applications* 103 (1988), pp. 103–118. ISSN: 0024-3795. DOI: [https://doi.org/10.1016/0024-3795\(88\)90223-6](https://doi.org/10.1016/0024-3795(88)90223-6).
- [44] R. K. Janev and D. Reiter. *Isotope effects in molecule assisted recombination and dissociation in divertor plasmas*. Jülich report - JUEL 4411. englisch. Jülich: Forschungszentrum Jülich GmbH, 2018, 1 Online-Ressource (37 Seiten).
- [45] Harold Jeffreys. “An invariant form for the prior probability in estimation problems”. In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 186.1007 (1946), pp. 453–461. DOI: 10.1098/rspa.1946.0056.
- [46] Yuzhong Jin et al. “Progress in the engineering design of ITER Langmuir probe”. In: *Fusion Engineering and Design* 184 (2022), p. 113306. ISSN: 0920-3796. DOI: <https://doi.org/10.1016/j.fusengdes.2022.113306>.
- [47] J Karhunen et al. “Experimental distinction of the molecularly induced Balmer emission contribution and its application for inferring molecular divertor density with 2D filtered camera measurements during detachment in JET L-mode plasmas”. In: *Plasma Physics and Controlled Fusion* 64.7 (2022), p. 075001.
- [48] J. Karhunen et al. “Estimation of 2D distributions of electron density and temperature in the JET divertor from tomographic reconstructions of deuterium Balmer line emission”. In: *Nuclear Materials and Energy* 25 (2020), p. 100831. ISSN: 2352-1791. DOI: <https://doi.org/10.1016/j.nme.2020.100831>.

- [49] J. Karhunen et al. “Spectroscopic camera analysis of the roles of molecularly assisted reaction chains during detachment in JET L-mode plasmas”. In: *Nuclear Materials and Energy* 34 (2023), p. 101314. ISSN: 2352-1791. DOI: <https://doi.org/10.1016/j.nme.2022.101314>.
- [50] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017.
- [51] A. Kirk et al. “Spatial and Temporal Structure of Edge-Localized Modes”. In: *Phys. Rev. Lett.* 92 (24 June 2004), p. 245002. DOI: 10.1103/PhysRevLett.92.245002.
- [52] B Kool et al. “First demonstration of Super-X divertor exhaust control for transient heat load management in compact fusion reactors”. In: *arXiv preprint arXiv:2407.07784* (2024).
- [53] Vladislav Kotov, Detlev Reiter, and Andrey S Kukushkin. *Numerical study of the ITER divertor plasma with the B2-EIRENE code package*. Tech. rep. Forschungszentrum Juelich (Germany). Inst. fuer Energieforschung (IEF), 2007.
- [54] S. I. Krasheninnikov and A. S. Kukushkin. “Physics of ultimate detachment of a tokamak divertor plasma”. In: *Journal of Plasma Physics* 83.5 (2017), p. 155830501. ISSN: 0022-3778. DOI: 10.1017/S0022377817000654.
- [55] A. S. Kukushkin et al. “Role of molecular effects in divertor plasma recombination”. In: *Nuclear Materials and Energy* 12 (2017), pp. 984–988. ISSN: 2352-1791. DOI: 10.1016/j.nme.2016.12.030.
- [56] A.S. Kukushkin et al. “Finalizing the ITER divertor design: The key role of SOLPS modeling”. In: *Fusion Engineering and Design* 86.12 (2011), pp. 2865–2873. ISSN: 0920-3796. DOI: <https://doi.org/10.1016/j.fusengdes.2011.06.009>.
- [57] AS Kukushkin and HD Pacher. “The role of “momentum removal” in divertor detachment”. In: *Contributions to Plasma Physics* 56.6-8 (2016), pp. 711–716. ISSN: 1521-3986.
- [58] Rachel Kurchin, Giuseppe Romano, and Tonio Buonassisi. “Bayesim: A tool for adaptive grid model fitting with Bayesian inference”. In: *Computer Physics Communications* 239 (2019), pp. 161–165. ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2019.01.022>.
- [59] Sehyun Kwak et al. “Bayesian modelling of Thomson scattering and multichannel interferometer diagnostics using Gaussian processes”. In: *Nuclear Fusion* 60.4 (Feb. 2020), p. 046009. DOI: 10.1088/1741-4326/ab686e.
- [60] John D Lawson. “Some criteria for a power producing thermonuclear reactor”. In: *Proceedings of the physical society. Section B* 70.1 (1957), p. 6.
- [61] Jarrod Leddy et al. *Single Gaussian Process Method for Arbitrary Tokamak Regimes with a Statistical Analysis*. 2022. DOI: 10.48550/ARXIV.2202.11557.
- [62] LL Lengyel. *Analysis of radiating plasma boundary layers*. Tech. rep. Max-Planck-Institut für Plasmaphysik, 1981.
- [63] B.L. Linehan et al. “Validation of 2D T_e and n_e measurements made with Helium imaging spectroscopy in the volume of the TCV divertor”. In: *Nuclear Fusion* 63.3 (Feb. 2023), p. 036021. DOI: 10.1088/1741-4326/acb5b0.
- [64] Bryan Lee Linehan. “Transport Properties of Divertor Edge Plasmas Measured with Multi-Spectral Imaging”. PhD thesis. Massachusetts Institute of Technology, 2024.
- [65] B Lipschultz et al. “Variation of the divertor geometry in Alcator C-Mod”. In: *Proc. 16th Int. Conf. on Fusion Energy (Montreal, Canada)*. Vol. contributed papers, Part I. IAEA, Vienna. 1996.
- [66] B. Lipschultz et al. “Ultrahigh densities and volume recombination inside the separatrix of the Alcator C-Mod tokamak”. In: *Physical Review Letters* 81.5 (1998), pp. 1007–1010. ISSN: 0031-9007. DOI: DOI10.1103/PhysRevLett.81.1007.

-
- [67] Bruce Lipschultz, Felix I. Parra, and Ian H. Hutchinson. “Sensitivity of detachment extent to magnetic configuration and external parameters”. In: *Nuclear Fusion* 56.5 (Apr. 2016), p. 056007. DOI: 10.1088/0029-5515/56/5/056007.
- [68] D. C. Liu and J. Nocedal. “On the limited memory BFGS method for large scale optimization”. In: *Mathematical Programming* 45.1436-4646 (Aug. 1989), pp. 503–528. ISSN: 1. DOI: 10.1007/BF01589116.
- [69] A. Loarte. “Understanding the edge physics of divertor experiments by comparison of 2D edge code calculations and experimental measurements”. In: *Journal of Nuclear Materials* 241-243 (1997), pp. 118–134. ISSN: 0022-3115. DOI: [https://doi.org/10.1016/S0022-3115\(97\)80035-9](https://doi.org/10.1016/S0022-3115(97)80035-9).
- [70] Bartosz Lomanowski et al. “Interpretation of Lyman opacity measurements in JET with the ITER-like wall using a particle balance approach”. In: *Plasma Physics and Controlled Fusion* 62.065006 (2020). DOI: 10.1088/1361-6587/ab7432.
- [71] N. Lonigro et al. “First 2D electron density measurements using Coherence Imaging Spectroscopy in the MAST-U Super-X divertor”. In: *EPS Conference of Plasma Physics 2023*. 2023.
- [72] N. Lonigro et al. *First 2D electron density measurements using Coherence Imaging Spectroscopy in the MAST-U Super-X divertor*. 2024.
- [73] A. G. Meigs et al. “Deuterium Balmer/Stark spectroscopy and impurity profiles: First results from mirror-link divertor spectroscopy system on the JET ITER-like wall”. In: *Journal of Nuclear Materials* 438 (2013), S607–S611. ISSN: 0022-3115. DOI: 10.1016/j.jnucmat.2013.01.127.
- [74] B Min et al. “Lost in the dark: A survey of energy poverty from space”. In: *Joule* 8.87 (July 2024), pp. 1982–1998. DOI: 10.1016/j.joule.2024.05.001.
- [75] William Morris et al. “MAST upgrade divertor facility: a test bed for novel divertor solutions”. In: *IEEE Transactions on Plasma Science* 46.5 (2018), pp. 1217–1226.
- [76] D. Moulton et al. “Super-X and conventional divertor configurations in MAST-U ohmic L-mode; a comparison facilitated by interpretative modelling”. In: *Nuclear Fusion* 64.7 (June 2024), p. 076049. DOI: 10.1088/1741-4326/ad4f9c.
- [77] David Moulton et al. “Using SOLPS to confirm the importance of total flux expansion in Super-X divertors”. In: *Plasma Physics and Controlled Fusion* 59 (Apr. 2017). DOI: 10.1088/1361-6587/aa6b13.
- [78] J. M. Muñoz Burgos et al. “Evaluation of emission contributions from charge-exchange between the excited states of deuterium with He⁺ during diagnostic of thermal helium gas beam injection and laser-induced fluorescence”. In: *Physics of Plasmas* 26.6 (June 2019), p. 063301. ISSN: 1070-664X. DOI: 10.1063/1.5088363.
- [79] Omkar Myatra. “Numerical modelling of detached plasmas in the MAST Upgrade super-X divertor”. PhD thesis. University of York, 2021.
- [80] Radford M Neal et al. “MCMC using Hamiltonian dynamics”. In: *Handbook of Markov Chain Monte Carlo* 2.11 (2011), p. 2.
- [81] Jef Ongena. “Fusion: A true challenge for an enormous reward”. In: *EPJ Web of Conferences* 98 (Jan. 2015), p. 05004. DOI: 10.1051/epjconf/20159805004.
- [82] N Osborne et al. “A novel understanding of the role of plasma-molecular kinetics on divertor power exhaust”. In: *arXiv preprint arXiv:2410.14403* (2025).
- [83] N Osborne et al. “Initial Fulcher band observations from high resolution spectroscopy in the MAST-U divertor”. In: *Plasma Physics and Controlled Fusion* 66.2 (2023), p. 025008.

- [84] Andrea Pavone et al. “Machine learning and Bayesian inference in nuclear fusion research: an overview”. In: *Plasma Physics and Controlled Fusion* 65 (Mar. 2023). DOI: 10.1088/1361-6587/acc60f.
- [85] A. Perek et al. “A spectroscopic inference and SOLPS-ITER comparison of flux-resolved edge plasma parameters in detachment experiments on TCV”. In: *Nuclear Fusion* 62.9 (July 2022), p. 096012. DOI: 10.1088/1741-4326/ac7813.
- [86] A. Perek et al. “MANTIS: A real-time quantitative multispectral imaging system for fusion plasmas”. In: *Review of Scientific Instruments* 90.12 (2019), p. 123514. DOI: 10.1063/1.5115569.
- [87] A. Perek et al. “Measurement of the 2D emission profiles of hydrogen and impurity ions in the TCV divertor”. In: *Nuclear Materials and Energy* 26 (2021), p. 100858. ISSN: 2352-1791. DOI: 10.1016/j.nme.2020.100858.
- [88] Richard A Pitts et al. “Physics basis for the first ITER tungsten divertor”. In: *Nuclear Materials and Energy* 20 (2019), p. 100696.
- [89] Xander Pope et al. “Bayesian Analysis of Conventional and Ultrafast Spectroscopy Data for Investigating Detachment in the MAST-Upgrade Super-X”. In: *arXiv preprint arXiv:2404.15072* (2024).
- [90] ADAS Project, University of Strathclyde, and the IAEA. *OPEN-ADAS: Atomic Data and Analysis Structure*. 2018. URL: <https://open.adas.ac.uk/> (visited on 01/07/2022).
- [91] A.A. Pshenov et al. “Divertor plasma opacity effects”. In: *Nuclear Materials and Energy* 34 (2023), p. 101342. ISSN: 2352-1791. DOI: <https://doi.org/10.1016/j.nme.2022.101342>.
- [92] F. Reimold et al. “Divertor studies in nitrogen induced completely detached H-modes in full tungsten ASDEX Upgrade”. In: *Nuclear Fusion* 55.3 (2015), p. 033004. ISSN: 0029-5515. DOI: Artn03300410.1088/0029-5515/55/3/033004.
- [93] Detlev Reiter. *The data file AMJUEL: Additional atomic and molecular data for EIRENE*. Tech. rep. Forschungszentrum Jülich GmbH, Jan. 2000.
- [94] Detlev Reiter, Martine Baelmans, and Petra Boerner. “The EIRENE and B2-EIRENE codes”. In: *Fusion science and technology* 47.2 (2005), pp. 172–186.
- [95] K -U Riemann. “The Bohm criterion and sheath formation”. In: *Journal of Physics D: Applied Physics* 24.4 (Apr. 1991), p. 493. DOI: 10.1088/0022-3727/24/4/001.
- [96] Paul J Schroeder et al. “Demonstration of a two-dimensional divertor Thomson scattering system on DIII-D”. In: *Review of Scientific Instruments* 94.12 (2023).
- [97] SA Silburn et al. “Coherence imaging of scrape-off-layer and divertor impurity flows in the Mega Amp Spherical Tokamak”. In: *Review of Scientific Instruments* 85.11 (2014). DOI: 10.1063/1.4891165.
- [98] VA Soukhanovskii, AG McLean, and SL Allen. “Near-infrared spectroscopy for divertor plasma diagnosis and control in DIII-D tokamak”. In: *Review of Scientific Instruments* 85.11 (2014).
- [99] P. C. Stangeby. “Basic physical processes and reduced models for plasma detachment”. In: *Plasma Physics and Controlled Fusion* 60.4 (2018), p. 044022. ISSN: 0741-3335. DOI: 10.1088/1361-6587/aaac6f.
- [100] P.C. Stangeby. *The Plasma Boundary of Magnetic Fusion Devices*. Series in Plasma Physics and Fluid Dynamics. Taylor & Francis, 2000. ISBN: 9780750305594.
- [101] F Subba et al. “SOLPS-ITER modeling of divertor scenarios for EU-DEMO”. In: *Nuclear Fusion* 61.10 (2021), p. 106013.

-
- [102] H. P. Summers et al. “Ionization state, excited populations and emission of impurities in dynamic finite density plasmas: I. The generalized collisional–radiative model for light elements”. In: *Plasma Physics and Controlled Fusion* 48.2 (2006), pp. 263–293. ISSN: 0741-3335 1361-6587. DOI: 10.1088/0741-3335/48/2/007.
- [103] J Svensson and A Werner. “Large scale Bayesian data analysis for nuclear fusion experiments”. In: *2007 IEEE International Symposium on Intelligent Signal Processing*. IEEE, 2007, pp. 1–6.
- [104] J Svensson, A Werner, and JET-EFDA Contributors. “Current tomography for axisymmetric plasmas”. In: *Plasma Physics and Controlled Fusion* 50.8 (May 2008), p. 085002. DOI: 10.1088/0741-3335/50/8/085002.
- [105] J Svensson et al. “An integrated data analysis model for the W7-AS stellarator”. In: *30th EPS Conference on Contr. Fusion and Plasma Phys., St. Petersburg*. Vol. 27. 2003.
- [106] J. Svensson, European Fusion Development Agreement, and JET Project. *Non-parametric Tomography Using Gaussian Processes*. EFDA-JET. EFDA, 2012.
- [107] J. Svensson et al. “Integrating diagnostic data analysis for W7-AS using Bayesian graphical models”. In: *Review of Scientific Instruments* 75.10 (Oct. 2004), pp. 4219–4221. ISSN: 0034-6748. DOI: 10.1063/1.1789611.
- [108] J. Svensson et al. “Modelling of JET Diagnostics Using Bayesian Graphical Models”. In: *Contributions to Plasma Physics* 51.2-3 (2011), pp. 152–157. DOI: <https://doi.org/10.1002/ctpp.201000058>.
- [109] JL Terry et al. “Volume recombination and opacity in Alcator C-Mod divertor plasmas”. In: *Physics of Plasmas* 5.5 (1998), pp. 1759–1766.
- [110] C Theiler et al. “Results from recent detachment experiments in alternative divertor configurations on TCv”. In: *Nuclear Fusion* 57.7 (2017), p. 072008.
- [111] P. M. Valanju et al. “Super-X divertors and high power density fusion devicesa”). In: *Physics of Plasmas (1994-present)* 16.5 (2009), p. 056110.
- [112] K Verhaegh et al. “A novel hydrogenic spectroscopic technique for inferring the role of plasma–molecule interaction on power and particle balance during detached conditions”. In: *Plasma Physics and Controlled Fusion* 63.3 (Jan. 2021), p. 035018. DOI: 10.1088/1361-6587/abd4c0.
- [113] K. Verhaegh et al. *Experimental data and simulations comparing the divertor and core performance of the MAST Upgrade Super-X, Elongated and Conventional Divertors in beam-heated L-mode conditions*. UK Atomic Energy Authority, Open Data Registry. Dataset. 2024.
- [114] K. Verhaegh et al. “Investigating the impact of the molecular charge-exchange rate on detached SOLPS-ITER simulations”. In: *Nuclear Fusion* 63.7 (May 2023), p. 076015. DOI: 10.1088/1741-4326/acd394.
- [115] K. Verhaegh et al. “Novel inferences of ionisation & recombination for particle/power balance during detached discharges using deuterium Balmer line spectroscopy”. In: *Plasma Phys. Control. Fusion* 61.125018 (2019). DOI: 10.1088/1361-6587/ab4f1e.
- [116] K. Verhaegh et al. “Spectroscopic investigations of detachment on the MAST Upgrade Super-X divertor”. In: *Nuclear Fusion* 63.1 (Dec. 2022), p. 016014. DOI: 10.1088/1741-4326/aca10a.
- [117] K. Verhaegh et al. “The role of plasma-molecule interactions on power and particle balance during detachment on the TCv tokamak”. In: *Nuclear Fusion* 61.10 (Sept. 2021), p. 106014. DOI: 10.1088/1741-4326/ac1dc5.

- [118] K. Verhaegh et al. “The role of plasma–atom and molecule interactions on power; particle balance during detachment on the MAST Upgrade Super-X divertor”. In: *Nuclear Fusion* 63.12 (Oct. 2023), p. 126023. DOI: 10.1088/1741-4326/acf946.
- [119] Kevin Verhaegh. “Spectroscopic investigations of detachment on TCV: Investigating the role of atomic physics on the ion current roll-over and the dynamics of detachment in TCV”. PhD thesis. University of York, 2018.
- [120] Kevin Verhaegh et al. “A study of the influence of plasma–molecule interactions on particle balance during detachment”. In: *Nuclear Materials and Energy* 26 (2021), p. 100922. ISSN: 2352-1791. DOI: <https://doi.org/10.1016/j.nme.2021.100922>.
- [121] Kevin Verhaegh et al. “An improved understanding of the roles of atomic processes and power balance in divertor target ion current loss during detachment”. In: *Nuclear Fusion* 59.126038 (2019). DOI: 10.1088/1741-4326/ab4251.
- [122] Kevin Verhaegh et al. “Divertor shaping with neutral baffling as a solution to the tokamak power exhaust challenge”. In: *Communications Physics* 8.1 (2025), p. 215.
- [123] Kevin Verhaegh et al. “Investigations of atomic and molecular processes of NBI-heated discharges in the MAST Upgrade Super-X divertor with implications for reactors”. In: *Nuclear Fusion* 64.8 (July 2024), p. 086050. DOI: 10.1088/1741-4326/ad5851.
- [124] Irina Veselova et al. “SOLPS-ITER drift modelling of ITER burning plasmas with narrow near-SOL heat flux channels”. In: *Nuclear Materials and Energy* 26 (2021), p. 100870. ISSN: 2352-1791. DOI: <https://doi.org/10.1016/j.nme.2020.100870>.
- [125] Fritz Wagner et al. “Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak”. In: *Physical Review Letters* 49.19 (1982), p. 1408.
- [126] L Wang et al. “Integration of full divertor detachment with improved core confinement for tokamak fusion plasmas”. In: *Nature communications* 12.1 (2021), p. 1365.
- [127] Kelly Warner and Gary M Hieftje. “Thomson scattering from analytical plasmas”. In: *Spectrochimica Acta Part B: Atomic Spectroscopy* 57.2 (2002), pp. 201–241.
- [128] R.P. Wenninger et al. “DEMO divertor limitations during and in between ELMs”. In: *Nuclear Fusion* 54.11 (Nov. 2014), p. 114003. DOI: 10.1088/0029-5515/54/11/114003.
- [129] M Wensing et al. “SOLPS-ITER simulations of the TCV divertor upgrade”. In: *Plasma Physics and Controlled Fusion* 61.8 (July 2019), p. 085029. DOI: 10.1088/1361-6587/ab2b1f.
- [130] M. Wensing et al. “SOLPS-ITER validation with TCV L-mode discharges”. In: *Physics of Plasmas* 28.8 (Aug. 2021), p. 082508. ISSN: 1070-664X. DOI: 10.1063/5.0056216.
- [131] John Wesson. *Tokamaks*. Vol. 149. Oxford University Press, 2011.
- [132] S. Wiesen et al. “The new SOLPS-ITER code package”. In: *Journal of Nuclear Materials* 463 (2015). PLASMA-SURFACE INTERACTIONS 21, pp. 480–484. ISSN: 0022-3115. DOI: <https://doi.org/10.1016/j.jnucmat.2014.10.012>.
- [133] T.A. Wijkamp et al. “Characterisation of detachment in the MAST-U Super-X divertor using multi-wavelength imaging of 2D atomic and molecular emission processes”. In: *Nuclear Fusion* 63.5 (Mar. 2023), p. 056003. DOI: 10.1088/1741-4326/acc191.
- [134] D. Wunderlich et al. “Yacora on the Web: Online collisional radiative models for plasmas containing H, H₂ or He”. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 240 (2020), p. 106695. ISSN: 0022-4073. DOI: <https://doi.org/10.1016/j.jqsrt.2019.106695>.
- [135] Dirk Wunderlich and Ursel Fantz. “Evaluation of State-Resolved Reaction Probabilities and Their Application in Population Models for He, H, and H₂”. In: *Atoms* 4.4 (2016). ISSN: 2218-2004. DOI: 10.3390/atoms4040026.

- [136] J.H. You et al. “Divertor of the European DEMO: Engineering and technologies for power exhaust”. In: *Fusion Engineering and Design* 175 (2022), p. 113010. ISSN: 0920-3796. DOI: <https://doi.org/10.1016/j.fusengdes.2022.113010>.

Appendix A

Mass Matrix Construction for Hamiltonian Monte Carlo

The efficiency of Hamiltonian Markov Chain Monte Carlo (HMC) is influenced by the choice of the mass matrix, M . In HMC [80], the current state of the Markov chain receives momentum, p , which is randomly sampled from

$$p \sim \mathcal{N}(0, M). \quad (\text{A.1})$$

The mass matrix also appears in the kinetic energy term of the approximately conserved Hamiltonian,

$$\underbrace{-\mathcal{L}(\theta \mid \mathcal{D})}_{\text{potential energy}} + \underbrace{\frac{1}{2}p^T M^{-1}p}_{\text{kinetic energy}}, \quad (\text{A.2})$$

where \mathcal{L} is the log-probability, θ are the inferred parameters, and \mathcal{D} is the measured data.

If the posterior distribution exhibits strong parameter correlations or large differences in parameter scale, a poor choice of mass matrix can lead to inefficient trajectories, high autocorrelation between samples, and low acceptance rates. Ideally, the mass matrix approximates the posterior covariance

$$M \approx \Sigma, \quad (\text{A.3})$$

such that the transformed posterior geometry becomes approximately isotropic. Therefore, approximating the covariance matrix enables for an estimation of the mass matrix.

A.1 Adaptive Estimation

One approach is to estimate the mass matrix adaptively during sampling. In this approach, a set number of samples, N , from the chain is used to empirically approximate

the covariance matrix:

$$\hat{\Sigma} = \frac{1}{N-1} \sum_{i=1}^N (\theta_i - \bar{\theta})(\theta_i - \bar{\theta})^T, \quad (\text{A.4})$$

where $\bar{\theta}$ is the sample mean. This estimation can be updated throughout the chain's progress. However, the determination of the hamiltonian, equation (A.2), requires the inversion of the mass matrix. Matrix inversion approximately scales with $\mathcal{O}(n^3)$, where n is the dimensionality of the parameter space. Therefore, repeated updates to the mass matrix are computationally expensive for large parameter spaces.

A.2 MAP Curvature Approximation

A second approach approximates the posterior locally as a multivariate Gaussian distribution in the vicinity of the maximum a posteriori (MAP) estimate. By first using an optimiser to determine the MAP estimate, θ_{MAP} , the local posterior covariance can be estimated once and subsequently used as an approximation to M throughout HMC sampling.

The local posterior covariance can be approximated through the inverse Hessian:

$$\Sigma_{\text{MAP}} \approx H^{-1}. \quad (\text{A.5})$$

This comes from a second-order Taylor expansion about the MAP estimate, θ_{MAP} , gives

$$\mathcal{L}(\theta|\mathcal{D}) \approx \mathcal{L}(\theta_{\text{MAP}}|\mathcal{D}) + \frac{1}{2}(\theta - \theta_{\text{MAP}})^T H (\theta - \theta_{\text{MAP}}), \quad (\text{A.6})$$

where

$$H = \nabla_{\theta}^2 \mathcal{L}(\theta|\mathcal{D}) \Big|_{\theta=\theta_{\text{MAP}}} \quad (\text{A.7})$$

is the Hessian of the negative log-posterior evaluated at the MAP.

This corresponds to the Gaussian approximation

$$P(\theta|\mathcal{D}) \approx \mathcal{N}(\theta_{\text{MAP}}, H^{-1}), \quad (\text{A.8})$$

such that the inverse Hessian approximates the local posterior covariance.

When analytic second derivatives are unavailable, the Hessian may be approximated numerically using finite differences. However, finite difference approximations may produce Hessians that are not positive definite due to numerical noise and an invalid local Gaussian assumptions. Since the covariance matrix used in HMC must

be invertible, a positive definite approximation is required. As with equation (34) of [104] (see [43] for additional details), eigendecomposition,

$$H = Q\Lambda Q^T, \quad (\text{A.9})$$

followed by eigenvalue thresholding,

$$\tilde{\Lambda}_{ii} = \max(\Lambda_{ii}, \epsilon), \quad (\text{A.10})$$

(where $\epsilon > 0$ is a small regularisation parameter) enables the reconstruction of the regularised, positive-definite Hessian:

$$\tilde{H} = Q\tilde{\Lambda}Q^T. \quad (\text{A.11})$$

This provides the covariance approximation

$$\tilde{\Sigma}_{\text{MAP}} = \tilde{H}^{-1}. \quad (\text{A.12})$$

A.3 *A Priori* Construction

In this work, the majority of inferred parameters corresponded to spatial plasma fields. In 5.3.2, it is discussed how Gaussian process (GP) priors can inform the structure of these plasma fields. Consequently, physically informed covariance structure could be specified *a priori* without requiring a Hessian estimation.

The GP covariance matrices naturally encode expected parameter correlations and characteristic spatial length scales. These matrices therefore formed the dominant block structure of the HMC mass matrix. Additional terms were included to account for inter-field correlations, diagnostic parameter couplings, and nuisance parameters such as camera correlation coefficients.

The approximate covariance matrix was constructed as

$$\Sigma_{ij} = d_{ij} \begin{cases} \Sigma_{ij}^{GP}, & i, j \text{ within the same inferred field} \\ C_{ij}^f, & i, j \text{ in different correlated fields} \\ c_{ij}^{cam}, & i, j \text{ corresponding to camera correlation parameters} \\ \left(\frac{upr_i - lwr_i}{2}\right)^2, & i = j \text{ for independent scalar parameters} \\ 0, & i \neq j \text{ for uncorrelated scalar parameters} \end{cases}, \quad (\text{A.13})$$

where d_{ij} represents optional diagnostic weighting factors, Σ^{GP} denotes the GP prior covariance for individual fields, C^f captures inter-field covariance structure (which,

to a first-order approximation, can be determined through analysis of SOLPS-ITER cases), and c^{cam} describes known camera parameter correlations.

This approach provided a computationally efficient approximation to the posterior covariance while embedding physically informed parameter correlations directly into the HMC dynamics.

Appendix B

Settling Aspects of the Bayesian Inference Approach of this Work

The integrated data analysis work of this thesis is reliant on the Balmer series. The synthetic testing of Chapter 3, which explicitly looks at the information content of the Balmer lines, enables a convenient method for investigating and deciding on certain nuances of Bayesian inference. This includes the hyperparameters associated with the likelihood calculation and the preferred characterisation of the posterior distribution. These aspects are discussed in more detail in this appendix.

B.1 Non Informative Prior Choice

Inherent to Bayesian inference is the acknowledgement that not including any additional prior information is itself a subjective choice that can alter the inference. This is exemplified when considering parameters that cover substantially different scales. In Figure 3.4 we observed a systematic offset between the true temperature and the inferred temperature for ionisation conditions. We saw in Figure 3.3 that similar line ratios exist for both molecular recombination conditions ($T_e \sim 1\text{-}3$ eV, vertical bands) and ionisation conditions ($T_e \gtrsim 3$ eV, horizontal bands). When finding the marginalised electron temperature probability distribution, an integral is performed over the n_e , n_D and Q_{mol} axes. Observing Figure 3.3, the narrow vertical bands around 1 eV cover a far larger range in electron density than the horizontal bands with the equivalent line ratio. Consequently, when marginalising over the electron density axis, there is a spurious increase in the probability around 1 eV (MAR associated emission) in comparison to the higher temperature found in ionisation conditions (aligned to electron-impact excitation, EIE). The large scale of the electron density




Figure B.1: Inferred marginalised electron temperature probability distributions through Bayesian inference with Balmer lines D_α , D_β , D_γ , and D_δ .

Upper row (a-c): Jeffreys Prior included

Lower row (d-f): Fulcher data included

Samples refer to multiple trials of the inference, each with different synthetic data owing to the expected 10% and 20% uncertainty on Balmer line and Fulcher measurement data respectively.

parameter is the reason for the systematic offset between the true and inferred temperature in ionisation conditions.

Jeffreys prior is an alternative non-informative prior designed to be scale invariant [45]. The Jeffreys prior for the scale parameter of the Gaussian distribution, σ , is given by $P(\sigma) \propto 1/\sigma$. This prior does not introduce any bias towards a particular scale of σ (it is scale-invariant). It reflects a non-informative prior belief, ensuring that all possible scales of σ are treated equally before observing any data. This is generalised to the logarithmic prior which, rather than assuming all parameters have a uniform probability distribution in linear-space, assumes that all parameters have a uniform probability distribution in log-space such that $P(\theta) \propto 1/\theta$. This negates the impact of parameters having substantially different scales.

The logarithmic prior is implemented in the upper row (a-c) of Figure B.1. However, we now observe a substantial increase in the inference error under molecular recombination conditions (which are now inferred to be more likely to be obtained from higher temperatures) in comparison to Figure 3.4.d-f. The subjective choice in non-informative priors (uniform parameter probabilities in linear or uniform parameter probabilities in log space) having such an influence on our inference is clearly undesirable. It is symptomatic of emission from MAR and ionisation conditions having Balmer line ratios that are indecipherable (within the uncertainty expected through the emission models). Additional data, such as the Fulcher band, is required to separate MAR and ionisation conditions (Figure B.1.d-f).

B.2 Likelihood Degrees of Freedom

The choice of likelihood distribution is extremely influential in Bayesian inference. The Student's t distribution, equation (2.22), contains the hyperparameter ν that relates to how outliers are handled. Low values of ν are far more lenient to the existence of outliers than the Gaussian distribution ($\nu \rightarrow \infty$). This is potentially of considerable importance when our Balmer emission models are known to contain uncertainties; we may not wish penalise outliers as strongly as the Gaussian distribution which was assumed throughout this thesis so far (equation (3.7)).

Repeating the analysis of Chapter 3 with a lower value of ν (than the earlier assumed Gaussian likelihood, $\nu \rightarrow \infty$) typically increased the inference uncertainty (the 95% highest density interval). This is because parameter combinations that result in emissivity predictions very different from the measured data were not prescribed with such a low probability. Consequently, there was a wider region of parameter space that *could* have given rise to the measured data. However, decreasing ν did not substantially improve the accuracy of inference (how often the true parameter was within the 95% highest density interval). This was because decreasing ν increased the relative influence of the prior choice. As discussed in Section B.1, the choice of non-informative prior is subjective and so we do not desire to increase its influence. Consequently, the Gaussian likelihood ($\nu \rightarrow \infty$) was found to be optimal.

B.3 Likelihood Scale

It is common to take the scale of the likelihood function (σ_n of equation (3.7)) as the uncertainty of the forward model. Our synthetic testing found that the inference accuracy was typically low which was considered to be due to an underestimate of the inference uncertainty (which is heavily influenced by the scale parameter). Subsequent testing found that increasing the scale parameter increased the 95% highest density interval (taken as the inference uncertainty), but it did not substantially improve the inference accuracy. As with decreasing the decreasing ν hyperparameter discussed in Section B.2, increasing the scale parameter increases the relative influence of the subjective, non-informative prior choice which is undesirable. Consequently, the optimal likelihood scale was taken to be the uncertainty in the emission model.

B.4 Comparing Inferred Parameter Estimation Methods

While it is common to consider the parameters that give the mode of the multidimensional posterior distribution (θ_{MAP}) as the inferred parameters, it is not always the best estimator. The mean of samples from the posterior distribution represents an alternative estimator. This has the benefit of capturing uncertainties in the various parameters within the estimator by averaging over them. This estimator is known as the marginalised expectation (MEX), θ_{MEX} .

The synthetic tests performed in Chapter 3 can be used to investigate these two estimators. θ_{MAP} and θ_{MEX} were found to differ considerably. The differences were found to be due to the posterior distribution being multimodal. However, neither θ_{MAP} nor θ_{MEX} was found to universally provide a better estimator of θ_{True} . Typically, θ_{MEX} resulted in higher electron temperature percentage error in regions A and B. This was due to θ_{MEX} being more strongly influenced by inferred uncertainties which were found to be of low accuracy in regions A and B. Due to its slightly lower percentage error, θ_{MAP} was taken to be the preferred estimator.

Appendix C

Helium Emission Models

MAST-U's MWI and TCV's MANTIS cameras also routinely capture Helium-I emission lines which has been shown [63] to provide information on n_e and T_e . The nature of integrated data analysis is such that a single inaccurate forward model can offset the overall inference. As such we explicitly include parameters for n_{He^0} and $n_{\text{He}^{1+}}$ at each mesh vertex (at a cost of $2V$ additional free parameters). Doing so avoids assumptions of ionisation balance models and permits inclusion of both electron-impact excitation (of He^0) and electron-ion recombination (of He^+) contributions to emission. He I PECs were provided by ADAS [90] and implemented as with equation (4.5).

In conditions with significant hydrogen atomic densities, charge-exchange between helium and hydrogen could occur. This can change both the ionisation fractions of He (which is included in our model through the n_{He^0} and $n_{\text{He}^{1+}}$ parameters) and lead to the emission of He I [78]. Although this is neglected in our emission model, since population coefficients for this interaction are not freely available, having neutral hydrogen density, n_D , as a parameter in our model means that charge exchange can trivially be included once population coefficients for this interaction become available.

The 668, 728 and 502 nm singlet He I lines were chosen for the helium forward models in Chapter 4. This followed from the work by [8], which used these transitions due to their insensitivity to transport of metastable states, which can impact singlet to triplet He I line ratios. Detailed studies on TCV using He I emission [63] have suggested that the magnetic field as well as molecular effects could impact He I emission. However, based on those results, and the relatively low magnetic field at MAST-U (0.7 T on-axis), we would not expect that this would impact the chosen He I singlet transitions significantly.

Appendix D

Priors for Chapter 4

D.1 Prior Distributions

Each prior requires: the calculation, from θ , of the physical quantity that the prior contribution is describing (the ‘argument’), $f_j(\theta)$; and a probability function describing the expected distribution of these arguments. Spatially-dependent priors on the smoothness of the parameter fields are particularly useful in our work. However, whilst a smooth field is more probable than a noisy field, perfectly smooth fields are not expected. This highlights the importance of the choice of the (log) prior probability function; the probability to attribute to different degrees of smoothness (quantified as the prior argument) is somewhat arbitrary but significantly impacts the inference.

An empirical approach was used to determine the prior probability functions. For each prior, its argument at each mesh vertex was evaluated across many SOLPS-ITER simulations (excluding those used to generate the synthetic data in this work). This resulted in a distribution of expected prior arguments for each prior. By fitting these distributed prior arguments with different differentiable probability distributions, we found that a heavy-tailed function was optimal. This allowed for a preference for compliance with the prior rationale without significantly penalising outliers. For example, the majority of vertices in the mesh would be expected to follow a monotonic T_e parallel to flux surfaces but, by having a heavy-tailed distribution, non-monotonic variations are permitted should the data suggest so.

For each prior, it was found that a Cauchy distribution best characterised the prior arguments (with each mesh vertex, k , independently treated). Consequently, each prior had a log-probability of

$$\mathcal{L}_j(\theta) = V \log(\gamma_j) - \sum_{k=1}^V \log \left((f_j(\theta_k, \dots))^2 + \gamma_j^2 \right). \quad (\text{D.1})$$

γ_j , the scale parameter, was empirically found for each individual prior and details of these are outlined in Appendix D.2.

D.2 Prior Details

For each prior, its argument was evaluated at all relevant mesh vertices, k , for multiple SOLPS-ITER cases. For all priors, the Cauchy distribution best encapsulated the prior argument distribution and the scale parameter, γ , was found.

Prior details for the Details of the spatially-independent priors are given in Table D.1 and details of the spatially-dependent priors are given in Table D.2.

Table D.1: The spatially independent priors implemented within D-MIBAS for Chapter 4. When the argument deploys an upper and lower limit, the scale parameters are respectively listed.

Prior (j)	Argument ($f_j(\theta_k)$)	γ_j	Justification
Static Electron Pressure, $P_e^k = T_e^k n_e^k$	$\max\left(\frac{P_e^k}{P_e^{\max}} - 1, 0\right)$	1.5×10^{-2}	Static electron pressure is not expected to exceed $P_e^{\max} = 7 \times 10^{20} \text{ eV m}^{-3}$ in the MAST-U Super-X divertor (Figure 4 of [9]).
Neutral Fraction, $f_{n_H}^k = \frac{n_H^k}{n_e^k + n_H^k}$	$\max\left(\frac{f_0}{f_{n_H}^k} - 1, \frac{f_{n_H}^k}{f(T_e)} - 1, 0\right)$	$(2.3 \times 10^{-3}, 4.1 \times 10^{-5})$	In the MAST-U Super-X divertor, the neutral fraction is expected to remain between $f_0 = 0.002$ and $f(T_e) = (1 - c) \exp(-T_e^k/l) + c$, where $c = 0.04$ and $l = 5 \text{ eV}$ (Figure 4 of [9]).

Continued on next page

Prior (j)	Argument ($f_j(\theta_k)$)	γ_j	Justification
Helium ionisation ratio, $I_{He^{1+}}^k = \frac{n_{He^{1+}}^k}{n_{He^0}^k}$	$\max\left(\frac{I_0(T_e)}{I_{He^{1+}}^k} - 1, \frac{I_{He^{1+}}^k}{I_1(T_e)} - 1, 0\right)$	$(1.4 \times 10^{-3}, 1.4 \times 10^{-3})$	Limit set on the ratio of singly charged helium density $n_{He^{1+}}$ to neutral helium density n_{He^0} . Temperature-dependent limits $I_0(T_e)$ and $I_1(T_e)$ were empirically found through SOLPS-ITER case analysis.
Helium concentration ratio, $\frac{C_{He}^k}{n_{He^0}^k + n_{He^{1+}}^k} = \frac{C_{He}^k}{n_e^k}$	$\max\left(\frac{C_0}{C_{He}^k} - 1, \frac{C_{He}^k}{C_1} - 1, 0\right)$	$(3.6 \times 10^{-2}, 3.6 \times 10^{-2})$	Limits on the expected ratio of helium density (neutral n_{He^0} and singly charged $n_{He^{1+}}$) to electron density in the MAST-U Super-X divertor, found empirically via SOLPS-ITER case analysis.
Parameter bounds, $\theta_F \in (T_e, n_e, n_H, Q_{\text{mol.}}, n_{He^0}, n_{He^{1+}})$	$\max\left(\frac{\log(A)}{\log(\theta_F^k)} - 1, \frac{\log(\theta_F^k)}{\log(B)} - 1, 0\right)$		$A \in (0.2, 1 \times 10^{16}, 1 \times 10^{15}, 1 \times 10^{-3}, 1 \times 10^{15}, 1 \times 10^{13})$ respectively. $B \in (60, 2.5 \times 10^{20}, 2 \times 10^{20}, 1 \times 10^3, 2 \times 10^{19}, 2 \times 10^{19})$. A normal distribution with $\sigma = 1 \times 10^{-4}$ was used in place of the Cauchy distribution.

Table D.2: The spatially dependent priors implemented within D-MIBAS for Chapter 4. Where multiple fields F are stated, the scale parameters are, respectively, listed. Derivatives $\partial_{s_{\parallel}}$, $\partial_{s_{\theta,\parallel}}$ and $\partial_{s_{\psi_N}}$ are with respect to distance parallel to total flux surfaces, distance parallel to surfaces of constant poloidal magnetic flux, and distance perpendicular to surfaces of constant poloidal magnetic flux, respectively. Gradients are approximated by means of finite differences with neighbouring mesh vertices.

Prior (j)	Argument ($f_j(\theta_k, \dots)$)	γ_j	Justification
Pressure Drop, $P_e^k = T_e^k n_e^k$	$\min\left(0, \frac{P_e^{k,\text{upstream}}}{P_e^k} - 2\right)$ δ_k at target $\delta_{T_e^{k,\text{upstream}} > 5 \text{ eV}}$	$(1.7 \times 10^{-2}, 1.7 \times 10^{-2})$	From the two-point model, we expect at least a factor-of-two reduction in static electron pressure from upstream to target along a flux tube [100]. Since this pressure-reduction argument loses validity if the upstream vertex is below the detachment threshold, the prior is applied only where $T_e^{k,\text{upstream}} > 5 \text{ eV}$ and only at mesh vertices along the target.
Parallel Gradient of log, $F^k \in (T_e^k, T_e^k n_e^k)$	$\max(0, \partial_{s_{\parallel}} \log F^k)$	$(2.4 \times 10^{-3}, 1.0 \times 10^{-2})$	Expected monotonic reduction in electron temperature and static electron pressure along a path to the target, s_{\parallel} , parallel to surfaces of constant magnetic flux.
Parallel Gradient of, $Q_{\parallel}^k = -\kappa T_e^k \partial_{s_{\parallel}} T_e^k$	$\max(0, \partial_{s_{\parallel}} Q_{\parallel}^k)$	$4.0 \times 10^{+2}$	Expected monotonic reduction in parallel heat flux along a path to the target, s_{\parallel} . The prior's argument form allows the thermal conductivity κ to take any positive value.

Continued on next page

Prior (j)	Argument ($f_j(\theta_k, \dots)$)	γ_j	Justification
Perpendicular gradient of, F^k ($T_e^k, n_e^k, T_e^k n_e^k, n_{He^{1+}}^k$)	$\max\left(0, \text{sgn}\left(F^k \in F^{k,\max}\right) \partial_{s_{\psi_N}} F^k\right)$	– (1.6 × 10 ⁻² , 6.9 × 10 ⁻² , 4.2 × 10 ⁻² , 1.6 × 10 ⁻¹)	Expected single peak in electron temperature, electron density, static electron pressure, and singly charged helium density along a path s_{ψ_N} perpendicular to surfaces of constant poloidal magnetic flux.
Parallel smoothness of log, F^k ($T_e^k, n_e^k, n_{He^{1+}}^k$)	$\partial_{s_{\parallel} s_{\parallel}}^2 \log F^k$	(1.2 × 10 ⁻¹ , 1.0 × 10 ⁻¹ , 1.7 × 10 ⁻¹)	Expected smoothness in the relative change of electron temperature, electron density, and singly charged helium density along s_{\parallel} .
Parallel smoothness of log, $F^k \in (n_H^k, n_{He^0}^k)$	$\partial_{s_{\theta, \parallel} s_{\theta, \parallel}}^2 \log F^k$	(1.9 × 10 ⁺¹ , 5.3 × 10 ⁺¹)	Expected smoothness in the relative change of neutral hydrogen and helium densities. As these species are not magnetised, we evaluate along $s_{\theta, \parallel}$, parallel to surfaces of constant <i>poloidal</i> magnetic flux.
Perpendicular smoothness of log, F^k ($T_e^k, n_e^k, n_H^k, n_{He^0}^k, n_{He^{1+}}^k$)	$\partial_{s_{\psi_N} s_{\psi_N}}^2 \log F^k$	(1.8 × 10 ⁺² , 2.2 × 10 ⁺² , 3.3 × 10 ⁺² , 7.9 × 10 ⁺² , 7.6 × 10 ⁺²)	Expected smoothness in the relative change of electron temperature and of electron, neutral hydrogen, neutral helium, and singly charged helium densities along s_{ψ_N} , perpendicular to surfaces of constant poloidal magnetic flux; larger scale parameters are applied in this direction.