

UNIVERSITY OF CALIFORNIA

Los Angeles

Study of turbulence and transport in and optimization of
mirror configurations in the Large Plasma Device

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

by

Philip Travis

2025

© Copyright by

Philip Travis

2025

ABSTRACT OF THE DISSERTATION

Study of turbulence and transport in and optimization of
mirror configurations in the Large Plasma Device

by

Philip Travis

Doctor of Philosophy in Physics

University of California, Los Angeles, 2025

Professor Troy A. Carter, Chair

The primary goal of this thesis is to work towards accelerating and automating fusion science. The combination of the physical flexibility of mirror machines with the optimization and information-exploitation potential of machine learning may be a potent one and is explored here. Progress towards this overarching goal is accomplished by studying turbulence in a mirror machine, optimizing mirror configurations using machine learning, and developing a highly extendable and flexible model for correlating and interpreting diagnostic signals.

In the Large Plasma Device (LAPD) at UCLA, a study of turbulence and transport in mirror configurations was undertaken. Using the flexible nature of the LAPD field configuration, several different mirror ratios from $M = 1$ to $M = 2.68$ were studied. Langmuir and magnetic probes were used to measure profiles of density, temperature, potential, and magnetic field. Particle flux measurements were also taken. The goal of this work was to see the interaction of interchange modes with drift waves, but no such interchange modes were observed likely because of the many stabilization phenomena present. This fact, along with reduced cross-field particle flux, indicate that a sufficiently cold edge of a simple mirror may have less cross-field transport than one would

expect.

For the purposes of machine learning, a partially-randomized dataset was collected in the LAPD mirror configurations. The goal was to maximize the diversity of data to cover the largest portion of machine operation space as possible. Using this collected dataset, neural network (NN) ensembles with uncertainty quantification were trained to predict time-averaged ion saturation current (I_{sat} — proportional to density and the square root of electron temperature) at any position within the dataset domain. This model was then used to optimize the device for strong, intermediate, and weak axial variation of I_{sat} . In addition, this model was used to infer trends in the effect on I_{sat} of LAPD controls. This model and optimization were validated on followup experiments, yielding qualitative and, at times, quantitative, agreement. This investigation demonstrated that, using ML techniques, insights can be extracted from experiments and magnetized plasmas can be globally optimized. The primary goals of this work were to provide an example of a solid, validated machine learning study and demonstrate how ML can be useful in understanding operating plasma devices.

Using this same randomized dataset, a generative model was trained to learn a probability distribution. In particular, energy-based models (EBMs) provide a powerful and flexible way of learning relationships in data by constructing an energy surface. In this work, a CNN- and attention-based multimodal EBM was trained on time series and single-dimensional data. This EBM learned all distributional modes of the data but with some differences in probability mass. Via conditional sampling of the model using a novel, auxiliary energy function technique, diagnostic reconstruction is demonstrated. In addition, the inclusion of additional diagnostics improved reconstruction error and generation quality, showing that even uncalibrated, unanalyzed diagnostics can provide useful information. Fundamentally, this work demonstrated the flexibility and efficacy of EBM-based generative modeling of laboratory plasma data, and demonstrated practical use of EBMs in the physical sciences.

The dissertation of Philip Travis is approved.

Eduardo Paulo Jorge da Costa Alves

Christoph Niemann

Jacob Bortnik

Troy A. Carter, Committee Chair

University of California, Los Angeles

2025

*To humanity
and Altair, my dog*

TABLE OF CONTENTS

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | Nuclear fusion via magnetic confinement | 1 |
| 1.2 | Confinement degradation from instabilities and turbulence | 3 |
| 1.2.1 | Importance of this turbulence and transport study | 4 |
| 1.3 | Accelerating research using machine learning | 5 |
| 1.3.1 | Importance of these ML models | 7 |
| 1.4 | Outline of the dissertation | 8 |
| 2 | Background | 10 |
| 2.1 | The Large Plasma Device at UCLA | 10 |
| 2.1.1 | Plasma source | 10 |
| 2.1.2 | Magnetic field | 12 |
| 2.1.3 | Gas fueling | 12 |
| 2.2 | Data acquisition | 14 |
| 2.3 | Diagnostics | 15 |
| 2.3.1 | Langmuir probes: I_{sat} , sweeps, triple probes | 15 |
| 2.3.2 | Magnetic flux (Bdot) probes | 20 |
| 2.3.3 | The 288 GHz heterodyne interferometer | 21 |
| 2.3.4 | Thomson scattering | 21 |
| 2.3.5 | Fast framing camera | 22 |
| 2.4 | Potential instabilities in the LAPD | 23 |
| 2.4.1 | Drift waves | 23 |

| | | |
|----------|---|-----------|
| 2.4.2 | Curvature interchange | 23 |
| 2.4.3 | Rotational interchange | 25 |
| 2.4.4 | Nonlinear instability and conducting wall mode | 26 |
| 2.5 | Machine learning and neural networks | 26 |
| 2.5.1 | Fundamentals of neural networks | 26 |
| 2.5.2 | Common layer types | 27 |
| 2.5.3 | Generative models | 28 |
| 3 | Turbulence and transport in mirror geometries in the Large Plasma Device | 30 |
| 3.1 | Introduction | 31 |
| 3.2 | The experiment and device configuration | 33 |
| 3.2.1 | The Large Plasma Device (LAPD) | 33 |
| 3.2.2 | Diagnostics | 35 |
| 3.3 | Mirror-induced changes | 35 |
| 3.3.1 | Profile modification | 35 |
| 3.3.2 | Reduced particle flux | 39 |
| 3.3.3 | Compensating for the T_e profile | 48 |
| 3.3.4 | Drift waves | 50 |
| 3.3.5 | Turbulence modification | 55 |
| 3.3.6 | Magnetic fluctuations | 56 |
| 3.4 | 2d Structure | 63 |
| 3.5 | Discussion | 69 |
| 3.5.1 | Lack of mirror-driven instabilities | 69 |
| 3.5.2 | Instabilities driving LAPD turbulence | 73 |

| | | |
|----------|--|-----------|
| 3.5.3 | Causes of particle flux reduction | 74 |
| 3.5.4 | Differences between DR1 and DR2 | 78 |
| 3.6 | Conclusions and future work | 80 |
| 4 | Creating a randomized dataset for machine learning tasks | 84 |
| 4.1 | Goal and introduction | 84 |
| 4.2 | Configuration of the LAPD | 85 |
| 4.3 | Signals collected | 87 |
| 4.4 | Data cleaning | 91 |
| 4.5 | Data bias | 92 |
| 4.6 | Azimuthal asymmetry of probe data | 97 |
| 4.7 | Applying and improving the dataset | 98 |
| 5 | Optimizing mirror configurations in the LAPD using machine learning | 99 |
| 5.1 | Introduction | 99 |
| 5.2 | Processing of I_{sat} signals | 102 |
| 5.3 | Model development and training | 105 |
| 5.3.1 | Model inputs | 105 |
| 5.3.2 | Training details | 106 |
| 5.3.3 | Validating the training pipeline | 106 |
| 5.3.4 | Baselines for mean-squared error | 107 |
| 5.3.5 | Effects of training set and model sizes | 111 |
| 5.3.6 | Improving performance with machine state flags | 111 |
| 5.3.7 | Learning rate scheduling | 112 |

| | | |
|----------|--|------------|
| 5.4 | Uncertainty quantification | 112 |
| 5.4.1 | β -NLL loss | 112 |
| 5.4.2 | Cross-validation MSE | 113 |
| 5.4.3 | Model calibration via weight decay | 115 |
| 5.5 | Evaluating model performance | 120 |
| 5.5.1 | Checking geometrical intuition | 120 |
| 5.5.2 | Directly comparing prediction to measurement | 122 |
| 5.5.3 | Comparison with Thomson scattering | 124 |
| 5.6 | Effect of I_{sat} calibration | 125 |
| 5.7 | Inferring trends | 127 |
| 5.8 | Optimizing profiles | 130 |
| 5.9 | Discussion | 136 |
| 5.9.1 | Key achievements | 136 |
| 5.9.2 | Current limitations | 137 |
| 5.9.3 | Future directions | 138 |
| 5.10 | Conclusion | 138 |
| 5.11 | The open dataset and repository | 139 |
| 6 | Energy-based models for diagnostic reconstruction and distribution modeling . . . | 140 |
| 6.1 | Goals and prior work | 141 |
| 6.1.1 | Goals | 141 |
| 6.1.2 | Generative modeling of plasmas | 141 |
| 6.2 | Introduction to energy-based models (EBMs) | 142 |
| 6.3 | Data preparation | 143 |

| | | |
|----------|--|------------|
| 6.4 | Training the EBM | 144 |
| 6.4.1 | The sampler | 145 |
| 6.4.2 | Replay buffer configuration | 147 |
| 6.4.3 | Sampling behavior | 149 |
| 6.5 | Architecture | 150 |
| 6.6 | Unconditional sampling | 153 |
| 6.7 | Diagnostic reconstruction via conditional sampling | 154 |
| 6.7.1 | Sampling interferometer signals | 157 |
| 6.8 | Symmetries and trends in the energy function | 159 |
| 6.9 | Conclusion | 162 |
| 6.10 | Future work | 162 |
| 7 | Conclusions | 164 |
| 7.1 | Mirror turbulence and transport | 164 |
| 7.2 | Machine learning on the LAPD | 165 |
| 7.3 | Future work | 167 |
| 7.4 | Accelerating fusion science | 169 |
| A | 0D mirror optimization | 171 |
| A.1 | List of assumptions / conditions | 171 |
| A.2 | User specified parameters | 172 |
| A.2.1 | Simple mirror endplug | 172 |
| A.2.2 | Tandem mirror | 173 |
| A.2.3 | Engineering parameters | 174 |

| | | |
|-------|--|-----|
| A.3 | Fusion | 174 |
| A.3.1 | Reactivity | 174 |
| A.3.2 | Fusion power | 175 |
| A.4 | General formulae | 178 |
| A.5 | Radial particle transport | 180 |
| A.5.1 | Classical diffusion | 180 |
| A.5.2 | Bohm diffusion | 181 |
| A.5.3 | Gyro-Bohm diffusion | 181 |
| A.5.4 | ETG-driven transport | 182 |
| A.6 | Mirror-specific derived quantities | 182 |
| A.6.1 | Temperatures and confinement time in a beam-heated mirror from Egedal et al 2022 [45] | 182 |
| A.6.2 | Confinement time given by classical transport | 184 |
| A.6.3 | End Cells/Plugs | 185 |
| A.6.4 | Tandem mirror — central cell | 188 |
| A.6.5 | Overall power balance and plant power estimates | 190 |
| A.6.6 | Instabilities | 191 |
| A.7 | Costs and economics | 192 |
| A.7.1 | Heating | 192 |
| A.7.2 | Magnets | 192 |
| A.8 | Optimization constraints | 193 |
| A.8.1 | Midplane fields regularization via alpha particle confinement penalties . . . | 193 |
| A.9 | Optimizing mirror configurations | 194 |

| | | |
|----------|--|------------|
| A.9.1 | Gradient descent using SymPy and JAX | 194 |
| A.9.2 | Example: optimizing Q in a simple mirror | 194 |
| A.10 | Conclusions | 195 |
| B | Wisdom acquired | 199 |
| | References | 200 |

LIST OF FIGURES

| | | |
|------|---|----|
| 1.1 | The mirror machine loss cone | 3 |
| 2.1 | Pictures of the LAPD | 11 |
| 2.2 | Picture of LAPD interior | 13 |
| 2.3 | Images of Langmuir probes used in the LAPD | 16 |
| 2.4 | Typical Langmuir probe I-V curve | 17 |
| 2.5 | Ceramic cap of a Bdot probe | 20 |
| 2.6 | Schematic of a drift wave | 24 |
| 2.7 | Schematic of an interchange mode | 25 |
| 3.1 | Cartoon of the mirror turbulence experiment setup | 34 |
| 3.2 | Raw data processing steps | 36 |
| 3.3 | Shot-to-shot variation of I_{sat} profiles | 39 |
| 3.4 | Midplane I_{sat} profile | 40 |
| 3.5 | Line-integrated density | 41 |
| 3.6 | T_e from Langmuir sweeps at the midplane | 42 |
| 3.7 | Plasma potential and derived $\mathbf{E} \times \mathbf{B}$ velocity profiles | 43 |
| 3.8 | I_{sat} and B_{\perp} fluctuation power profiles | 45 |
| 3.9 | Cross-field, $\tilde{E} \times B$ particle flux | 46 |
| 3.10 | Phase and coherency of I_{sat} current and V_f near x_{PF} | 47 |
| 3.11 | Swept vs triple probe measurements | 49 |
| 3.12 | Triple probe T_e and T_e fluctuation profiles | 49 |
| 3.13 | Diffusivity relative to D_B | 50 |

| | | |
|------|--|----|
| 3.14 | Diffusivity with T_e compensation relative to D_B | 51 |
| 3.15 | I_{sat} gradients under varying profile smoothing methods | 51 |
| 3.16 | I_{sat} (density) fluctuation power | 52 |
| 3.17 | k_{\parallel} and coherency γ | 53 |
| 3.18 | Normalized density fluctuations vs ρ_s/L_n | 55 |
| 3.19 | I_{sat} scaled by L_n^2/n^2 | 56 |
| 3.20 | Fluctuation power summed for each k_y | 57 |
| 3.21 | B_{\perp} fluctuation power spectra at various locations | 58 |
| 3.22 | B_{\perp} fluctuation power profiles for three regions | 60 |
| 3.23 | Summed fluctuation power of B_{\perp} in the core | 61 |
| 3.24 | B_{\perp} , flat field 500G vs 400G | 62 |
| 3.25 | B_{\perp} fluctuations at $x=0$ for different mirror lengths | 63 |
| 3.26 | B_z fluctuations in the core and x_{PF} | 64 |
| 3.27 | B_{\perp} fluctuation power profiles for low frequencies | 65 |
| 3.28 | B_z fluctuation power profiles for low frequencies | 66 |
| 3.29 | 2d plane of the perpendicular magnetic field and the derived current density for $M = 1$ | 67 |
| 3.30 | B_{\perp} and the derived current density for the $M = 1.9$ and $M = 2.68$ cases | 68 |
| 3.31 | Azimuthal mode number m amplitudes of I_{sat} fluctuations | 70 |
| 3.32 | k_y averaged about x_{PF} | 71 |
| 3.33 | Gradient scale length L_n | 75 |
| 3.34 | Particle flux: breakdown into components | 76 |
| 3.35 | I_{sat} decorrelation time | 77 |
| 3.36 | ExB shearing rate | 78 |

| | | |
|------|---|-----|
| 3.37 | Discharge power vs mirror ratio | 79 |
| 3.38 | I_{sat} profiles ($M=1$), DR1 vs DR2 | 80 |
| 3.39 | Line integrated density, DR1 vs DR2 | 81 |
| 3.40 | I_{sat} calibration factor over runs | 82 |
| 4.1 | A demonstration of Latin-hypercube sampling vs random sampling | 86 |
| 4.2 | LAPD configuration and diagnostics for the ML datarun | 89 |
| 4.3 | Example machine state information and diagnostic signals | 90 |
| 4.4 | Example interferometer skips | 92 |
| 4.5 | Cleaning Thomson scattering spectra | 93 |
| 4.6 | Time-averaged I_{sat} distribution over shots | 94 |
| 4.7 | Distribution of probe x-coordinates in the dataset | 95 |
| 4.8 | y-axis profile before and after shifting the y-coordinate | 97 |
| 5.1 | Cartoon of the experiment setup | 102 |
| 5.2 | I_{sat} traces from the swept probe | 103 |
| 5.3 | Gas puff timings and example I_{sat} time series at three different z-axis locations from three different dataruns. Note that some discharges do not achieve steady state in I_{sat} . . | 104 |
| 5.4 | Examples of I_{sat} profiles from DR2 | 105 |
| 5.5 | Training and validation losses when overfitting the model | 107 |
| 5.6 | Linear predictions of I_{sat} profiles | 109 |
| 5.7 | The loss and MSE for the training, validation, and test sets | 114 |
| 5.8 | Cross-validation test set error | 116 |
| 5.9 | Cross-validation training set error | 116 |
| 5.10 | Z-score distributions for the training and test sets for $\lambda = 0$ | 117 |

| | | |
|------|---|-----|
| 5.11 | Model performance and calibration for different weight decays | 118 |
| 5.12 | Model extrapolation performance and uncertainty | 119 |
| 5.13 | Mirror configuration predictions scaled to the cathode radius | 121 |
| 5.14 | Predictions compared with validation dataruns | 123 |
| 5.15 | Comparison of original DR2 profiles with the profiles from DR0pt | 125 |
| 5.16 | Comparison of I_{sat} predictions with Thomson scattering measurements | 126 |
| 5.17 | Discharge voltage scan: effect on x and z profiles | 128 |
| 5.18 | Predictions of axial gradient | 129 |
| 5.19 | Optimized axial profiles, predicted and measured | 133 |
| 6.1 | EBM training curves | 148 |
| 6.2 | Distribution of batches in replay buffer samples | 149 |
| 6.3 | EBM neural network architecture | 151 |
| 6.4 | Layer blocks used in the EBM architecture. Residual layers work well. | 152 |
| 6.5 | MCMC energies, gradients, and integrated trajectory length for unconditional samples | 154 |
| 6.6 | Unconditional samples – diode 3 and mirror field | 155 |
| 6.7 | Unconditional samples – all inputs | 156 |
| 6.8 | Histograms of scaled values and energy for unconditional samples | 157 |
| 6.9 | Interferometer time series reconstruction | 160 |
| 6.10 | Energy function scan along probe x coordinate | 161 |
| A.1 | DT reactivities | 175 |
| A.2 | DD reactivities | 176 |
| A.3 | DT, DD, and D-He3 reactivity comparison | 176 |
| A.4 | Simple mirror: cost function and gradient magnitude for each step | 196 |

| | | |
|-----|---|-----|
| A.5 | Simple mirror: magnetic fields for each step | 197 |
| A.6 | Simple mirror: fusion power and Q for each step | 197 |
| A.7 | Simple mirror: temperatures for each step | 198 |

LIST OF TABLES

| | | |
|-----|--|-----|
| 3.1 | Magnetic mirror lengths and ratios | 34 |
| 3.2 | LAPD plasma parameters | 37 |
| 3.3 | x_c and x_{PF} locations for each mirror ratio | 38 |
| 4.1 | Data breakdown by class and dataset (percent) | 96 |
| 5.1 | Summary of test set losses for different training data and ensembles | 110 |
| 5.2 | Machine inputs and actuators for model inference | 131 |
| 5.3 | Machine inputs and actuators for optimized axial profiles | 134 |
| 6.1 | RMSE and 2σ of conditional interferometer samples for test set and DR2_02 | 158 |

ACKNOWLEDGMENTS

I am fortunate to have many to thank for making my PhD such an enjoyable and satisfying part of my life. Shout out to my homies Gurleen Bal, Kamil Sklodowski, and Yhoshua Wug for providing much-needed support and camaraderie over all these years – graduate school would have been very different (maybe painful!) without you guys.

I would also like to extend gratitude to all those in the nfmfmhdut community and the SULI mafia for their support and friendship, particularly Becky Masline and Shawn Zamperini for building a community for us over covid. In no particular order: Kunal Sanwalka, Sam Frank, Nick McGreivy, Ben Israeli, Garrett Prechel, Andrew Maris, Allen Wang, Kenny Gage, Ryan Chaban, Mason Yu, Tyler Cote, Genevieve DeGrandchamp, Alex Leviness, Nirbhav Chopra (UIUC gang), Oak Nelson. If I forgot anyone, I apologize, and be sure to grab me at the next APS for a drink.

Thank you to my labmates, classmates, and collaborators for fostering a warm community: Jeff Robertson, Mel Abler, Tom Look, Javi Serrano, Jesus Perez, Josh Larson, Henry Wong, Erik Everson, Tim Vanhoomissen, Ryan Buckley, Sarah Chase, Tri Luong, Lukas Rovige, Yuchen Qian, Bobby Dorst, Jess Pilgram, Jia Han, Seth Dorfman, Angelica Ottaviano, Patrick Crandall, Gary Wan, Jess Gonzalez, Graeme Sabiston, Mary Konopliv, Ani Thuppul, and Adam Collins. And in particular, shoutout to Gio Rossi for helping pass the comprehensive exam with a hilariously imbalanced competition and Kyle Callahan and Tom Neiser for being such great roommates at conferences. To those in the astronomy department, thank you for accepting a wayward physicist as one of your own (at least on the camping trips). Thanks to the lab staff for their willingness to help with running experiments and answering late-night texts: Shreekrishna Tripathi, Steve Vincena, Pat Pribyl, Zoltan Lucky, Marvin Drandell, and Tai Sy Ly. On behalf of all of us, I'd like to thank lab mom Meg Murphy for the administrative and emotional support (and food!) over all these years – it's hard to imagine what the lab would be like without her.

Friends and family, thank you for providing support and homes to crash at when I needed to: Lydia and Andy Wilson (and Savannah and Shepard), Tim and Allison Travis, Sean Lyon, Claire

Baum, Erik Wessel, Rory Bentley (spacebros), Ada Morral, Antonett Prado, and most of all Piper (Bernese) and Blake (terrier mix). And anti-collaborator Altair (Golden-Bernese mix), thank you for keeping me sane in all the important ways and driving me nuts in all the others.

Thank you to professors Walter Gekelman, George Morales, and Derek Schaeffer for your instruction and insight. To committee members, Jacob Bortnik and Paulo Alves, thank you for your insight and support in my exploration of machine learning techniques. And thank you to Chris Niemann for answering my late-night texts when troubleshooting the Thomson scattering system. Lastly, and most importantly, I'd like to thank my advisor, Troy Carter, for supporting me despite my wanderings in ML territory. I am certain that the work I have done here would not have been possible anywhere else under any other advisor, and for that I am immeasurably grateful.

This work was performed at the Basic Plasma Science Facility, which is a DOE Office of Science, FES collaborative user facility and is funded by DOE (DE-FC02-07ER54918) and the National Science Foundation (NSF-PHY 1036140). Thank you US taxpayers.

Chapter 3 is based on the following publication: P. Travis, T. Carter, *Turbulence and transport in mirror geometries in the Large Plasma Device*, Journal of Plasma Physics. 91(1):E40 (2025). doi:10.1017/S0022377825000029.

Chapters 4 and 5 are based on a publication currently under review: P. Travis, J. Bortnik, T. Carter, *Machine-learned trends in mirror configurations in the Large Plasma Device*, Physics of Plasmas. (2025).

VITA

- 2017 B.S. in Engineering Physics, University of Illinois at Urbana-Champaign
- 2018 M.S. in Physics, University of California, Los Angeles
- 2017-2025 Graduate student researcher, University of California, Los Angeles

PUBLICATIONS AND PRESENTATIONS

P. Travis, T. Carter, *Turbulence and transport in mirror geometries in the Large Plasma Device*, Journal of Plasma Physics. 91(1):E40 (2025). doi:10.1017/S0022377825000029

Y. Qian, W. Gekelman, P. Pribyl, T. Sketchley, S. Tripathi, Z. Lucky, M. Drandell, S. Vincena, T. Look, P. Travis, T. Carter, G. Wan, M. Cattelan, G. Sabiston, A. Ottaviano, R. Wirz, *Design of the Lanthanum hexaboride based plasma source for the large plasma device at UCLA*, Rev. Sci. Instrum. 94, 085104 (2023). doi:10.1063/5.0152216

P. Travis, J. Bortnik, T. Carter, *An open dataset from the Large Plasma Device for machine learning and profile prediction*, poster presentation, 66th Annual Meeting of the APS Division of Plasma Physics, Atlanta, GA (2024)

P. Travis, T. Look, L. Rovige, C. Niemann, P. Pribyl, T. Carter, *Predicting profiles in LAPD mirror configurations*, poster presentation, 65th Annual Meeting of the APS Division of Plasma Physics, Denver, CO (2023)

P. Travis, S. Vincena, P. Pribyl, T. Carter, *Developing a generative ML model for LAPD trend inference and profile prediction*, 64th Annual Meeting of the APS Division of Plasma Physics, Spokane, WA (2022)

P. Travis, T. Carter, *Upgrading LAPD diagnostic pipelines for training generative ML models*, poster presentation, 63rd Annual Meeting of the APS Division of Plasma Physics, Pittsburg, PA (2021)

P. Travis, T. Carter, *Automated Langmuir sweep analysis using machine learning*, contributed talk, 62nd Annual Meeting of the APS Division of Plasma Physics, online, FL (2020)

P. Travis, T. Carter, *Study of turbulence and transport in magnetic mirror geometries in the LAPD*, poster presentation, Transport Task Force, San Diego, CA (2018)

CHAPTER 1

Introduction

Nuclear fusion is potentially a fantastic long-term energy source. On Earth, the fuel is abundant, the waste manageable, and is otherwise clean. The most easily accessible fusion reaction for us is the fusion of two heavy hydrogen isotopes: deuterium and tritium, with one and two neutrons, respectively. This reaction produces a neutron with 14.1 MeV of energy and a helium-4 nucleus with 3.5 MeV. In terms of mass of the reactants, the energy density of fusion is roughly seven orders of magnitude higher than hydrocarbon fuels. Reaching the conditions for this fusion reaction requires of the plasma sufficiently high density n , temperature T , and confinement time τ . Since the inception of the controlled fusion program, progress has been made towards improving these parameters [166], but much work is left to be done to put the power of the sun in the palm of our hands.

1.1 Nuclear fusion via magnetic confinement

The current dominant configuration for fusion plasmas is the donut-shaped tokamak [82]. This device operates by first imposing a toroidal field with external magnets, like bending a solenoid into a circle. Toroidal devices make intuitive sense: the magnetic field can only confine plasmas perpendicular to the field via the Lorentz force $F = q(E + v \times B)$, so the third axis is wrapped in a circle so that the magnetic field lines are closed. A current is induced in this toroidal direction to create an orthogonal poloidal field component, a net helical field, to cancel out particle drifts. Currently, tokamaks are in the lead for plasma performance in magnetic confinement fusion [166].

Tokamaks, however, have two major drawbacks: the large toroidal current and the complicated geometry. At minimum, a breakeven reactor requires many millions of Amperes of circulating current [29]. This current is a source of free energy that, in situations called disruptions, can be dumped into an electron beam that causes catastrophic damage to the wall as well as structural damage via induced currents. These disruptions are thought by some to be a showstopper for the tokamak program. These disruptions can be avoided altogether by instead building a toroidal device that has a (mostly) externally set field. Such a device is called a stellarator [18]. These stellarators are, structurally speaking, even more complicated than a tokamak. The requirement for purely externally defined fields and sufficiently good confinement properties leads to very complicated coil geometries that are difficult and expensive to assemble. Devices that are difficult and expensive to build will not be economically competitive – costs are very important to the viability of fusion reactors [140]. Geometrically simple linear devices, such as mirror machines, may be the best path forward despite intrinsically worse confinement.

The magnetic mirror approach to fusion [124] operates on the principle of conservation of the magnetic moment $\mu = \frac{W_{\perp}}{B}$. When a particle enters a region of higher field, W_{\perp} increases by conservation of μ . Given sufficiently high W_{\perp} relative to W_{\parallel} , by conservation of energy ($W_{\perp} + W_{\parallel} = \text{constant}$) W_{\parallel} must decrease and, if the parallel energy is sufficiently low, then particle stops and reverse direction. Thus, some particle trapping can be achieved by creating a solenoid with high-field magnets at each end. Some particles can escape, and if the loss rate is rapid enough, deplete a portion of the velocity space leading to a “loss cone” distribution which can then drive many severe instabilities. A schematic of this loss cone can be seen in fig. 1.1. The primary focus of magnetic mirror research has been on plugging this hole in the distribution function and mitigating the interchange instability [124].




Figure 1.1: Schematic of the velocity distribution loss cone in mirrors (for ions). Ions are lost for all perpendicular velocities below a certain threshold because of the ambipolar potential created by fast-escaping electrons.

1.2 Confinement degradation from instabilities and turbulence

All fusion plasmas are contained within a vacuum vessel and are hot relative to the vessel wall, which implies the existence of temperature and density (and thus pressure) gradients. These gradients provide free energy to drive instabilities, which in some cases cascade to smaller length scales and dissipate. This process is turbulence, and it occurs inside every fusion-relevant plasma that we know of. These turbulent eddies can be large and are the fundamental limit on cross-field transport for fusion reactors. Effort must be made to study and characterize this turbulence, and if possible, suppress it.

Many instabilities contribute to transport and turbulence in fusion devices. Instabilities that impact mirror machines will be briefly reviewed here. A detailed look at instabilities present in the LAPD can be found in chapters 2 and 3.

Drift waves are a common density gradient-driven instability, driven unstable by dissipation

[68, 74, 154]. These drift waves are commonly seen in magnetic confinement fusion devices given the minimal requirements of a density gradient and magnetic field. Other drift-like modes also exist, such as ion temperature gradient modes (ITG), trapped electron modes (TEM), and electron temperature gradient modes (ETG), summarized in Doyle 2007 [120].

Mirror machines have been primarily concerned with stabilization of flute-like interchange modes. These modes are driven by pressure gradients aligned with the magnetic field curvature vector [124, 50, 165]. These modes have been stabilized in a variety of ways, such as line-tying [52], using a good-curvature expander tank [134, 77], or non-axisymmetric magnetic field configurations [108]. Ballooning modes [44] may also be driven unstable in mirror machines, and are commonly seen in tokamaks [148, 28].

Mirror machine research has also focused on velocity space instabilities driven by the loss cone distribution. The drift-cyclotron loss cone instability (DCLC) [13, 49, 92] is the coupling of ion drift waves to the ion cyclotron motion, driven by the density gradient and loss cone distribution. This mode can be stabilized by filling the loss cone, typically with warm plasma, so it is not expected to see this mode in highly collisional plasmas such as those produced in the LAPD. The Alfvén ion cyclotron (AIC) instability is similar – Alfvén waves interact with ion cyclotron motion, scattering the ions and degrading confinement [23].

1.2.1 Importance of this turbulence and transport study

As enumerated above, many instabilities may be present in a mirror machine, and understanding how modes may couple and influence transport is critical to design and operation of fusion reactors. Using the LAPD, we seek to understand turbulence and transport at edge-like conditions of mirror machines, and attempt to observe the interaction of interchange modes with drift waves or other instabilities. This study also highlights the ability for the LAPD to operate with mirror configurations, which may be useful for future studies since mirrors are once again being considered for a fusion reactor [46, 51, 53]. In addition, characterization of modes present on the LAPD would give a

greater understanding in interpreting data in other LAPD studies. Distinguishing the modes present could illuminate promising directions for further studies on LAPD turbulence and transport.

The magnetic field of the LAPD was configured to create several mirror ratios and lengths, with a focus on primarily short mirror cases. These mirrors were diagnosed with Langmuir probes and magnetic probes. Spectra from these diagnostics were analyzed to determine the modes present and calculate cross-field particle flux. Additional data were collected to reconstruct the 2d structure of the modes present using correlation techniques. Drift-Alfvén waves were observed at 12 kHz and up with the peak having strong dependence on machine-averaged magnetic field, as expected. Lower frequency modes could not be confidently identified, and could be a mixture of drift waves, rotational interchange, a nonlinear instability, or the conducting wall mode. Particle flux measurements were unexpected: an increased mirror ratio led to a decrease in the particle flux. Magnetic curvature has a destabilizing effect, so an increased $E \times B$ particle flux was expected. Nevertheless the particle diffusivity estimate was consistent with Bohm diffusion. The decrease in particle flux could be explained in-part by the decreased gradient scale length caused by the larger plasma radius at higher mirror ratios. No evidence for the interchange mode was observed, likely because of many stabilization methods present.

Although the results are unexpected, they are promising when considering the cold edge of a fusion reactor. If the edge is in contact with a conducting surface, it will likely be stable and not contribute too greatly beyond typical Bohm (or turbulent) transport.

1.3 Accelerating research using machine learning

Machine learning, though a nascent field for quite some time, has exploded in popularity since the advent of deep neural networks trained on GPUs [94]. Machine learning is effectively curve fitting, though, as compute is becoming cheaper and GPUs more powerful, larger models can be trained with greater capabilities. Fundamentally, machine learning, and in particular deep learning, scale very well with data and dimensionality. Plasma devices and simulations are producing increasing

amounts of data which can be exploited using these machine learning techniques.

Recently, machine learning techniques have become increasingly used in the plasma physics and nuclear fusion community. One of the most common applications is building surrogate models for expensive simulations. Surrogates train on simulation output and, once learned, provide a fast way to predict in the domain covered by the training set. Surrogate modeling [149] has been applied to turbulence and transport computations [107, 99, 54, 70], profile prediction [109], global tokamak simulations [35], uncertainty quantification [168], and in inertial confinement fusion (ICF) [115, 7]. ICF in particular has had great success enhancing predictive capabilities by combining experimental data and theoretical models using machine learning [58, 76, 65].

Reinforcement learning has been used for trajectory control on tokamaks [146, 144, 145], notably on the TCV tokamak [32, 151] for exploring new magnetic configurations. Predicting instabilities, particularly disruptions, has also been a major application of machine learning in the field [130, 86, 117, 57, 110, 129, 132]. A comprehensive review of disruption prediction can be seen in Vega et al. [158].

Profile prediction in tokamaks has also been performed using machine learning. Electron density, temperature, and other quantities in tokamaks have been predicted [4, 36], and can be adaptable using reservoir NNs [79]. Temporal evolution of parameters has been successfully modeled using recurrent neural networks (RNNs)[24, 162, 146, 144].

Machine learning techniques can also be used to solve differential equations by parameterizing the solution using a neural network. These “physics-informed neural networks” (PINNs) have also been increasingly used in fusion [133, 8, 143].

A deeper review of machine learning with respect to profile prediction and generative modeling can be found in chapters 5 and 6, respectively.

1.3.1 Importance of these ML models

The study of plasmas is characterized by inconsistent experiments, incomplete and untrustworthy diagnostics, and difficult theory. Machine learning provides a new tool to tackle these challenges. As demonstrated in the referenced works above, machine learning can be used to great effect to accelerate simulations, control plasmas, and make predictions. In order to successfully accomplish these tasks, these models learn some structure (such as a lower-dimensional manifold) over the data, but this structure is usually exploited only for the downstream task, such as predicting a plasma parameter at some point in time or an actuator state. My goal is to make this structure explicit and explorable; the model has likely learned relationships that can improve our understanding of fusion plasmas. If a model learns relevant and useful relationships from data, then this information can be exploited without requiring painstaking manual data analysis. Exploiting this information without humans in the loop could dramatically accelerate the rate of progress in fusion science.

This work pushes the frontier of extracting insight from a plasma device using machine learning. A diverse set of machine configurations were collected on the Large Plasma Device (LAPD) by selecting machine settings via Latin hypercube sampling. This randomization of machine configurations is a first for magnetized plasma research and yielded a set of 44 randomized dataruns (67 dataruns in total) spanning a wide range of LAPD parameter space, totalling 132,000 discharges. The diversity of this dataset is unique for plasma physics research: usually most machine parameters are held constant and vary one or a few in a grid-like pattern, but in this work all parameters are changed simultaneously.

Using a small neural network (around 200,000 parameters), time-averaged ion saturation current (I_{sat}) can be predicted at virtually any point in the LAPD given the machine state. Information is extracted from this model by inferring trends: scanning along particular inputs and observing how the predicted I_{sat} changes. Trends inferred this way agree with intuition and arguments from geometry, demonstrating that the model has accurately learned the underlying relationships. In addition, this model is used to optimize the axial variation of I_{sat} in LAPD discharges. Axial

variation in density and temperature is an important problem affecting many physics studies on the LAPD; the ability to optimize the axial variation can alleviate this issue. Machine learning is required to perform this trend inference and optimization of LAPD plasmas: the dimensionality of the problem is too high for a conventional grid search. Fundamentally, this work demonstrates the new capability of searching a high-dimensional space on a plasma device that would be very difficult, if not impossible, to do without machine learning.

This work was extended to include additional diagnostics and time-series data, and was expanded to generative modeling. An energy-based model (EBM) was trained on these additional diagnostics, learning a joint distribution over the diagnostics and machine parameters. This model can be used to reconstruct diagnostic signals given any combination of inputs. In this case the interferometer signal was conditionally sampled on only machine parameters and compared the case where additional diagnostic signals were included. Additional unanalyzed and uncalibrated signals improved reconstruction performance, indicating that other diagnostics contain useful information even in a primitive form. This diagnostic reconstruction using generative modeling is a first for plasma science, and this reconstruction was performed by modifying the energy function in a novel way. In addition, the energy surface was directly examined by scanning along the probe x-axis dimension, demonstrating a novel ability to find symmetries in the data from the trained model. These EBMs provide a very flexible way of representing the data, and open up many possibilities for combining with additional datasets and simulations.

Fundamentally, the machine learning techniques demonstrated in this thesis provide a pathway towards automating fusion science, and demonstrates a new way of representing and analyzing data over a wide range of experimental conditions.

1.4 Outline of the dissertation

Chapter 2 describes the Large Plasma Device, the configuration of the machine, and the diagnostics used in this thesis. It also provides background on some relevant instabilities and a very brief

introduction to neural networks.

Chapter 3 details a study of mirror turbulence and transport using the older barium oxide cathode. Multiple mirror ratios and lengths were analyzed, with the conclusion being that cross-field particle flux reduced at higher mirror ratios, and no evidence of the interchange instability was seen.

Chapter 4 describes the process of collecting data from randomly-set LAPD configurations, the diagnostics collected, and the biases inherent in the data. The dataset collected here is used in the two subsequent chapters. The chapter also covers the data cleaning process for some of the diagnostics.

Chapter 5 is a thorough machine learning study on predicting ion saturation current – I_{sat} – in LAPD mirrors using neural networks. This is possibly the first time a global optimization like this has been performed in magnetized plasma device. In addition, trends were inferred by scanning along various combinations of machine inputs. The important features of this work are the validation of model predictions and thorough uncertainty quantification. Fundamentally, this is a demonstration that machine learning can be used to extract insights from data in a complicated, multivariate physical system.

Chapter 6 takes a different angle on machine learning than what was done in chapter 5. Instead of predicting a single output given machine inputs, we instead learn a probability distribution over the data using an energy-based model. Using this model, we can reconstruct any arbitrary combination of diagnostics, machine settings, or anything in the input space – including hallucinating discharges altogether. The flexibility of the learned energy function is demonstrated. This flexibility provides a way to combine theory and other experiments to improve predictions.

Appendix A describes a small 0d reactor optimization project in collaboration with Kunal Sanwalka, based on a spreadsheet by Cary Forest. This project was undertaken to learn how to use Jax and SymPy, and also gain some intuition on mirror-based fusion reactors.

CHAPTER 2

Background

2.1 The Large Plasma Device at UCLA

The Large Plasma Device (LAPD)[60, 126] is a 26 meter long linear plasma device located at the Basic Plasma Science Facility (BaPSF) at the University of California, Los Angeles. This device is built for basic plasma science and can create quiescent, and long-lived, i.e., longer than the timescales of interest, plasmas. The LAPD produces up to 18 m long, 1 m diameter plasmas. A cartoon of the LAPD and the canonical coordinate system can be seen in fig. 2.1.

2.1.1 Plasma source

Typically, plasmas are produced using a hot cathode and anode at the south end of the device. This cathode was originally barium oxide (BaO)-plated nickel [60], but was recently upgraded to a segmented lanthanum hexaboride (LaB₆) source [126]. The BaO cathode was 72 cm in diameter, which mapped to 60 cm in a flat magnetic field configuration. A 72 cm diameter, 50% transparent molybdenum anode was used to accelerate electrons from the cathode down the length of the machine. This BaO source could reach densities of $4 \times 10^{12} \text{ cm}^{-3}$ and temperatures up to 8 eV. The LaB₆ cathode is 35 cm across and electrons are accelerated through a 64.4 cm diameter, 66% transparent molybdenum anode. The LaB₆ cathode can form hotter, denser plasmas with densities up to $3 \times 10^{13} \text{ cm}^{-3}$ and temperatures up to 20 eV, though typical operation yields temperatures around 5 eV, and is also more robust to accidental atmospheric exposure. The LaB₆ cathode is heated to $\approx 1700^\circ\text{C}$ using a $\approx 2 \text{ kA}$ heater. Both of these cathodes were used in the work presented




Figure 2.1: Top: picture from the LAPD taken from the south end of the device, next to the cathode and anode region. Bottom: picture taken in the middle section closer to the north end. Right: cartoon of the LAPD with the coordinate system used.

here. An insertable, smaller (20 cm diameter) LaB_6 source also exists at the north end of the machine, but is not used in this work. An interior shot of the LAPD can be seen in fig. 2.2.

The voltage applied across the cathode and anode is supplied by a 4.2 Farad capacitor bank switched by group of IGBTs. The discharge voltage is configurable up to 180V before triggering the over-voltage protection, though the capacitors are rated up to 200V. Current through the cathode can exceed 10 kA. Discharges can last as long as 70 ms, though a typical duration is around 15-20 ms. Discharge duration, power, and repetition rate are governed by the size of the capacitor bank and the charging power supply. The discharge repetition rate is configurable between 0.1 and 1 Hz.

2.1.2 Magnetic field

The LAPD has 13 independently-configurable magnet power supplies to shape the geometry of the axial magnetic field. Two of the supplies control the source region field, one controls the north end field where the smaller LaB_6 source resides, and the remaining 10 supplies control the field of the main plasma column. The source field can reach up to 8 kG and the main plasma column field up to 1.6 kG. A 1 kG field leads to an ion gyroradius of 2 mm at 1 eV, and an electron gyroradius of 50 μm at 5 eV, so these plasmas are highly magnetized. The source field is set manually on the power supplies themselves, but the middle 10 fields can be programmed using the LabView housekeeping system.

2.1.3 Gas fueling

There are two main ways of providing the neutral gas necessary for producing plasmas: the static fill system and gas puffing. The static fill system utilizes mass flow controllers to fill the chamber to the desired pressure, usually between 10^{-5} and 5×10^{-4} Torr. The LAPD can be filled with a variety of (nonreactive) gases, helium being the most common, followed by hydrogen and argon. The gas puff system utilizes piezo valves to puff gas into the chamber halfway between the cathode and anode. Gas puff duration and valve voltage can be set, which influence the total amount of




Figure 2.2: The interior of the LAPD taken during a discharge from the end of the machine. The hot LaB_6 cathode can be seen at the far end. Inside the plasma chamber are, from left to right, a carbon iris (used in Josh Larson's experiments) for controlling the width of the plasma from the smaller north-end source, an electric dipole probe, and the traveling wave antenna. The pink glow is the plasma formed by the far, main LaB_6 source.

gas puffed and thus plasma density. Plasma breakdown using gas puffing is very reliable. Without fueling the LAPD has a base pressure of 5×10^{-7} Torr. The pressure is constantly monitored by the housekeeping system.

2.2 Data acquisition

Probes can sample virtually any point in this plasma through unique ball valves placed every ≈ 32 cm along the length of the device, enabling the collection of time series data with high spatial resolution. Four probe drives can simultaneously be used to move probes in the x-y plane. An x-y-z probe drive is also available for collecting volumetric data. There are also 3 permanently attached probe drives mounted 45° up from the -x axis on which Langmuir probes are currently mounted. These drives have a limited motion compared to the standard x-y probe drives used during dataruns and the signals are digitized on 8-bit oscilloscopes. Typically many shots are taken at one position to obtain good sample statistics.

Primarily, probe data acquisition is handled through the main data acquisition system, simply referred to as the “DAQ”. The DAQ consists of SIS 3302 digitizers (theoretically 32 channels total) capable of sampling signals between ± 2.5 V at 100 MHz at 16 bits. Typically sample averages are taken (16 samples for my data) to reduce data transfer and file size. The DAQ is set up through a LabView-based control system.

This LabView system manages “dataruns” which are a series of discharges with a particular LAPD configuration (including the DAQ, probe motion control, and other device configurations such as function generators). Another LabView system controls the probe movements. All these devices are enabled and commands issued in a particular order via the run sequencer.

Another LabView system manages the machine state information (MSI). This system collects information on the discharge (current, voltage), auxiliary diagnostic signals (and formerly interferometer traces), gas total pressure and RGA pressures. The time series diagnostics are read by a National Instruments PCIe-6346 data acquisition card which samples with 16 bits over 8 analog

inputs at 500 kS/s per input. The MSI also collects information on the state of the magnetic field and cathode heater from the housekeeping system. An auxiliary python system can also be used to collate diagnostics from multiple sources, including the machine state information.

2.3 Diagnostics

LAPD can field many different diagnostics and other equipment through the ball-valve ports as well as much larger box-shaped ports. The most common and useful diagnostics used are discussed.

2.3.1 Langmuir probes: I_{sat} , sweeps, triple probes

Langmuir probes (LPs) are a workhorse of diagnostics in low temperature plasmas, such as the LAPD, and are also used at the (lower temperature) edge of fusion devices. These probes can be used to measure density, temperature, and potential of the plasma. LPs are essentially a conductive tip (typically tungsten in our case) inserted into the plasma. Three different types of these LPs can be seen in fig. 2.3.

The LP can be biased to measure different portions of the current-voltage relationship. Common biasing schemes are: applying a strong negative bias, sweeping the bias along a set range, free-floating the tip, and biasing another tip via the ion saturation current from another tip. The I-V relationship of a LP and the quantities deduced from it can be seen in fig. 2.4 which will be detailed below.

2.3.1.1 Ion saturation current (I_{sat})

When the LP tip is biased negatively, the probe tips rejects all electrons and only collect ions, and is thus called the ion saturation current. This current, derived from the Bohm sheath criterion for cold ions ($T_i \ll T_e$), is:

$$I_{\text{sat}} = -e^{-1/2} S q_i n_i \sqrt{\frac{T_e}{m_i}} \quad (2.1)$$





Figure 2.3: Three different types of Langmuir probes used in the LAPD: flat protruding tips (top left), flat, flush tips (bottom left), and cylindrical tips (right). Based on measurements using a camera and macro lens, probe areas are often different from what is assumed.

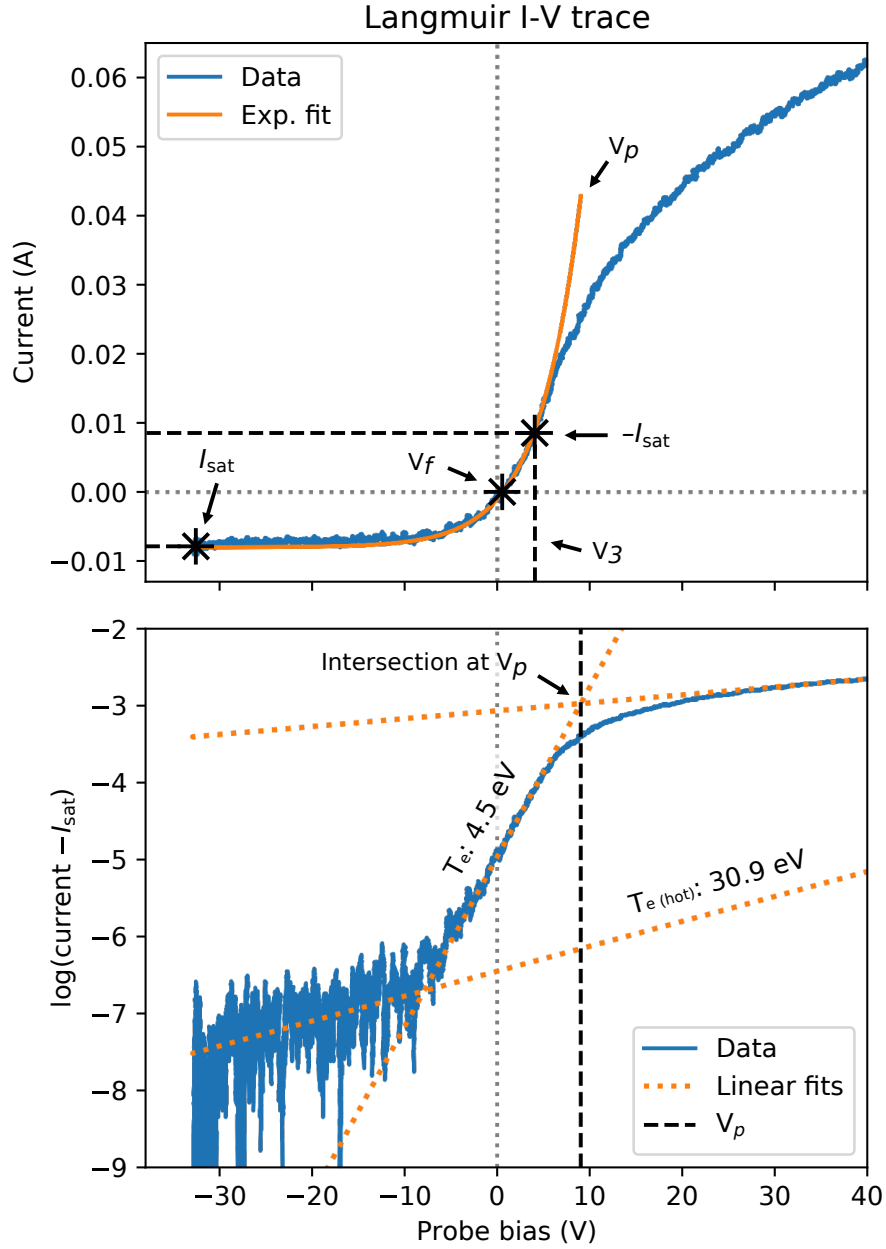


Figure 2.4: Typical Langmuir probe I-V curve on top, with a logarithmically-scaled version on the bottom (with I_{sat} subtracted). The log-scaled version is typically used for fitting the temperature(s). Useful points on the curve are labeled.

where S is the effective area, q_i is the charge of the ion, and n_i is the ion density, T_e the electron temperature, and m_i the ion mass. For a probe tip area of 1 mm^2 , a current of 1 mA corresponds to $n_e \approx 1\text{--}2 \times 10^{12} \text{ cm}^{-3}$ for a T_e from 4 to 1 eV. The I_{sat} on the I-V curve can be seen on the left in the top plot of fig. 2.4. A bias that is too negative can cause arcing among tips on the same probe shaft. A typical bias is around -60 V .

2.3.1.2 Floating potential (V_f)

The floating potential (V_f) is the voltage where the probe tip has zero net current, often accomplished by placing a large-valued resistor near the probe tip. This V_f is useful for measuring electrostatic fluctuations, and can also be used as a proxy for plasma potential V_p if the electron temperature T_e is known via the relationship

$$V_f = V_p - \frac{1}{2} T_e \ln \left(\frac{2m_i}{\pi m_e} \right) \quad (2.2)$$

In addition, two vertically-arranged V_f tips can be used to approximately measure the local electric field by dividing the potential difference by the distance between the two probes: $E = (V_{f,\text{top}} - V_{f,\text{bottom}})/\Delta x$, where Δx is the displacement between the probes. This electric field is useful for calculating the local $E \times B$ particle flux.

2.3.1.3 Sweeps for T_e

Sweeping the bias applied to a LP tip (within a reasonable range) yields the exponential portion of I-V curve, seen at the top of fig. 2.4. Below the plasma potential, the ion contribution to the current is I_{sat} (eq. 2.1). The electron contribution is (assuming a Maxwellian):

$$I_e(V_B) = S q_e n_e \sqrt{\frac{T_e}{2\pi m_e}} \exp \left(\frac{-q_e(V_p - V_B)}{T_e} \right) \quad (2.3)$$

where q_e is the electron (elementary) charge, n_e the electron density, m_e the electron mass, and V_B the applied bias.

Electron temperature can be determined by fitting the exponential portion, which is typically

done using the I_{sat} -subtracted log-scaled plot seen in the bottom of fig. 2.4. The electron temperature is then the reciprocal of the fitted slope, in this case 4.5 eV. In the LAPD, given the electron beam used to break down the plasma, the electron distribution can have a hot tail which can be fit with a separate curve, which yields 30.9 eV in this case. The electron saturation current region – the portion of the sweep around $V_B \approx V_p$ and higher – can be also be fit. The intersection of this line and the temperature fit line is typically considered to occur at the plasma potential V_p .

Plasma potential is useful for calculating the azimuthal velocity profile. The radial electric field can be calculated using a finite differences along a radius, and using the background electric field, the azimuthal $E \times B$ velocity can be calculated. From this the shearing rate can also be calculated.

2.3.1.4 Triple probe T_e measurements

Time-resolved T_e measurements can be obtained by measuring the difference between a probe tip biased by the I_{sat} from another probe, called V_3 here, and the floating potential. This process effectively measures T_e using three points on the LP I-V curve. These three points (hence, triple probe) can be seen in the top plot of fig. 2.4. At V_f , the net current is zero:

$$0 = Sq_e n_e \sqrt{\frac{T_e}{2\pi m_e}} \exp\left(\frac{-q_e(V_p - V_f)}{T_e}\right) + I_{\text{sat}} \quad (2.4)$$

At V_3 , the current is $-I_{\text{sat}}$:

$$-I_{\text{sat}} = Sq_e n_e \sqrt{\frac{T_e}{2\pi m_e}} \exp\left(\frac{-q_e(V_p - V_3)}{T_e}\right) + I_{\text{sat}} \quad (2.5)$$

(note that I_{sat} is a negative quantity). Combining these two equations and, through some algebra, one obtains T_e as a function of V_f and V_3 :

$$T_e = \frac{e(V_3 - V_f)}{\ln(2)} \quad (2.6)$$

Given the sensitivity to fluctuations, sweeps are typically more reliable than these triple probe T_e measurements, but they can be relatively accurate. In this case, using the three points from fig. 2.4 yields $T_e = 4.48$ eV which is within measurement error of the swept measurement of 4.54 eV. Triple



Figure 2.5: Ceramic cap of a Bdot probe which is inserted into the LAPD to measure magnetic fluctuations

probe measurements require much less post-processing effort than swept measurements because of the difficulty of automating sweep fits.

2.3.2 Magnetic flux (Bdot) probes

Magnetic measurements are performed using magnetic flux probes, typically called “Bdots”. Using Faraday’s law, a changing magnetic field induces an EMF: $V = -A_{\text{eff}} \frac{dB}{dt}$, where A_{eff} is the effective area of the probe which depends on the area of the loop(s) and the number of turns. These Bdots are formed from three orthogonal loops to capture magnetic fluctuations along all axes. These loops are differentially wound and amplified so that electrostatic effects are removed. The wire is coiled on a ceramic tube which is then covered with a ceramic cap, which can be seen in fig. 2.5. These probes are calibrated using a Helmholtz coil to measure the spectral response and crosstalk using a network analyzer. For all the signals used in this work, the response of the probe is linear (and remains linear until well into the \sim MHz regime). The construction, calibration, and operation of these probes is covered by Everson et al. [47].

Converting these Bdot signals, after calibration, simply requires integrating $\frac{dB(t)}{dt}$ over time,

which is easily accomplished in the frequency domain and avoids accumulating errors. Integrating over t in the frequency domain simply requires dividing by $-i\omega$:

$$\int \dot{B}(t)dt = \int \sum_{\omega} b_{\omega} e^{-i\omega t} dt = \sum_{\omega} \frac{b_{\omega}}{-i\omega} e^{-i\omega t} \quad (2.7)$$

where $\sum_{\omega} b_{\omega} e^{-i\omega t}$ is the Fourier decomposition of $\dot{B}(t)$

2.3.3 The 288 GHz heterodyne interferometer

The LAPD is equipped with two 288 GHz heterodyne interferometers which measure line-integrated density. These interferometers function by applying a voltage shift to the source of the 288 GHz signal (96 GHz, tripled) using a sawtooth waveform at 750 kHz, leading frequency changes in the range of tens of MHz. Density fluctuations measurable by the interferometer occur much slower than this 750 kHz sweep frequency. When this wave is launched into the LAPD and returned, this 750 KHz wave is phase-shifted relative to a 750 kHz reference caused by both the free-space propagation time and the plasma-induced path length increase. A higher-density plasma leads to an increased path length, which leads to a larger observed phase shift. This phase shift can then be unwrapped to obtain a density, given some calibration factor. The reference and detector signals are recorded by an oscilloscope, which can be read from LabView (as part of the MSI system) or python, or written to a file.

2.3.4 Thomson scattering

Thomson scattering is the gold standard of temperature measurements. In the LAPD, the Thomson scattering has been implemented in the non-collective regime using a 532 nm, 460 mJ Nd:YAG laser to measure T_e and n_e [62]. Thomson scattering is the process of electrons scattering and Doppler shifting incident light, where the Doppler shift is caused by the electron thermal motion. This scattered light is then collected by a spectrometer where the Doppler shift is directly measured. The density can also be measured by counting the detected photons after absolute calibration of the

entire optical system, including scanning of the fiber coupling optics to maximize signal. Between shots, background spectra are taken and subtracted from the signal to reduce the influence of optical imperfections and spurious light. A notch filter is also used to block reflections from the 532 nm laser.

Analysis is fairly straightforward: the spectrum can be fit assuming a Gaussian distribution if the plasma distribution is Maxwellian. The temperature is determined by the width of the distribution: $T_e = 0.4513 \cdot \Delta\lambda_{1/e}^2$, where $\lambda_{1/e}$ is the half width where the spectrum reduces by $1/e$ [62]. The density can be calculated by integrating this fit distribution, but without a recent and accurate absolute calibration this value is not reliable and should not be used.

The Thomson scattering system benefits from higher density (above 10^{13} cm^{-3}) and lower temperatures, otherwise many shots (a few thousand) may be required to get a spectrum with a desirable fitting accuracy. Presently, this system measures temperature on-axis at port 32 at a single point in time (span of 4 ns) in the plasma. The measurement volume can be moved to elsewhere in the device with great effort. Coincidentally, a Helium ion line ($7 \rightarrow 4$ transition) at 541 nm is occasionally visible on the spectrum, but the resolution of the spectrometer (0.28 nm) is not sufficient to calculate a thermal Doppler shift for reasonable ion temperatures, but it could provide an upper bound.

2.3.5 Fast framing camera

The facility currently has a Phantom v7.3 fast framing camera. At 800 by 600 resolution, it can record at 6,688 frames per second, but can record at over 35,000 frames per second at a 256 by 256 resolution. The camera sensor is 14 bit monochrome and capable of resetting pixels and reacquiring light on event of sensor saturation (the “extreme dynamic range” feature). Although this footage is not analyzed in a quantitative way, it is nonetheless useful for building intuition on the structures and dynamics in the device. The light collected by the camera has been shown to correlate fairly well with I_{sat} , at least in the plasma bulk (according to unpublished data by Daniel Guice).

2.4 Potential instabilities in the LAPD

A brief overview of potential instabilities is given here. The applicability of and evidence for the existence of these instabilities is further discussed in chapter 3.

2.4.1 Drift waves

Drift waves exist in any magnetized plasma that has a density gradient, and are driven unstable in the presence of dissipation (such as finite resistivity) [154]. For this reason, these drift waves are often called a “universal” wave or instability. A schematic of the wave can be seen in fig. 2.6. The mechanism behind the wave is as follows: assuming the plasma follows a Boltzman relationship, that is $\delta n/n \approx e\delta\phi/T$, a density perturbation results in a potential perturbation. This resulting electric field creates an $E \times B$ force that push portions of the perturbation out and other portions inward. These forces lead to a propagating wave. If there is dissipation in the electron response, it leads to a phase shift ε such that the density and potential fluctuations are related by $\delta n/n \approx e\delta\phi/T(1 + i\varepsilon)$. This phase shift leads to a growing mode – an instability.

The currents created by the electron response can couple to an Alfvén wave, leading to drift-Alfvén waves which have been observed in the LAPD in the past [103, 159], and are discussed further in chapter 3.

2.4.2 Curvature interchange

Curvature-induced interchange modes occur in the presence of a pressure gradient in line with the curvature vector. A schematic of these modes can be seen in fig. 2.7. The curvature drift (and in many cases, ∇B drift) lead to charge separation in the plasma. If there is a positive density perturbation, there will be a large current. A nearby density depletion will have a similar current, but at a lower level. On average, this leads to an electric field which is then reinforced by the $E \times B$ force. This interchange instability was a major challenge for early mirrors, and resulted in




Figure 2.6: Schematic of the drift wave. A density perturbation is propagated by the resulting $E \times B$ force, where the electric field is created by the Boltzman relationship




Figure 2.7: Schematic of the interchange mode. A density perturbation, combined with charge separation, leads to a growing mode.

non-axisymmetric configurations to stabilize the mode [124]. The large amplitude of the mode would make it very obvious in a density fluctuation measurement. These modes are also called “flute modes” in mirror machine literature.

2.4.3 Rotational interchange

Similar to drift and curvature-induced modes, the rotational interchange mode is caused by charge separation resulting from the curvature force. Like the interchange mode, these can be flute modes, but can also have $n \neq 0$ and could couple to other instabilities as well [56].

2.4.4 Nonlinear instability and conducting wall mode

Other instabilities may be present in the LAPD as well. Simulations and statistical comparison with experiments suggest the presence of a nonlinear instability with the largest growth rate [55]. These modes are flute-like, but can three-wave couple to non-flute modes. This instability has not been directly observed, but the statistics of fluctuations are consistent with experiment. Simulations also suggest the possible presence of temperature gradient- and conducting wall-driven modes (CWM) [55, 16]. This mode has not been directly observed on the LAPD, but nonetheless should be considered when evaluating fluctuation spectra given the presence of strong gradients and at least one conducting boundary.

2.5 Machine learning and neural networks

Fundamentally, neural networks (NNs) are a function that, given sufficient capacity, can represent any function – this is the “universal approximation theorem”. Stated succinctly, NNs are flexible curve fits. As we will see in this work, these curve fits can be very useful.

2.5.1 Fundamentals of neural networks

Neural networks are built one ‘layer’ at a time. A layer in an NN is a vector \vec{x}_i multiplied by some weight matrix \mathbf{W} followed by a nonlinearity to get the activation x_{i+1} :

$$\vec{x}_{i+1} = f(\mathbf{W}\vec{x}_i) \quad (2.8)$$

where \vec{x}_i is the previous activation (or input). A bias b (a learnable offset) is often concatenated to the input vector beforehand. The activation function f can be any nonlinearity. Typical choices for f include the rectified linear unit – ReLU (zero when the input is less than zero, linear otherwise), tanh, sigmoids, sigmoid linear unit (SiLU), and so on. ReLU, and a variant of which called Leaky ReLU, are popular choices given the simplicity of computation.

Stacking these layers one after another leads to the ability to express very complex curves. The

true innovation here is the ability to express very high dimensional curves which was otherwise impossible until the advent of modern NN techniques, enabled by fast matrix multiplication hardware (read: GPUs).

The weights of these layers are trained via gradient descent of some loss function \mathcal{L} . The loss is effectively the penalty for the network predicting the incorrect answer. A typical loss function for regression tasks is the mean-squared error:

$$\mathcal{L} = \frac{1}{M} \sum_{i=1}^M (x_i - g(x_i))^2 \quad (2.9)$$

where $g(x)$ is the output of the neural network and x_i is a training example in a batch size M . Training over a dataset is often done in batches to reduce the memory required and to introduce stochasticity into the training process which can improve generalization.

Gradient descent is the process of modifying the NN weights by calculating the gradient of the loss function with respect to those weights:

$$\theta \leftarrow \theta - \nabla_{\theta} \mathcal{L} \quad (2.10)$$

Advanced purpose-built algorithms exist to do this gradient descent in a fast manner, the most popular being Adaptive Moment Estimation, known as Adam [88].

2.5.2 Common layer types

In this work, we use three types of NN layers. First is the dense layer, also referred to as a “fully-connected” layers or “multilayer perceptron” if the entire model uses dense layers. In this layer every input is connected to every output. This fully-connected topology can lead to very large parameter counts if these layers are repeatedly stacked with a large width.

Second: the convolutional layer, or convolutional neural network (CNN) [96]. This layer scans along the input with a smaller NN. The input dimension of this layer is the “kernel size”, and the different networks scanning across the input are the “filters”. These CNNs have been used to great effect in image processing and time series analysis and are relatively parameter-efficient.

Third: the attention layers. These layers have three inputs called the query, key, and value. When these are all the same the mechanism is called “self-attention”. The query and key are matrix-multiplied and converted to a probability distribution via a softmax function. This softmax result is then applied to the value, yielding scaled dot product attention [157]:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad (2.11)$$

where d_k is the dimension of the query and key. The result of this combination is mask that selects important parts of the value vector. Stacking many of these layers, with multiple self-attention mechanisms for the same input vectors, creates a “transformer” [157] which is the basis for the recent advancements in large language models.

2.5.3 Generative models

Generative models are models that learn a joint probability distribution over the inputs X and outputs Y , $p(X, Y)$ instead of learning the conditional distribution $p(y|X = x)$ which are “discriminative” models. Models that produce a single point estimate are also considered to be discriminative. The learned joint distribution can typically be sampled both unconditionally (generating realistic-looking data) and conditionally to generate samples when given additional information. A brief overview of the commonly used generative models in fusion will be given here. Common generative architectures are variational autoencoders (VAEs), generative adversarial networks (GANs), diffusion models, autoregressive models, and energy-based models (EBMs).

A VAE [89] is composed of an encoder network and a decoder network. The model is trained by mapping the encoder output to a latent probability distribution. This probability distribution can then be sampled by the decoder. The model is trained by minimizing the difference between the data distribution and the sampled distribution (such as using the Kullback-Leibler divergence) and also minimizing the reconstruction error.

A GAN [64] is composed of an explicit generator network that generates plausible samples, and a discriminator network that distinguishes between generated samples and data. The generator is

trained on how well it fools the discriminator, and the discriminator is trained on how well it can discriminate between generated samples and real data. These two optimizations compete.

Diffusion models [71] differently from GANs and VAEs. The diffusion model functions by iteratively adding noise to an image such that the image converges to a Gaussian. An NN is then trained to invert this noising process, enabling sampling from a distribution from pure noise.

Autoregressive models, usually for time-series data or language modeling, predict the next item in a sequence given an earlier (set of) items. This model is the main idea behind training of large language models – so called “generative pretrained transformers”, or “GPT” [128].

Energy-based models parameterize the probability distribution via an energy function where the probability is defined as $p = \exp(-E)/Z$, where Z is the partition function. These are trained by evaluating the negative log-likelihood of the probability distribution via contrastive divergence [69, 135, 150]:

$$-\log p(x|\theta) = E(x|\theta) - \log Z_\theta \quad (2.12)$$

where θ are the parameters of the model. Taking the gradient with respect to model parameters yields

$$-\frac{\partial \log p(x|\theta)}{\partial \theta} = \frac{\partial E(x|\theta)}{\partial \theta} \Big|_{x \sim p} - \frac{\partial E(x|\theta)}{\partial \theta} \Big|_{x \sim q} \quad (2.13)$$

where the terms on the right hand side are averaged over samples drawn from the data distribution p and model distribution q , respectively. These are called the “positive” and “negative” samples, respectively.

CHAPTER 3

Turbulence and transport in mirror geometries in the Large Plasma Device

In this chapter we study turbulence and cross-field particle transport in LAPD mirror configurations. Mirror machines are once again rising in prominence as a candidate for commercial fusion reactors with the advent of highly-funded commercial ventures and high-field high-temperature superconducting magnets [46, 51], so development of a functional understanding of cross-field transport in mirrors is imperative. Using the LAPD, multiple mirror ratios from $M = 1$ to $M = 2.68$ and three mirror-cell lengths from $L = 3.51\text{m}$ to $L = 10.86\text{m}$ were examined. Langmuir and magnetic probes were used to measure profiles of density, temperature, potential, and magnetic field. The fluctuation-driven $\tilde{E} \times B$ particle flux was calculated from these quantities. Two probe correlation techniques were used to infer wavenumbers and two-dimensional structure. Cross-field particle flux and density fluctuation power decreased with increased mirror ratio. Core density and temperatures remain similar with mirror ratio, but radial line-integrated density increased. The physical expansion of the plasma in the mirror cell by using a higher field in the source region may have led to reduced density fluctuation power through the increased gradient scale length. This increased scale length reduced the growth rate and saturation level of rotational interchange and drift-like instabilities. Despite the introduction of magnetic curvature, no evidence of mirror driven instabilities — interchange, velocity space, or otherwise — were observed. For curvature-induced interchange, many possible stabilization mechanisms were present, suppressing the visibility of the instability. This chapter is based on my 2025 publication in the Journal of Plasma Physics titled “Turbulence and transport in mirror geometries in the Large Plasma Device” [3], with some additions and changes.

3.1 Introduction

Historically, mirror research has prioritized the main issues with mirror confinement: stabilizing the interchange instability, stabilizing velocity-space (loss-cone-driven) instabilities, and minimizing axial electron heat losses. Nevertheless cross-field transport remains an important topic in magnetic-confinement fusion reactor development, in both linear and toroidal geometries. Insight into edge-relevant turbulence and its coupling to interchange and other mirror-driven instabilities, performed in a basic plasma science device, may be useful for a mirror-based reactor. Although not at fusion-relevant core temperatures or densities, the Large Plasma Device (LAPD) operates at conditions similar to the edge of fusion devices and can provide insight into the physical processes in that region. A characterization of edge fluctuations has been undertaken, with emphasis on interpreting these fluctuations within the context of a mirror.

Non-classical cross-field particle transport is often caused by low-frequency, large-amplitude fluctuations. These fluctuations are the result of various instabilities. One such process is the “universal” drift instability, which appears in the presence of a density gradient and a dissipation mechanism, such as resistivity. Drift wave turbulence and the effect on transport has been extensively studied in the past [74, 154]. In the presence of sufficiently high rotation or sheared flow, rotational interchange and the Kelvin-Helmholtz instabilities also contribute or couple to these fluctuations.

Various gradient-, rotation-, and shear-driven instabilities (and suppression of such) have been studied previously in the LAPD experimentally [137, 138, 139] and in simulations using BOUT, a 3d fluid turbulence code, and an eigenvalue solver [122]. The LAPD has a sufficiently high spontaneous rotation rate that rotation-driven instabilities may be excited without artificial drive. Simulations using BOUT++ [55] have also suggested that a rapidly growing nonlinear instability may dominate over all other linear instabilities.

Imposing a magnetic mirror configuration introduces magnetic curvature. The alignment of the curvature vector with a pressure gradient vector component causes the flute-like interchange instability if no stabilization mechanism is present. This interchange mode could couple to finite k_{\parallel}

drift waves. The coupling of drift waves to curvature-induced interchange modes has been studied in toroidal devices such as TORPEX [121, 48], where curvature was seen as the driving component for the unstable drift-interchange modes. Drift-like fluctuations have also been observed in the GAMMA-10 mirror [106, 167]. Flute-like modes and drift waves have been studied in other linear devices, such as Mirabelle [20], where the appearance of flute-like modes or drift waves were controlled by varying the field and limiter diameter.

The rotational interchange and curvature interchange can both be flute-like modes. Rotational interchange (also called the “centrifugal instability”) is driven by the aligned centrifugal force and pressure gradient vectors, but curvature-driven interchange is instead driven by magnetic curvature and is typically referred to as simply the “flute” or “interchange” instability. Rotational interchange [81] has been observed in the LAPD in the past [137, 139], and the curvature-driven interchange instability has been observed in many other mirror machines [165, 50, 124].

Biasing or modifying the electrical connection of the plasma with the end wall has proven to be an important actuator in many mirror machines such as TMX-U [72], GAMMA-10 [106], and GDT [10, 9, 15], and will be utilized on WHAM [46]. Active biasing was not attempted in this study, but the intrinsic rotation and strong electrical connection to the source region may provide a useful analog for edge biasing in other mirror machines.

Because the LAPD exhibits fully-developed, broadband turbulence, it is difficult to identify the presence of a particular instability by the linear properties. Nevertheless, the LAPD has good coverage of perpendicular spectra using correlation-plane techniques, and some measure of parallel spectra using the correlation between two axially-separated probes. A space-time spectral characterization of the many instabilities present in this low beta, moderate aspect ratio, gas-dynamic trap regime is attempted.

The goal of this study was to investigate the changes to turbulence and transport in LAPD mirror configurations. Of particular interest were the potential coupling of the interchange instability with drift waves or other modes, and the effect of the mirror geometry on cross-field particle flux. Presented is a characterization of the observed modes and the effect of introducing a mirror

geometry.

This chapter is organized as follows. Sec. 3.2 discusses the configuration of the LAPD and the diagnostics used. Sec. 3.3 covers the changes seen when imposing a magnetic mirror configuration on profiles, particle flux, drift waves, turbulence, and magnetic fluctuations. Sec. 3.4 explores the changes in 2d (x-y plane) structure. Sec. 3.5 discusses the active and expected instabilities and reasons for their modification. Sec. 3.6 summarizes the study and discusses the requirements for a deeper investigation.

3.2 The experiment and device configuration

3.2.1 The Large Plasma Device (LAPD)

In this study, the plasma was formed using an emissive, 72 cm diameter barium-oxide (BaO) cathode [61] (mapped to 60 cm in a flat field) and a 72 cm diameter, 50% transparent molybdenum anode that accelerate electrons across a configurable 40 – 70V potential; voltages of 60 and 63V were used in this study. The source has since been upgraded to a lanthanum hexaboride (LaB6) cathode [127] that enables access to higher-density, higher-temperature regimes, but all the data in this study are from plasma formed by a BaO cathode.

The flexible magnetic geometry of the LAPD was used to construct a variety of magnetic mirror configurations. The discharge current, fill pressure, and other machine parameters were held constant. The typical plasma parameters observed in this study can be seen in table 3.2. Data in several mirror ratios and lengths were collected (see table 3.1) but emphasis is placed on the short cell because the highest mirror ratio possible ($M = 2.68$) with a 500 Gauss midplane field could be accessed and probes were able to be placed outside of the mirror cell. An overview of the axial magnetic field for the the short mirror configurations and probe locations can be seen in fig. 3.1. 2- or 3-cell mirror configurations were also explored but are not examined in this study. All results presented below are from the short mirror cell configuration unless otherwise specified.




Figure 3.1: Cartoon of the Large Plasma Device and the coordinate system used. Only four of the eleven mirror configurations studied are plotted for clarity (mirrors of the same length have similar shapes and simply scale with mirror ratio). Diagnostic set varied by datarun; unlabeled diagnostics were used in both dataruns.

| Mirror length | Mirror ratios (M) | | | |
|-----------------|-----------------------|------|------|------|
| Flat | 1 | | | |
| 3.51 m (short) | 1.47 | 1.90 | 2.30 | 2.68 |
| 7.03 m (medium) | 1.49 | 1.98 | 2.46 | |
| 10.86 m (long) | 1.47 | 1.97 | 2.44 | |

Table 3.1: Magnetic mirror lengths and ratios. The lengths are measured where the curvature changes sign and the ratio is the maximum divided by the minimum. Approximately 3.5m must be added to the length if the good-curvature region is included. In the case of small asymmetries, the field strengths were averaged before calculation of the mirror ratio.

3.2.2 Diagnostics

All diagnostic signals were recorded with an effective sampling rate of 6.25 MHz (16-sample average at 100 MSPS) and a spatial resolution of 0.5 cm. When necessary, averaging over time is done in the approximate steady-state period of the plasma discharge (4.8 to 11.2 ms from the 1 kA trigger signal). Unless otherwise noted, all data presented will be from probes inside the mirror region ($z \approx 7\text{m}$). An example of a raw I_{sat} signal and processing steps can be seen in fig. 3.2. The raw signals are detrended by subtracting the mean across shots to obtain the fluctuations only. FFTs are then taken of these fluctuations for calculating power spectra and cross-correlated quantities. Large-amplitude plasma signals (with the largest contribution to transport) are seen at frequencies below 200 kHz and are thus the focus of this study. Plasma signals at frequencies above 200 kHz are very small and are dominated by noise (from electronics and instrumentation) and are thus ignored. Fluctuation power profiles can then be constructed.

The data presented were collected in two phases. The first phase (“datarun”), DR1, collected Langmuir probe (I_{sat} and V_f) and magnetic fluctuation (“Bdot”) [47] traces. 50 shots were taken at each position for every configuration. The second phase, DR2, was conducted with a similar set of diagnostics focused on temperature measurements (swept and triple probe) and 2d x-y structure. 15 shots were taken at each position, except for Langmuir sweeps with 64 shots. When appropriate, all data for each position were shot-averaged. “ I_{sat} ” will be used interchangeably with “density” and be presented with units of density (assuming a flat $T_e = 4.5\text{ eV}$ profile).

3.3 Mirror-induced changes

3.3.1 Profile modification

Because the field at the plasma source increases with M , the midplane plasma expands by a factor of \sqrt{M} . This bulging of the midplane is a direct result of magnetic flux conservation: $S_1 B_1 = S_2 B_2$, where S is the area and B is the field for two locations in the LAPD. Thus, the ratio of radii at




Figure 3.2: Raw data and basic processing steps for LAPD probe diagnostics as demonstrated by an I_{sat} trace from a DR1, $M = 1$ mirror at 26 cm. Data are truncated from 4.8 to 11.2 ms (a) and detrended (b). Power spectral density is calculated (c), and a power profiles can be constructed (d). The shaded regions are excluded from this analysis.

| | | | | |
|-----------------------------|--------------------------|-------------------------|----------------------|---------------------|
| Cathode radius (M=1) | x_c | 30 | | cm |
| Machine radius | R | 50 | | cm |
| Plasma length | L | ~ 17 | | m |
| Primary species | | He-4 1+ | | |
| Electron-helium mass ratio | | 1.37×10^{-4} | | |
| Neutral pressure | | $6 - 20 \times 10^{-5}$ | | Torr |
| Quantity | | Core | $x = x_{\text{PF}}$ | Unit |
| Density | n_e | 1.25×10^{12} | 0.6×10^{12} | cm^{-3} |
| Ion temperature | T_i | ~ 1 | — | eV |
| Electron temperature | T_e | 4 | 5 | eV |
| Beta (total) | β | 9×10^{-4} | 6×10^{-4} | |
| Midplane magnetic field | B_{mid} | 500 | — | G |
| Plasma freq | Ω_{pe} | 10 | 7.1 | GHz |
| Ion cyclotron freq | Ω_{ci} | 200 | — | kHz |
| Electron cyclotron freq | Ω_{ce} | 1.4 | — | GHz |
| Debye length | λ_D | 0.013 | 0.021 | mm |
| Electron skin depth | λ_e | 30 | 43 | mm |
| Ion gyroradius | λ_{ci} | 5.8 | — | mm |
| Electron gyroradius | λ_{ce} | 0.13 | 0.15 | mm |
| Ion thermal velocity | \bar{v}_i | 6.94 | — | km/s |
| Electron thermal velocity | \bar{v}_e | 1190 | 1330 | km/s |
| Sound speed | c_s | 13.0 | 13.9 | km/s |
| Alfvén speed | v_a | 446 – 1140 | –1620 | km/s |
| Ion sound radius | ρ_s | 65 | 69 | mm |
| Ion-ion collision freq | ν_{ii} | 730 | 380 | kHz |
| Electron-ion collision freq | ν_{ei} | 6.77 | 2.59 | MHz |
| Electron collision freq | ν_{ee} | 9.57 | 3.66 | MHz |
| Ion mean free path | $\lambda_{i,\text{mfp}}$ | 26 | 50 | mm |
| Electron mean free path | $\lambda_{e,\text{mfp}}$ | 175 | 512 | mm |
| Spitzer resistivity | η | 192 | 146 | $\mu\Omega\text{m}$ |

Table 3.2: LAPD machine information and plasma parameters in the core and peak-fluctuation region ($x = x_{\text{PF}}$) at the midplane in this study. Dashed quantities are assumed to be identical to core quantities.

| | | | | | |
|---------------|----|------|------|------|------|
| Mirror ratio | 1 | 1.47 | 1.90 | 2.30 | 2.68 |
| Scale factor | 1 | 1.21 | 1.38 | 1.52 | 1.64 |
| x_c (cm) | 30 | 36 | 41 | 45 | 49 |
| x_{PF} (cm) | 26 | 32 | 36 | 40 | 43 |

Table 3.3: x_c and x_{PF} locations for each mirror ratio when scaled by the expected magnetic expansion.

locations 2 and 1 is proportional to $\sqrt{B_1/B_2}$. The physical locations of the peak fluctuation region – x_{PF} (maximum gradient) – and the cathode radius x_c can be seen in tab. 3.3. This expansion leads to broader plasma profiles and decreased core density but are similar in the core and at x_{PF} when magnetically-mapped to the cathode radius x_c as seen in fig. 3.4. Dips between the core ($x/x_c = 0$) and the peak fluctuation region ($x = x_{PF}$) are seen, but fluctuation power from this region ($x/x_c = 0.5$ to 0.7) is not significant (fig. 3.8) so this region is not the focus of this study. The line-integrated density as measured by a 56 GHz heterodyne interferometer increases up to $\sim 35\%$ from the $M=1$ case of $\approx 8 \times 10^{13} \text{ cm}^{-2}$ (fig. 3.5) but does not increase past a mirror ratio of 2.3.

The error of the I_{sat} profiles as represented by the standard deviation (scaled by the time-averaged profiles) can be seen in fig. 3.3. The error is relatively small and should not play a factor in our analysis.

Discharge power increases only slightly (3%) at higher mirror ratios, so discharge power remains similar for all mirror cases. Langmuir sweeps and triple probe measurements of T_e (DR2) show slightly (less than 25%) depressed core and slightly elevated edge T_e with increasing mirror ratio (fig. 3.6) but otherwise remains unaffected. The temperature affects I_{sat} measurements through the $\sqrt{T_e}$ term so small changes are not significant relative to changes in density (a 25% change in T_e would yield a 12-13% change in I_{sat}). The low temperatures indicate that the plasma is collisional given the length scales of the system (as seen in table 3.2) and isotropic. Plasma potential decreases across the plasma (fig. 3.7) when the mirror ratio exceeds 1.9. This drop in plasma potential may be caused by the grounding of the anode to the wall, which should begin at $M = 1.93$ given the 72 cm

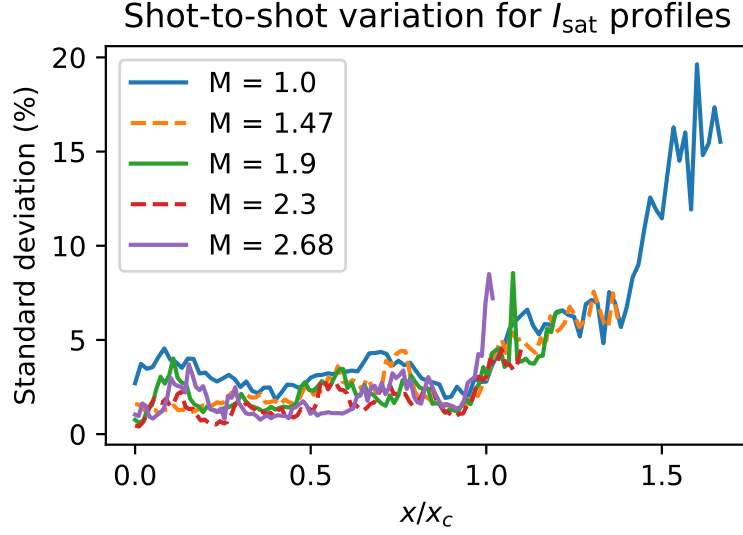


Figure 3.3: Shot-to-shot variation of I_{sat} profiles

anode and 100 cm vessel diameters. The reason for the local minimum in the $M=2.68$ is unknown. This potential profile creates a sheared $\mathbf{E} \times \mathbf{B}$ velocity profile (fig. 3.7) limited to 500 m/s in the core and exceeding ~ 3 km/s at the far edge. The flow does not exceed 4% of the sound speed (tab. 3.2) in the core or gradient ($x = x_{PF}$) region. The mirror ratio does not appear to significantly alter azimuthal flow. The floating potential (V_f) profile also exhibits behavior similar to the plasma potential (fig. 3.7), but is modified by the presence of primary electrons.

3.3.2 Reduced particle flux

The density fluctuation power peaks at the steepest gradient region ($x_{PF} = x/x_c \sim 0.88$) as expected as seen in fig. 3.8. x_{PF} occurs at nearly the same magnetically-mapped coordinate for each mirror ratio. These density fluctuations are a large driver of changes in the cross-field particle flux (eq. 3.1). V_f fluctuations also peak at the same location, but the total power across mirror ratios are similar and, relative to density fluctuations, much lower in the core. Core density fluctuations below 2 kHz are substantial in the core at lower mirror ratios, possibly caused by hollow profiles, nonuniform




Figure 3.4: Midplane I_{sat} profile, shot-averaged and time-averaged from 4.8 to 11.2 ms (assumed of $T_e = 4.5$ eV based on triple probe and Langmuir sweep measurements). Effective area was calibrated using a nearby interferometer. Profile shape remains similar in the core and gradient region when mapped to the cathode radius x_c . The dips in profiles at higher M below $x = x_{\text{PF}}$ are of unknown origin and are not the focus of this study. Shot-to-shot variation is less than 5% for $x \leq 0.95x_c$ and less than 9% for $x \leq 1.4x_c$ for all cases.

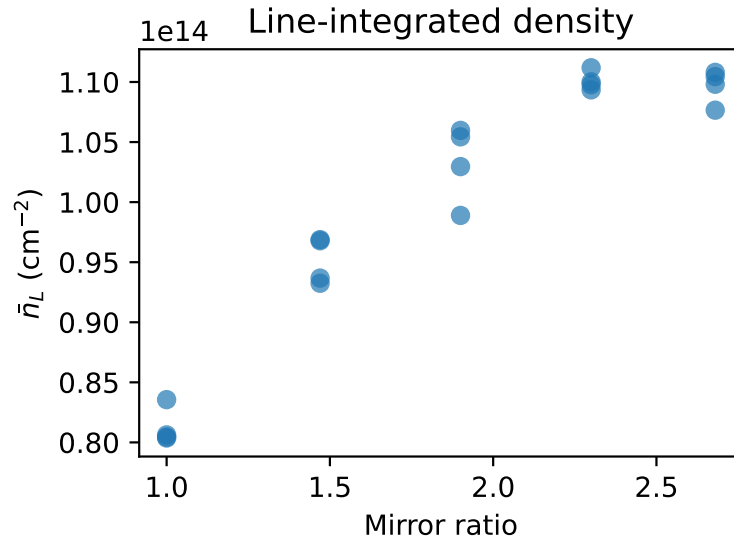


Figure 3.5: Line-integrated density as measured by a 56 GHz heterodyne interferometer as a function of mirror ratio, taken from four discharges for each mirror configuration. Density increases up to a mirror ratio of 2.3 where it appears to level off. The interferometer is located in the mirror cell bad-curvature region at 9.59m, 1.3m closer to the cathode from the midplane.




Figure 3.6: T_e from Langmuir sweeps (DR2) at the midplane. Triple probe results are nearly identical. The increased temperatures directly at the plasma edge, indicated by dotted portions of the curves, are likely artifacts caused by sheath expansion in lower densities. Changes in mirror ratio lead to at most 25% change in T_e . The plasma is collisional and isotropic.

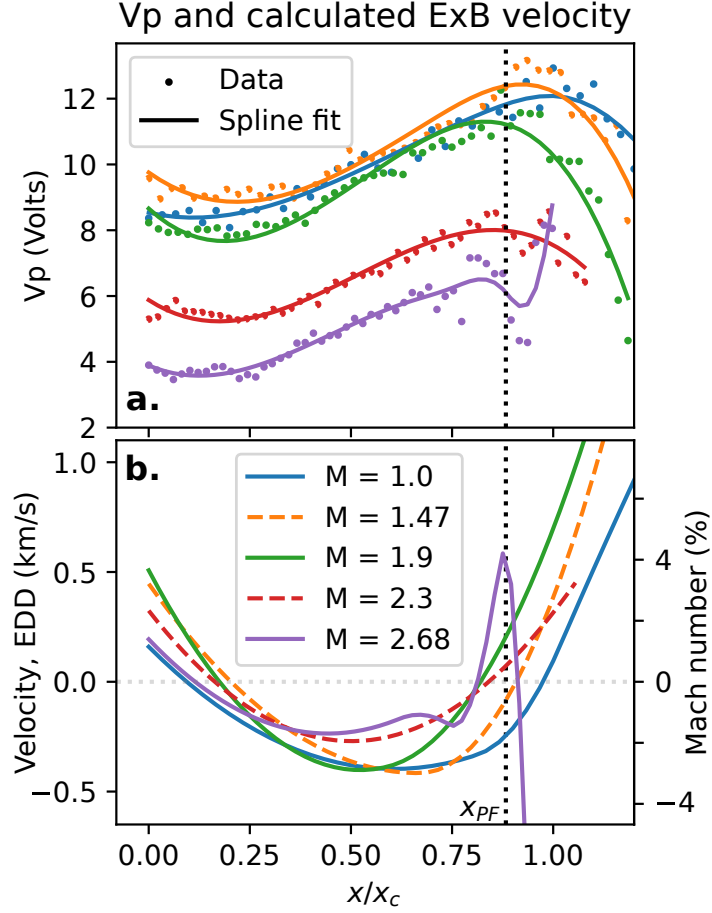


Figure 3.7: Plasma potential (a) and derived $\mathbf{E} \times \mathbf{B}$ velocity profiles (b) from Langmuir sweeps at the midplane. $x/x_c > 1.2$ has been excluded from the graph for greater clarity in the core and gradient region. The electric field was calculated by taking the gradient of the spline-smoothed plasma potential profile. The Mach number (in percent) is calculated using the approximate sound speed evaluated at $x = x_{PF}$ (tab. 3.2). The overall structure of the flows does not appreciably change when mirror ratio is varied.

cathode emissivity, or probe perturbations, but are outside the scope of this study.

A spectral decomposition technique is used to calculate the time-averaged particle flux [125] as seen in fig. 3.9:

$$\Gamma_{\tilde{E} \times B} = \langle \tilde{n} \tilde{v} \rangle = \frac{2}{B} \int_0^\infty k(\omega) \gamma_{n\phi}(\omega) \sin(\alpha_{n\phi}) \sqrt{P_{nn}(\omega) P_{\phi\phi}(\omega)} d\omega \quad (3.1)$$

where k is the azimuthal wavenumber, γ is the coherency, α is the cross-phase, and P the power spectrum. This method is more robust than the naive time-integration of $n(t) \tilde{E}(t)$ because it accounts for the coherency of the density-potential fluctuations. This representation also enables inspection of each contributing term in the event of surprising or problematic results. A plot of the $I_{\text{sat}}\text{-}V_f$ phase can be seen in fig. 3.10. The flattened particle flux in the core is likely caused by primary electrons emitted by the cathode. These electrons have long mean free paths (greater than a few meters) and sample fluctuations along the length of the machine, mixing the phases of these fluctuations. Since floating potential is set by the hotter electron population, the measured V_f fluctuations are no longer related to the local plasma potential fluctuations of a wave by bulk T_e [22]. These primary electrons have a significant effect in the core within the region mapped to the cathode $x \lesssim x_c$. I_{sat} fluctuations are not affected.

Azimuthal wave number is measured by two V_f probe tips 0.5 cm apart. This wavenumber estimation technique yields good agreement with correlation plane measurements (fig. 3.31). Note that \tilde{E} is not directly measured – it is instead calculated through the $k(\omega) \sqrt{P_{\phi\phi}(\omega)}$ terms. The $\tilde{E} \times B$ particle flux clearly decreases with mirror ratio; most of this decrease is attributed to the decrease in density fluctuation power. The particle flux for each mirror ratio was normalized to the $M = 1$ case via the plasma circumference to compensate for the increased plasma surface area at the same magnetically-mapped coordinate x/x_c . This particle flux is consistent with diffusive flux at the Bohm rate $D_B = \frac{1}{16} \frac{T_e}{B} \approx 6.25 \text{m}^2 \text{s}^{-1}$ as observed in other transport studies [102].

T_e profiles and fluctuations may affect particle fluxes but measurements of both were not taken in the same datarun; nevertheless, a quantification of the effect of T_e on particle flux is attempted. T_e fluctuations affect I_{sat} -based density measurements through the $T_e^{-1/2}$ term, and triple probe and




Figure 3.8: I_{sat} (a) and B_{\perp} (b) fluctuation power profiles for signals 2 kHz and up at $z=8.3\text{m}$ (midplane) and $z=7.7\text{m}$, respectively. The lower frequency components in I_{sat} are associated with bulk profile evolution, dominate the core region, and are not the focus of this study.




Figure 3.9: Cross-field, $\tilde{E} \times B$ fluctuation-based particle flux (calculated using eq. 3.1) with respect to mirror ratio. A monotonic decrease in particle flux is observed with increasing mirror ratio at the midplane. Particle flux is normalized by plasma circumference to the $M = 1$ case to account for the geometry-induced decrease in particle flux caused by a larger-diameter plasma.




Figure 3.10: Phase (a) and coherency (b) of I_{sat} current and V_f near x_{PF} at the midplane, smoothed. Positive phase means I_{sat} leads V_f . Peaks in coherency occur between 3-5 kHz and at the drift-Alvén wave peaks between 12 and 25 kHz. These coherency peaks tend to have larger phase shifts than other nearby frequencies.

Langmuir sweep T_e measurements suggest that temperature gradients have a negligible impact. A naive incorporation of temperature fluctuation data from DR2 into particle fluxes from DR1 suggest that cross-field particle flux may be underestimated by up to 50% via the I_{sat} temperature term, but the trend and relative fluxes across mirror ratios remain unchanged. Such a naive incorporation should be treated with skepticism because of the sensitive nature of the flux with respect to the gradient and the differences in profiles between DR1 and DR2. These difference in profiles may be caused by cathode condition, deposits on the anode, or a different gas mix and are difficult to account for.

3.3.3 Compensating for the T_e profile

Electron temperature (T_e) compensation for the I_{sat} measurement can be done in several ways. One way is to account for the average temperature (i.e., steady state) when calculating the density from I_{sat} . T_e can be gathered from triple probe or swept measurements. Triple probe measurements are generally less reliable than swept probe measurements. The difference between swept and triple probe T_e measurements can be seen in fig. 3.11. The two techniques have roughly good agreement, though the triple probe appears to slightly underestimate the temperature. The spikes in the edge are likely from sheath expansion of the probe in the swept measurements exacerbated by low density [153] (see fig. 3.6).

T_e fluctuations can affect I_{sat} fluctuation measurements through the $\sqrt{T_e}$ term. In this case, T_e measurements are difficult to compensate for in DR1 because of the changes in profiles between DR1 and DR2, so the T_e fluctuations were included by finding the ratio in DR2 of I_{sat} fluctuations before and after including these T_e fluctuations. This ratio was then applied to DR1. The issue of mismatched profiles still persists but this method allows for changes in fluctuation power between the two datarun sets. In general, \tilde{T}_e/T_e fluctuations are at most than 30% (near the edge), and much lower in the core seen in fig. 3.12.

This T_e compensation becomes particularly important when calculating the I_{sat} profile gradients




Figure 3.11: Difference between swept and triple probe temperature measurements. The triple probe appears to slightly underestimate the temperature and temperature gradient.




Figure 3.12: T_e and T_e fluctuation profiles from the triple probe. RMS electron temperature fluctuations are not particularly large.

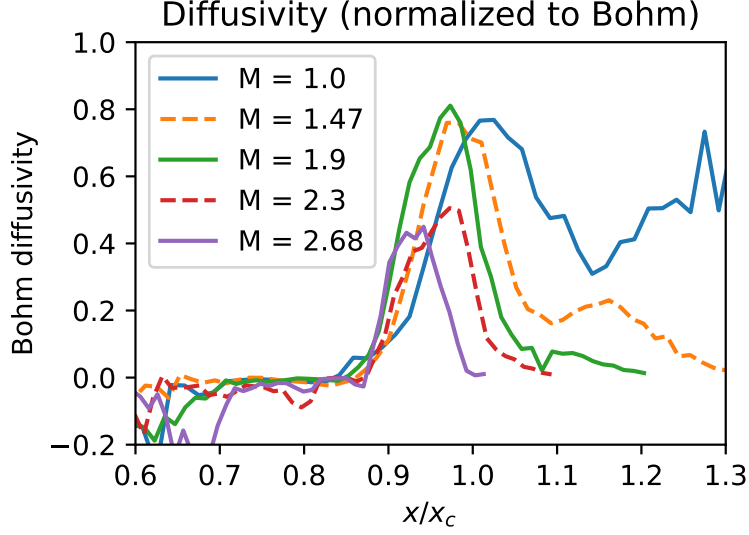


Figure 3.13: Diffusivity relative to D_B using a tanh fit for the density profile and the particle flux measurement assuming a constant T_e of 4.5 eV across the profile.

which is needed when calculating the diffusivity. A calculation of the diffusivity scaled to the Bohm diffusivity $D_B = \frac{1}{16} \frac{T}{eB}$ can be seen in fig. 3.13. This calculation uses the particle flux calculated earlier (in the paper) and tanh fit on the density profile for a density smooth gradient. In general, mirror ratios higher than two have a lower diffusivity. When the particle flux is compensated for T_e fluctuations, the temperature profile used in for the Bohm diffusion coefficient, and the density profile is smoothed convoluting a $\sigma = 2$ cm gaussian, the diffusion coefficient relative to D_B are roughly 2.5 times greater, seen in fig. 3.14. The trend, however, remains relatively the same: higher mirror ratios tend to have a lower diffusivity. The impact of different profile smoothing methods on the density gradient can be seen in fig. 3.15.

3.3.4 Drift waves

The I_{sat} fluctuation power spectra in the region of peak power $x \sim x_{\text{PF}}$, also where the density gradient is strongest, can be seen in fig. 3.16. Notably, the fluctuation peaks shift to higher




Figure 3.14: Diffusivity with T_e compensation relative to D_B . The particle flux is compensated for T_e fluctuations, and the swept-probe temperature profile is used for T_e . The diffusivity is around 2.5 times higher than without compensation, but the trend remains similar.




Figure 3.15: I_{sat} gradients under varying profile smoothing methods




Figure 3.16: I_{sat} (density) fluctuation power averaged over a 1 cm region around x_{PF} at the midplane. The fluctuation power is largely featureless below 2 kHz and beyond 40 kHz aside from electronics noise.

frequencies and decrease in total fluctuation power. The shift in frequency may be a result of the the Doppler shift caused by the change $\mathbf{E} \times \mathbf{B}$ plasma rotation seen in fig. 3.7 at the location $x/x_c \approx x_{\text{PF}}$. The shift in frequency is somewhat smaller than what would be expected from the field line-averaged increase in Alfvén speed at the longest possible wavelength. The phase angle of I_{sat} and V_f fluctuations provides insight into the nature of the driving instability. Including a nonzero resistivity η in the drift wave leads to a small phase shift δ between density and potential. This phase shift δ in a collisional plasmas is on the order of $\delta \approx \omega v_e / k_{\parallel}^2 \bar{v}_e^2$ [74]. Estimating this quantity using measured and typical values ($k_{\parallel} = 0.18$ rad/m, $\bar{v}_e = 1300$ km/s, $v_e = 3.7$ MHz, $\omega = 12$ kHz) yields a substantial phase shift of $\delta \approx 46^\circ$, which roughly agrees with the phase shifts in fig. 3.10, though the implied increased phase shift at higher frequencies does not agree with measurements. As seen in fig. 3.10, the phase shift between I_{sat} and V_f fluctuations are larger below 10 kHz, implying the presence of additional modes beyond or significant modification of resistive drift wave fluctuations. The phase difference between two V_f probes, 3.83 m apart, was used to calculate




Figure 3.17: k_{\parallel} (a) and coherency γ (b) as a function of frequency. Only results from the $M = 1$ case are available, but it is clear that there are long ($\gtrsim 34$ m) wavelength modes at 3 and 12 kHz. The probes used for calculating k_{\parallel} were located at the midplane ($z=8.31$) and $z=12.14$ m, 3.83 m apart.

the parallel wavelength $\frac{2\pi}{\lambda} = k_{\parallel} = \phi_{Vf1, Vf2}/\Delta z$ assuming the wavelengths are greater than 7.66 m. The two probes mapped to the same field line only in the $M = 1$ configuration, so parallel wavenumbers are available only for the flat case. Parallel wavenumbers are theoretically calculable from 2d correlation planes but the coherency dropped dramatically when a mirror geometry was introduced. A 34 m wavelength mode likely contributes to the measured k_{\parallel} from 3 to $\gtrsim 10$ kHz (fig 3.17). Drift waves are long-wavelength modes so coherent density and potential fluctuations along the flux tube are expected. The coherency is a measure of similarity of the spectral content of two signals, in this case V_f probes 1 and 2. The coherency is defined as $\gamma = \frac{|\langle P_{1,2} \rangle|}{\langle |P_{1,1}|^2 \rangle \langle |P_{2,2}|^2 \rangle}$ where $P_{x,y}$ is the cross-spectrum between signals x and y and the angle brackets $\langle \rangle$ denote the mean over shots. The coherency between the two V_f probes drops off with increasing frequency, with a slight bump at around 12 kHz. There are several candidates for the driving mechanism of the 3-5 kHz mode, but the 12 kHz mode is most likely a drift-Alfvén wave.

Experimental measurements (and various theories) [100] suggests that the normalized density fluctuation level \tilde{n}/n should fall near $3-10 \rho_s/L_n$. A plot of this relation using experimental data can be seen in fig. 3.18. However, comparison of \tilde{n}/n with the mixing length $1/(k_y L_n)$ [100] show that the normalized density fluctuations are about an order of magnitude too small for the $1/(k_y L_n)$ observed which is unexpected if these fluctuations correspond to drift waves. The normalized density fluctuation amplitude is expected to scale with the gradient scale length L_n , so the fluctuation power then scales with L_n^2 . A plot of scaling I_{sat} fluctuation power – assuming constant k_y – can be seen in fig. 3.19. The 3 kHz mode appears to follow this L_n scaling.

Another issue with this drift-wave interpretation of results is that the electron thermal diffusion along the field line is too high. In the fluid picture, the plasma must be collisional enough that thermal equilibrium is guaranteed (i.e., the temperature is Maxwellian), but if the collision rate is too high then thermal gradients can develop along the field line [63]. This condition on thermal diffusivity condition for the drift wave ω and k_z is $\omega \ll k_z^2 v_{e,th}^2 / \nu_{ei}$. Plugging in values from the experiments yields frequencies at least 5 times greater than mandated by the diffusivity condition and the condition is violated. This condition violation may be responsible for the unexpected phase

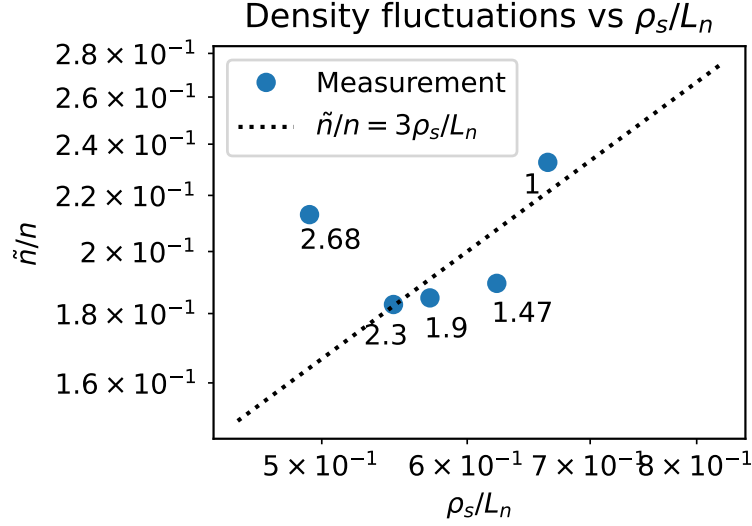


Figure 3.18: Normalized density fluctuations vs ρ_s/L_n . The measured values fall close to the $\tilde{n}/n = 3\rho_s/L_n$ line which is consistent with theory.

shifts seen between the density and potential fluctuations if drift waves were the only instability present.

3.3.5 Turbulence modification

The wavenumber-power relation in fig. 3.20 shows decreased fluctuation power when a mirror configuration is introduced. However, there is no discernible trend when the mirror ratio is increased further. The exponential nature of the curve also remains unchanged. The greatest decrease in fluctuation power occurred in low and high k_y 's, around 10 and 70 rad/m. The shape of the power- k_y curves follow an exponential distribution, and is inconsistent with a 2d drift-wave turbulent cascade (Wakatani Hasegawa k^{-3}) [161]. The steep dropoff in fluctuation power with k_y suggests that higher-wavenumber fluctuations do not have a significant effect on transport.

Previous simulations in a flat field [56] predicted frequency and wavenumber spectra that can be fit with many power laws or exponentials, but the data presented here (figs. 3.16, 3.21, 3.20) appear




Figure 3.19: I_{sat} fluctuation power when scaled by the square of the gradient scale length and the squared density. We expect this value to be constant assuming the same k_{\perp} .

to follow an exponential relationship within measurement variation.

Core fluctuations appear to decrease dramatically as seen in the I_{sat} fluctuation power (fig. 3.8). The I_{sat} decorrelation time increases from ~ 0.7 ms for $M = 1$ to ~ 2.5 ms for $M = 2.68$. At $x = x_{PF}$, decorrelation times for all mirror ratios remained at 0.2 ms.

3.3.6 Magnetic fluctuations

The perpendicular magnetic fluctuation (B_{\perp}) component of the drift-Alfvén wave can be seen in fig. 3.21. These B_{\perp} fluctuations are spatially and spectrally coincident with the electrostatic fluctuations (fig. 3.16). Drift-Alfvén waves have been studied in the LAPD in the past [103, 159]; strong coupling is observed for $\beta_e > m_e/m_i$ which is satisfied in this study. The Alfvén speed $\omega/k_{\parallel} = v_A = B/\sqrt{4\pi nM}$ (given $\omega \ll \Omega_{ci}$) when averaged over the entire column ranges from ~ 450 to ~ 1600 km/s. A k_{\parallel} corresponding to a wavelength $\lambda = 34$ m roughly falls within the bound established by the kinetic and inertial Alfvén wave dispersion relations at the frequency




Figure 3.20: Fluctuation power summed for each k_y for frequencies up to 100 kHz, smoothed. The contribution to fluctuation power is negligible past 100 kHz. The fluctuation power decreases substantially when a mirror configuration is introduced, but no trend is seen otherwise and the k_y spectra remain exponential. Note the logarithmic scale.




Figure 3.21: B_{\perp} fluctuation power averaged at the core from 0 to 3 cm (a) and around the peak fluctuation point ($x \sim x_{PF}$) (b). Fluctuation power decreases across the board with mirror ratio except for core frequencies close to Ω_{ci} . Peaks around 10 – 30 kHz at x_{PF} are consistent (region 2) with drift-Alfvén waves and the near-cyclotron frequency features in the core may be resonating Alfvén waves created by the magnetic mirror. Frequencies below 2 kHz and dominated by instrumentation noise and thus excluded.

peaks observed at $x \sim x_{PF}$ seen in fig. 3.21. The lengthening of field lines caused by curvature accounts for at most 10% of the change in frequency.

The spatial extent of the B_{\perp} features identified in fig. 3.21 are plotted in fig. 3.22. Feature 1 at ≈ 6 kHz shows increased fluctuation amplitudes at $x = 0$ for mirror ratios 1.9 and above, but for $M = 1$ and $M = 1.47$ there is no increase in fluctuation power. A similar feature, but at a much smaller level, is observed in I_{sat} fluctuation power in the core as well. This core feature may be caused by the density depletion in the core seen in the I_{sat} profile (fig. 3.4) driving low-amplitude waves or instabilities. Feature 2 in fig. 3.22 is the magnetic component of the drift-Alfvén wave. The fluctuation power peaks at the gradient region and corresponds with the peak in density fluctuations (fig. 3.8).

Feature 3 is particularly interesting because this is the only fluctuating quantity to *increase* with mirror ratio, seen in fig. 3.23. This feature may be broad evanescent Alfvénic fluctuations from the plasma source. These fluctuations have been observed in the LAPD in the source region alongside an Alfvén wave maser at $0.57 f_{ci}$ [104].

The sub-2 kHz modes in B_{\perp} and its harmonics are nearly constant in power across the entire plasma; these features are likely perturbations from the magnet power supplies and thus ignored. The lack of radial, azimuthal, and axial structure in these magnetic signals below 2 kHz and narrow bandwidth indicate a non-plasma origin. Significant radial and azimuthal structure in B_{\perp} fluctuation power starts to appear in frequencies larger than 4 kHz.

The drift-Alfvénic nature of the 12 kHz Bdot feature is confirmed by changing the flat field from 500G to 400G: the feature shifts down in frequency from 12 to 10 kHz seen in fig. 3.24. From the drift wave and Alfvén wave dispersion relations the frequency is expected to be $400\text{G} / 500\text{G} = 0.8$ of the original, which is approximately what is observed. The k_y of the drift-Alfvén wave also has an effect and may be responsible for a $10\text{ kHz} / 12\text{ kHz} = 0.83$ factor instead.

There may be some sort of resonator made by the mirror cell and its interaction with Alfvén waves. In fig. 3.25, the behavior of the B_{\perp} spectrum in the core changes dramatically between 1




Figure 3.22: B_{\perp} fluctuation power profiles for the three regions shown in fig. 3.21: region 1 (6 kHz) (a), region 2 — where frequencies are taken from the peaks of the drift-Alfvén waves for each mirror ratio (b), and region 3 (114 kHz) (c).




Figure 3.23: Summed fluctuation power of B_{\perp} in the core ($x/x_c \leq 0.3$) as a function of mirror length and ratio. Top (a): the fluctuation power is normalized by the sum of the full-spectrum summed power. Bottom(b): the frequency of the power distribution > 50 kHz weighted by the fluctuation power.




Figure 3.24: B_{\perp} flat field for 500G and 400G flat fields. The frequency of the identified drift-Alfvén wave at 12 kHz drops when the field is lowered, as expected.

and 10 kHz in the short mirror when compared with the medium and longer mirrors. It's unlikely that this is an Alfvénic fluctuation because the wavelength is an order of magnitude too large to fit in the machine.

For completeness, B_z fluctuation measurements are seen in fig. 3.26. The peaks in the 10 kHz region are potentially caused by slight coil misalignment of the probe (picking up B_{\perp}), or a compressive component of the drift-Alfvén wave.

The low frequency fluctuations in the Bdot spectra may seem important but plotting the spectra as a function of position (fig. 3.27) clearly shows the harmonics of the signal and the narrow bandwidth of them. This spectral feature is present regardless of mirror ratio, but changes in magnitude in approximate proportion with the field, i.e., the magnet power supply current. This power supply-induced field fluctuation can easily be seen in the ≈ 625 Hz mode in B_z , seen in fig. 3.28. The fluctuation power is largely constant across the entire plasma column, with the




Figure 3.25: B_{\perp} at $x=0$ for different mirror lengths. The origin of fluctuations between 1 and 10 kHz is unknown.

fluctuation power increasing with increased mirror fields. The taper of the fluctuation power at the edge could be caused by the background field vector no longer pointing in the z direction as the probe approaches the magnet coil. In general, the probe valves are not centered between the magnet coils, leading to rotation of the of the field vector as the probe is pulled out.

3.4 2d Structure

The perpendicular magnetic field structure is measured by collecting x-y planar \dot{B} ($dB_{\{x,y,z\}}/dt$) data alongside a stationary, axially separated I_{sat} reference probe (DR2). This probe provides a phase reference for the magnetic field fluctuations, allowing a 2d map of relative phase to be constructed over many shots. Only the region around x_{PF} was measured because of constraints on probe movement. The amplitude and phases for each magnetic field component are then used to reconstruct the local magnetic fluctuation vector \mathbf{B} . The axial current density structure, j_z , can be derived from this vector field. \mathbf{B} and the corresponding j_z for the flat-field ($M = 1$) case can be seen in fig. 3.29. Two main current channels can be seen with the magnetic fields circulating around




Figure 3.26: B_z fluctuations in the core and at x_{PF} . The peaks near the edge could be from coil misalignment or potentially a compressible component of the drift-Alfvén wave.




Figure 3.27: B_{\perp} fluctuation power for mirror ratios of 1 and 2.68. Lower frequencies are shown and the colorbar clipped to show detail in what appears to be power supply fluctuations.




Figure 3.28: B_z fluctuation power profiles for all mirror ratios at 625 Hz.




Figure 3.29: Perpendicular magnetic field and the derived current density for the flat-field ($M = 1$) case using a Bdot probe with an axially-separated I_{sat} reference (DR2). The x-y plane was centered near x_{PF} .

them. This structure quickly decoheres in time as expected in a turbulent plasma. At higher mirror ratios, the field magnitude and corresponding current density decrease (which was also seen in DR1: fig. 3.21). Similar structure is seen in the mirror configurations; the $M = 1.9$ and $M = 2.68$ cases can be seen in fig. 3.30.

Using two, axially-separated, correlated I_{sat} measurements (DR2), with one collecting x-y planar data, the azimuthal mode number m (radially integrated) was calculated. Higher-frequency and higher- m features are seen with increasing mirror ratio (fig. 3.31). The increased frequencies may be caused by a change in Doppler shift by the $\mathbf{E} \times \mathbf{B}$ flow. This higher- m trend suggests that azimuthal structures do not scale with increased plasma radius but instead remain roughly the same size. The




Figure 3.30: Perpendicular magnetic field and the derived current density for the $M = 1.9$ and $M = 2.68$ cases computed in the same manner as fig. 3.29. The x-y planes were centered near x_{PF} , and the view size was kept constant across the plots. The structure is much less obvious in the mirror cases, but all exhibit the expected Alfvén wave pattern

limited planar probe movement caused an increase in the lower bound on m in higher mirror ratios. At mirror ratios 1.47 and higher, the lower frequency component (< 10 kHz) appears to decrease significantly in amplitude. Calculating k_{\perp} from m evaluated at $x \sim x_c$ yields similar k_y values as the two-tip technique (fig. 3.32). The average k_y for a given frequency can be calculated using two V_f tips on the same probe by calculating the phase difference and dividing by the spatial separation of 5 mm: $k_y = \phi_{vf1, vf2} / \Delta y$ [14]. The maximum $|k_y|$ measurable before aliasing is $\pi / \Delta y \approx 628$ rad/m. As seen in fig. 3.32, the k_y spectrum remains similar across mirror ratios, but the wavenumber extends further into higher frequencies with increasing mirror ratio. These azimuthal mode numbers and gradient scale lengths are consistent with linear simulations using the 3d fluid code BOUT [122] in the flat, unbiased case.

3.5 Discussion

3.5.1 Lack of mirror-driven instabilities

No evidence is seen for mirror-driven instabilities — curvature, loss-cone, or otherwise. Given the LAPD parameters in this study (tables 3.1 and 3.2), the collision frequencies are sufficiently high such that the mirror is in the gas-dynamic regime: losses out of the mirror throat are governed by gas-dynamic equations rather than free streaming through the loss cone. To be in the gas-dynamic regime, the mirror length must exceed the mean free path of the ions [77]:

$$L > \lambda_{ii} \ln M / M \quad (3.2)$$

where L is the mirror length, λ_{ii} is the ion mean free path, and M is the mirror ratio. These collisions populate the loss cone and maintain a (cold) Maxwellian distribution, eliminating the possibility of loss-cone-, ion-driven instabilities like the AIC [23] or DCLC [147, 84] instabilities that have been observed in other (historic) devices.

The paraxial, approximate interchange growth rate is [124, 134]

$$\Gamma_0 = \frac{c_s}{\sqrt{L_M L_P}} \quad (3.3)$$




Figure 3.31: Azimuthal mode number m amplitudes calculated from two axially-separated, correlated, I_{sat} probes. Increasing mirror ratio (a to e) leads to increased m at higher frequencies. (DR2)




Figure 3.32: k_y averaged about x_{PF} and smoothed for each mirror ratio calculated using two vertically-separated V_f tips on the same probe. Little change is seen in k_y at lower frequencies but higher frequencies tend towards larger k_y at higher mirror ratios.

which yields $\Gamma_0 \approx 1.2$ kHz using $L_M \approx 7$ m and $L_P = 17$ m. c_s is used instead of \bar{v}_i because $T_i \ll T_e$ and mirror length L is split to distinguish between the contributions of the plasma length and mirror length to inertia and to curvature drive, respectively. Interchange is not visible in-part because the aspect ratio of these mirrors is quite large, limiting the growth rate of interchange. The length of the mirror (3.5 m), radius of curvature (6-7 m), and plasma column (17 m) are much larger than the radius of the plasma (0.5 m maximum), so the plasma inertia is large relative to the instability drivers. Line-tying to the cathode may further lower the growth rate. The hot cathode used for plasma formation could function as a thermionic endplate that can supply current to short out the flute-like interchange perturbations. Line-tying has been seen in flux rope experiments on the LAPD using a hotter, denser source [156], also in other devices [52], and is why interchange was not seen in the earliest mirror machines [124]. Note that the plasma terminates on the cathode or end plates before the magnetic field flares out, so there is no contribution to stability from an expander tank as seen in other GDTs [134, 77]. Finite Larmor radius (FLR) effects may provide a stabilizing effect for larger azimuthal mode numbers. At the highest mirror ratio, assuming a plasma radius of $a_0 = \sqrt{2.68} * x_{PF} = 43$ cm, the FLR stability criterion $\frac{m}{2} \frac{\rho_i L}{a_0^2} > 1$ [134] suggests a stabilizing effect may be present for azimuthal mode numbers $m > 4$.

If the curvature-induced interchange instability were observable, then introducing a mirror configuration would lead to new features in I_{sat} and, potentially, B_{dot} fluctuations. In particular, low-frequency mode(s) – likely less than 10 kHz given the low m -number and plasma rotation rates – would be observed growing from the pressure gradient region. For onset of the interchange instability, the mirror curvature or plasma pressure would need to be increased but the precise conditions required for this onset are not yet known for the LAPD.

Interchange could also be at least partially stabilized by the continuous production of electrons in the core that are electrostatically trapped by the ambipolar potential [66]. The intuition behind this stabilization mechanism is as follows: electrons are continuously produced via ionization of neutrals, and any change in the local potential will cause more or fewer electrons to be lost out the ends of the device along that field line, counteracting the potential change. This stabilization

mechanism has been experimentally demonstrated to completely suppress interchange when the ambipolar potential $\Phi \gtrsim 6T_e$ [90].

The $\mathbf{E} \times \mathbf{B}$ shear flow present (fig. 3.7) may also make a contribution to the stabilization of interchange [134, 10, 9, 15]. The estimated shearing rate is between 3 and 10 kHz, which is greater than the estimated ≈ 1.2 kHz growth rate of the interchange mode.

3.5.2 Instabilities driving LAPD turbulence

The 12 kHz and higher modes are most likely drift-Alfvén waves driven by the density gradient. However, there are several possibilities for the low-frequency (3-6 kHz) modes observed in the I_{sat} and magnetic fluctuation spectra. Drift waves, rotational interchange, the conducting wall mode, and a nonlinear instability. The long-wavelength nature of the mode – neither $k = 0$ or $k = \pi/L$ – points to a mixture of modes contributing to the fluctuations in that frequency range.

This rotational interchange mode has the following attributes, as summarized by [81]: flute-like ($k_{\parallel} = 0$), $|e\tilde{\phi}/T_e|/|\tilde{n}/n| \gtrsim 1$, radial potential phase variation 45 to 90°, maximum possible $|e\tilde{\phi}/T_e| < 1$. All of these attributes are seen for the lower frequency (3 kHz) mode. The V_f radial phase variation when $M > 1$ is not clearly seen because the coherency is dramatically reduced along the field line. The rotational interchange mode could couple with the drift wave at $k_{\parallel} = \pi/L \sim 0.37$ rad/m ($n = 0.5$), which has been observed in the past [139] and likely present here. Estimates of shearing rate from the $\mathbf{E} \times \mathbf{B}$ flow velocity profile (fig. 3.7), calculated fluctuation ratios, and radial phase shift variation suggest that Kelvin-Helmholtz-driven turbulence is not significant, if present at all. Historically, biasing a limiter has been required to clearly observe the Kelvin-Helmholtz instability [75, 137, 139]. Rotational interchange can be significant driver of the broadband turbulence spectrum when this limiter is biased, seen in both linear simulations [122] and biasing experiments [139].

Low frequency density fluctuations may also be driven by a flute-like conducting-wall temperature-gradient instability which only requires an electron temperature gradient to grow (even with straight

field lines) [16]. Simulations of turbulence in the LAPD suggest the possible presence of these conducting wall modes (CWM) which have the highest growth rate for $m \leq 20$ [55]. This lower- m mode could be responsible for the peak around 3 kHz in the $M = 1$ I_{sat} fluctuation (fig. 3.16) and azimuthal mode numbers (fig. 3.31) and for the low-frequency low- k_{\parallel} or flute-like behavior (fig. 3.17). This CWM may also be responsible for flatter electron temperature profiles seen in previous studies [119, 139] (fig. 3.6).

These linearly unstable modes may be outgrown by a rapidly-growing nonlinear instability that couples to drift-like modes as suggested by simulations [55]. This nonlinear instability is driven by the density gradient at an axial mode number of $n = 0$ and nonlinearly transfers energy to $n \neq 0$ fluctuations.

Neither the conducting wall mode nor the rapidly-growing nonlinear instability have been studied in detail on the LAPD. Precise identification the modes present requires further study.

3.5.3 Causes of particle flux reduction

The reduction in particle flux explained by a reduction in density fluctuations can be in-part explained by an increased gradient scale length $L_n = \frac{n}{\nabla n}$ (fig. 3.33), decreasing the linear drift wave growth rate and saturation level seen in sec. 3.3.5. This gradient length reduction may also reduce the growth rate of the rotational interchange instability, which could contribute to the low-frequency large-amplitude density fluctuations. The influence of this density fluctuation reduction appears reduced at higher mirror ratios past $M = 1.9$, where the wavenumber and phase angle appear to decrease in magnitude. The plot showing this breakdown in particle flux can be seen in fig. 3.34. The changes in I_{sat} fluctuation power is the most obvious driver, but the $I_{\text{sat}} - V_f$ phase difference, coherency and wavenumber also seem to have an effect. The V_f fluctuation power remains largely consistent across the different mirror ratios. Note that this particle flux appears somewhat different because this is using the uncalibrated I_{sat} values and the flux is not scaled by solid angle. This flux also does not use temperature-compensated I_{sat} measurements.




Figure 3.33: Gradient scale length L_n and the associated term in the drift wave growth rate L_n^{-2} . This scale length was calculated over a 3 cm region around x_{PF} (peak fluctuation region) at the midplane. Increasing the mirror ratio increases the gradient scale length, which suggests weakening of the underlying instability driver.

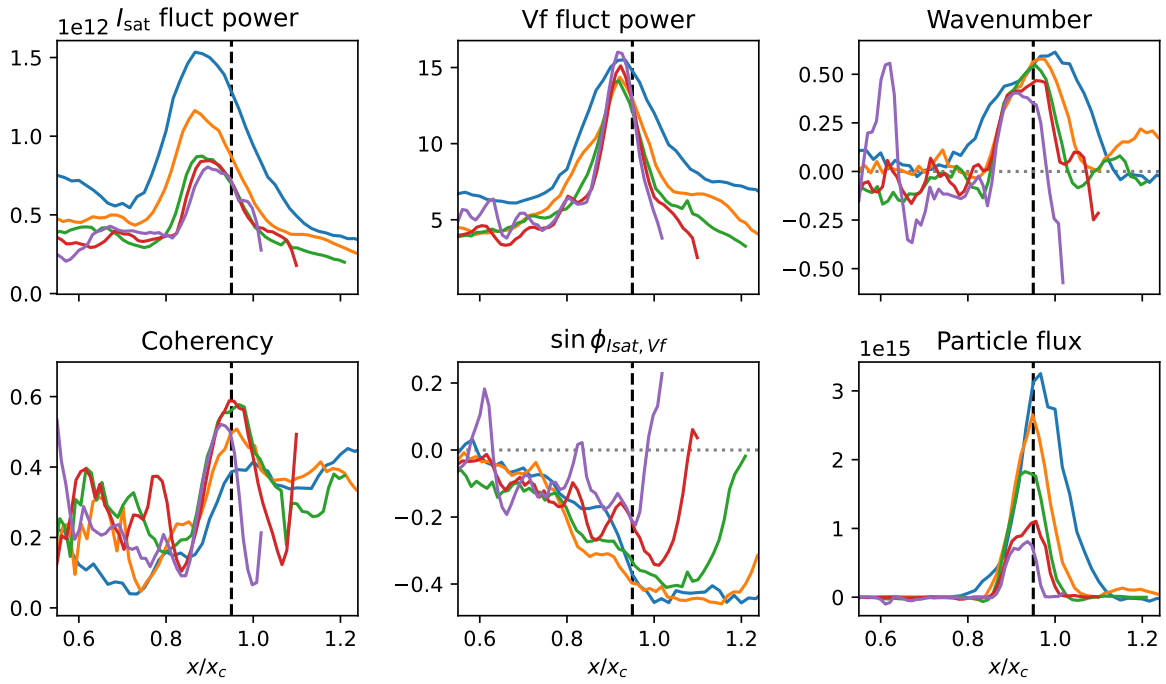


Figure 3.34: The particle flux broken down into the components used to calculate it. The dashed black line is simply a visual reference near the peak particle flux at $x/x_c = 0.95$. The I_{sat} fluctuation power appears to be the largest driver in changes in particle flux. The colored lines correspond to mirror ratio as seen in earlier plots.




Figure 3.35: Decorrelation time from I_{sat} time series data for different mirror ratios. All of the mirror ratios have a minimum decorrelation time at x_{PF} and much longer times (slower rate) in the core.

The decorrelation time of I_{sat} time series data is around 0.15 ms at x_{PF} . An estimate of the $\mathbf{E} \times \mathbf{B}$ flow shear from fig. 3.7 (DR2) yields a shearing time between 0.1 and 0.3 ms at x_{PF} . These times suggest that spontaneous flow shear may be important for suppressing turbulence, as seen in other studies [138, 27], at all mirror ratios. However, no clear trend in shearing strength is seen with mirror ratio.

The decorrelation time of a signal is calculated by taking the autocorrelation of a signal – I_{sat} in this case – and finding the full-width half-max of the envelope using a Hilbert transform. This decorrelation time can be seen in fig. 3.35. The decorrelation is minimized at x_{PF} and maximized in the core, further confirming the turbulent nature of the fluctuations at x_{PF} .




Figure 3.36: ExB shearing rate based on smoothed plasma potential profiles. This shearing rate is comparable to the decorrelation rate at x_{PF} , seen in fig. 3.35.

The estimated shearing rate from DR2 can be seen in fig. 3.36. The rate is plotted instead of time because of the singularity when the flow reverses. At around x_{PF} ($x/x_c \approx 0.87$), the shearing rate is around 2 to 8 kHz meaning the shearing time is around 0.5 to 0.125 ms. This is fairly close to the decorrelation time from the I_{sat} time series measurements (fig. 3.35). These similar times/rates suggests that ExB shearing may set the limit on cross-field transport.

It is important to note that the electron thermal diffusion time along the field line is very long compared to the frequency of the drift wave ($\omega \gtrsim k_{\parallel} \bar{v}_e^2 / \nu_{ei}$) [63] so the electron temperature along the field line may not be constant on the drift wave timescale. This factor is not taken into account in this analysis but may have substantial impact on interpretations of the measured phase shift.

3.5.4 Differences between DR1 and DR2

Directly applying signals between these two dataruns is not quite appropriate because the profiles/plasmas changed appreciably. These changes could have been caused by differences in cathode




Figure 3.37: Discharge power vs mirror ratio. DR2 had roughly 5% lower discharge power than in DR1 for unknown reasons.

temperature, emissivity, or other properties. The discharge power for DR2 was roughly 10% smaller than what was seen in DR1 seen in fig. 3.37. Since the discharge voltages were similar (DR1: 62.5 vs DR2: 60.5) we expect to see less dense plasmas in DR2.

Changes in the I_{sat} profiles between the two dataruns (and between two separate measurements in DR1) can be seen in fig. 3.38. Interestingly, there is some difference in the profiles *within the same datarun* which could be caused by probe shadowing. Probe shadowing effects should be less important in mirrors because the probe closest to the cathode magnetically maps to a region further outside than the probes in the mirror cell. This difference in density can also be seen in the line-integrated density from the 56 GHz interferometer (port 23): fig. 3.39. These differences in density could also be caused by different hydrogen and helium pressures in the runs. Helium pressure was roughly the same for both dataruns (6e-6 to 3e-6 for DR1, 6e-6 to 2e-5 for DR2), but the hydrogen pressure was an order of magnitude higher for the DR2, on the order 7e-6 instead of 1e-7 for DR1. These differences in pressures could have had an effect on plasma formation and transport, thus affecting profiles. Hydrogen fraction is known to have an effect on breakdown characteristics




Figure 3.38: I_{sat} profiles ($M=1$), DR1 vs DR2 in the mirror cell. DR2 has reduced I_{sat} across the entire profile. In identical machine settings in DR1, data taken a few days apart are also different.

in the newer Lanthanum-hexaboride (LaB6) cathode.

Differences could also occur within dataruns. Calibrating the effective area of the I_{sat} probes can be done using the 56 GHz interferometer, but this calibration factor drifted over time and seen in fig. 3.40. This could be deposits being removed or added to the probe, affecting the effective area. This calls into question the reliability of absolute I_{sat} measurements, but we proceed regardless because there's no easy way to fix this issue.

3.6 Conclusions and future work

Turbulence and transport was studied in mirrors with varying lengths and ratios using the flexible magnetic geometry of the LAPD. Particle flux and fluctuation amplitudes decreased up to a factor of two when mirror ratio was increased. The primary drivers of turbulence could be the rotational interchange mode, caused by spontaneous rotation, a nonlinear instability, drift instabilities, and drift-Alfvén instabilities driven by the density gradient. A temperature gradient-driven conducting




Figure 3.39: Line integrated density from the 56 GHz, DR1 vs DR2. On average, DR2 has a lower density than what's seen in DR1. The interferometer is located near the region of good curvature closest to the cathode.

wall mode may also be present. The decrease in density fluctuation amplitudes may be attributed to an increase in the gradient scale length caused by the dimensionally wider plasma at the mirror midplane. Despite imposing a mirror configuration, no signs of mirror-driven instabilities were observed. The highly-collisional, GDT-like plasma produced suppressed any velocity space instabilities. The interchange growth rate was likely suppressed to an undetectable level by line-tying, in-cell electron production, and shear flow.

Future experiments in hotter regimes with the new LaB6 cathode [127] will need to be performed to evaluate the robustness of these results, particularly concerning the stabilization of curvature-induced interchange. Additionally, the source field should be matched to the mirror midplane field so that the plasma remains the same radius to isolate geometric effects. Simultaneous measurements using flux and/or vorticity probes and I_{sat} are needed to concretely determine if azimuthal flow shear is modified by the mirror field, and to quantify the effect of flows on rotational interchange and drift wave instability drive in general. Multiple simultaneous axial measurements of potential would




Figure 3.40: I_{sat} calibration factor over different dataruns from the same run sets. Datarun number is monotonically increasing, so in this case it's used as a proxy for time. A rather large variance is seen in this calibration factor, which suggests the I_{sat} probes having time-varying characteristics affecting the measurement.

enable better understanding of the axial wavenumber and identification of possible modes.

CHAPTER 4

Creating a randomized dataset for machine learning tasks

4.1 Goal and introduction

The goal of collecting this dataset was to maximize the diversity of data coming from the LAPD. Previous datasets – even one made of 29 million passively-collected shots over three years – did not contain sufficient diversity to conduct an interesting ML study. In particular, the data must be sufficiently diverse to allow an optimization study without the need to collect more data. In addition, many diagnostics were recorded so that the signals could be correlated on the same shot, either in the machine learning model itself or as a preprocessing step. This chapter describes the process of collecting these data, example signals, and biases within the dataset. All of the data from this campaign (several terabytes) is available upon request.

The LAPD has many experimental control parameters for various physics studies. While the device can accommodate various insertable components, this dataset focuses on the parameters fundamental to the operation of the main cathode. Specifically, half way between the cathode and anode are three gas puff valves: East, West, and top. The aperture, duration, and triggering of these valves has a large impact on plasma formation. A static gas fill system also exists but it is not used. The cathode-anode voltage (and consequently, discharge power) strongly influences plasma density and temperature downstream of the source. Additionally, the magnetic field configuration substantially shapes the plasma column. One crucial variable not considered in this dataset is the cathode temperature, as its adjustment and equilibration requires many hours, limiting dataset diversity. This combination of diagnostic coverage, high repetition rate, and extensive configurability

renders the LAPD particularly suitable for machine learning studies.

4.2 Configuration of the LAPD

Data collection was conducted in two campaigns separated by 14 months. The initial run set is designated as DR1 and the subsequent run set as DR2. These run sets are further broken down into *dataruns* which are series of discharges (“shots”) with identical operational machine parameters. A total of 67 *dataruns* were collected over both campaigns. These two *datarun* sets had significant intrinsic differences: DR1 had two turbomolecular pumps offline, leading to much higher background pressures. Furthermore, the cathode condition in terms of emissivity or asymmetries is unquantified, so there may be intrinsic differences in the plasma produced regardless of machine configuration.

The LAPD control parameters varied in this dataset were the source field, mirror field, midplane field, gas puff valve voltage, gas puff duration, and discharge voltage. The magnetic field regions are labeled in fig. 4.2 and effectively control the width of the plasma relative to the cathode in their respective regions. The gas puff voltage governs gas flow rate into the chamber, though this relationship is not yet quantified, and the gas puff duration defines the piezo valve activation period. For DR1, the discharge voltage is applied across the cathode and anode at the same time as the gas puff, but for DR2 applied 10 ms after gas puff initiation. This difference between runs was not known at experiment time. While discharge voltage correlates to discharge current (and thus power), the current depends on the machine configuration and downstream conditions and cannot be predetermined.

These machine parameters – with the exception of gas puff duration – were randomly sampled via Latin-hypercube sampling (LHS) for 44 of the *dataruns*. LHS is a pseudorandom sampler that guarantees that each machine setting is set at least once. An example of LHS vs random sampling can be seen in fig. 4.1. It is possible for random sampling to miss certain machine settings, or entire portions of configuration space altogether. This fact is particularly important when the number of samples is small, such as in this case with 44 samples. Data were then collected with these settings.




Figure 4.1: An example of Latin Hypercube Sampling compared to a potential random sample of five points. LHS hits all rows and columns, but random sampling may leave some sections of parameter unsampled space altogether.

Gas puff duration was reduced for the last seven runs to 20, 10, or 5 ms (see fig. 5.3 for timings relative to I_{sat} signals). The breakdown of each setting in the dataset is given in appendix 4.5, Table 4.1. The top gas puff valve was used for only the first nine dataruns of DR2 because of equipment issues. 23 of the dataruns in the dataset are not random: they were chosen to be similar to common machine configurations used in more conventional studies, usually using flat fields (or different cathode fields) around 1 kG. These data were taken while other diagnostics were being configured.

I_{sat} and other probe-based measurements were acquired along $y=0$ lines (51 dataruns total) or x-y grids (16 dataruns total) with spatial resolutions varying between 1.5 to 2 cm. The fixed axial locations of the probes were 895 cm and 831 in DR1 and 1150, 1022, 863, and 639 cm for DR2 (Fig. 5.1). Six shots were recorded at each position except for the first four dataruns in DR1 with five shots each.

4.3 Signals collected

DR1 and DR2 had considerable overlap in diagnostics recorded, with some minor differences. A summary of the diagnostics and their locations on the LAPD can be seen in fig. 4.2. Some of the raw diagnostics signals and machine state information (MSI) can be seen in fig. 4.3. Some dataruns may not contain all diagnostics, as some data were collected while other diagnostics were being set up. The diagnostics and machine state information (MSI) recorded for this dataset are the following:

- **DR1 probes:** three probes were inserted into the LAPD. One had Langmuir sweeps, another “flux probe” had I_{sat} and two floating potential (V_f) tips, and the last “triple probe” had I_{sat} , V_f and electron temperature (T_e). These signals were digitized at 6.25 MHz (100 MHz, 16 sample average).
- **DR2 probes:** four probes were inserted, namely a flux probe, triple probe, Langmuir sweeps with I_{sat} on a separate tip, and another flux probe. These signals were digitized at 6.25 MHz (100 MHz, 16 sample average).
- **Diodes:** five diodes, axially distributed, were recorded as well. The one closest to the cathode had a He-II line filter. The diodes were uncalibrated, have a nonlinear response, and are sensitive beyond the visible spectrum. These diodes were a part of the MSI system and were recorded at 25 kHz. Each diode (besides the one with the He-II filter) had 8 layers of 1-stop (50% transmission) neutral density filter in front of the diode.
- **Interferometer:** signals from the 288 GHz heterodyne interferometer were recorded on an oscilloscope at 10 MHz, which was then downsampled to 100 kHz before analysis so that the processing computer could keep pace.
- **Thomson scattering:** a single point was measured on-axis at port 32, triggered at 8 ms into the plasma for DR1 or 12 ms for DR2. Periodically the collection optics were scanned to maximize the light collected during both run sets.

- **Spectrometer:** an Ocean Optics HR4000 spectrometer recorded spectra integrated over the duration of the shot. The spectrometer has a very narrow slit, leading to good spectral resolution but requiring many shots for a clean spectrum.
- **Monochromator (DR2 only):** three Helium neutral lines were recorded, namely 587, 667, and 707 nm, using an oscilloscope sampling at 1 MHz.
- **Diamagnetic loop:** the loop sits between ports 34 and 35 and consists of one large loop and two smaller concentric loops equaling the area of the large one. These signals were digitized using an oscilloscope at 500 kHz, but are strongly influenced by magnet power supply noise making analysis difficult.
- **Fast framing camera:** a Phantom v7.3 fast framing camera recorded plasma dynamics from the end of the machine, pointing towards the cathode. The images are monochrome, 14 bit, 14 μ s exposure, 256×256 pixels, and 2,500 fps using a 105 mm lens. The camera is capable of 36,697 fps at that resolution, but a lower one was used to lessen file transfer times and storage requirements.
- **Discharge current and voltage:** as part of the MSI system, time series of discharge current and voltage are recorded at 25 kHz.
- **Magnetic field profiles:** theoretical on-axis magnetic field values are calculated using the magnet power supplies. Both are recorded as part of the MSI. For the work here, we simply use the programmed field values for the cathode, mirror, and midplane regions. Occasionally the calculated field would be incorrect since the power supply currents for the cathode are set manually, which is the case in some dataruns here, but the profiles are unused in these studies so it isn't an issue.
- **Pressures:** total pressure and pressure breakdown by atomic mass unit are recorded by an ion gauge and RGA, respectively. The RGA takes approximately two minutes to complete a scan but the data should be reasonably accurate given the slow time-evolution of pressure.

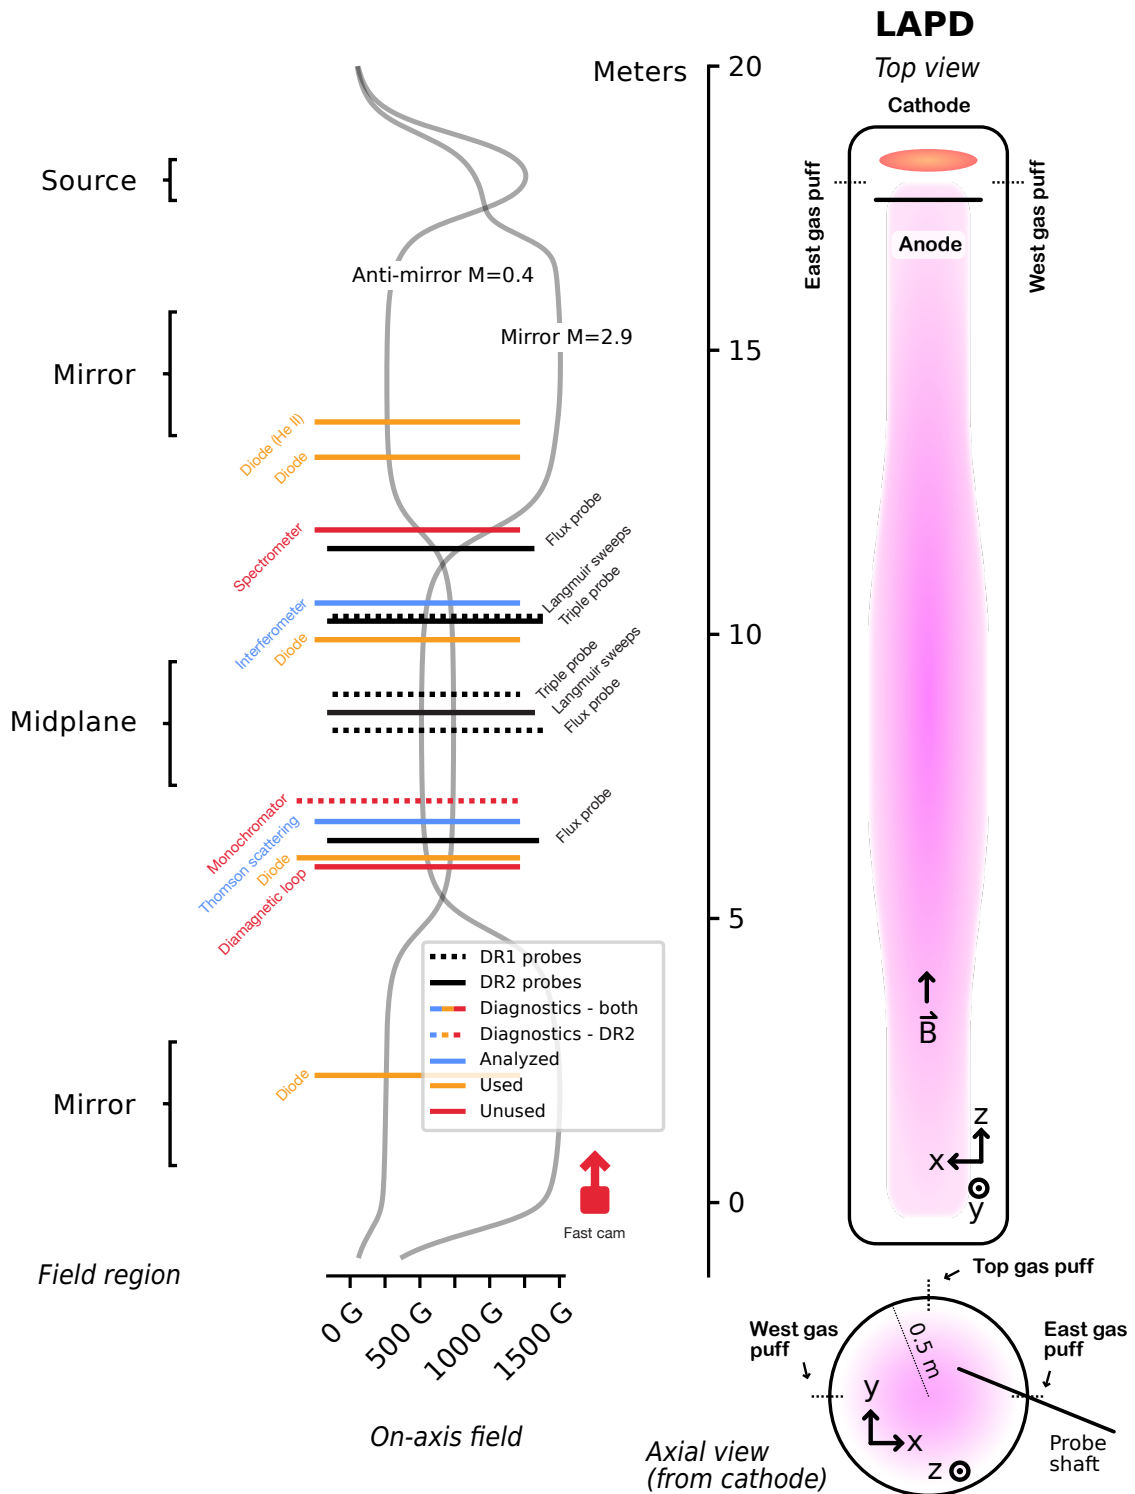


Figure 4.2: Diagnostics and example field configurations for DR1 and DR2

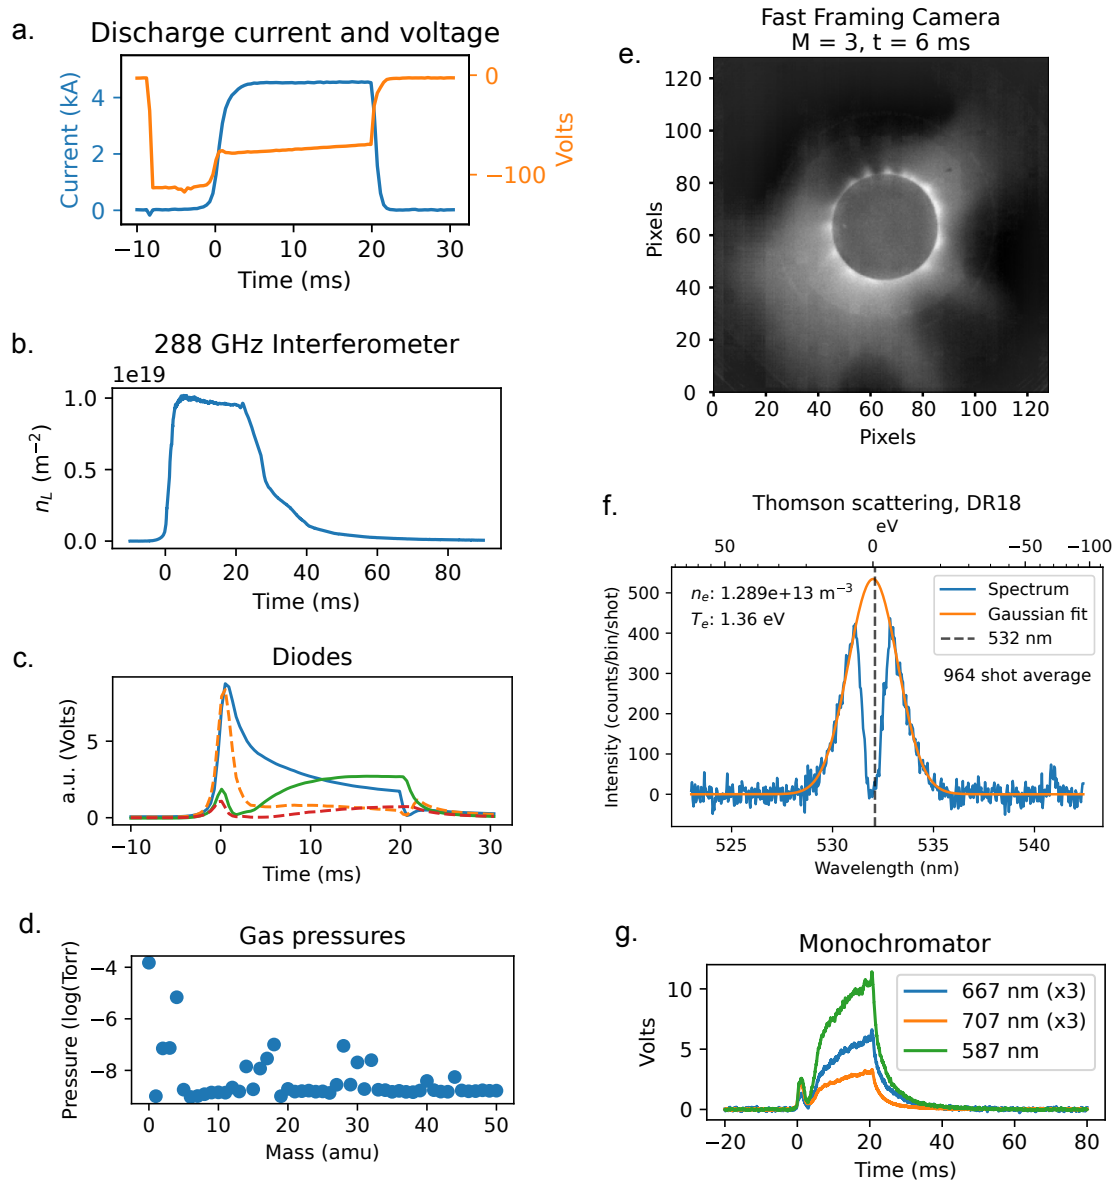


Figure 4.3: Example diagnostic signals and machine state information from a variety of discharges.

Of the probes, only I_{sat} was analyzed and used. The interferometer and Thomson scattering signals were also analyzed. The diode signals were unanalyzed but used in a downstream machine learning study. The spectrometer, monochromator, and diamagnetic loop remain unanalyzed and unused, but the raw signals could be useful for ML studies as will be shown with the diode signals (see chapter 6). The fast framing camera was useful for checking probe alignment and visualizing plasma structure, but it was otherwise not used or analyzed for the downstream ML studies.

4.4 Data cleaning

I_{sat} measurements in DR1 that saturated either the isolation amplifier or digitizer are excluded from the dataset. Only 484 shots were removed out of $\approx 132,000$, so the impact on the aggregate dataset is minimal. This signal saturation was detected while data were being taken and was corrected quickly.

Interferometer skips were occasionally seen, likely caused by large $\delta n/n$ structures combined with downsampling before conversion of the signal into a density measurement. Attempts were made to unwrap these skipping traces (see fig. 4.4) but without much success, so these shots were cut from the dataset.

The Thomson scattering (TS) diagnostic was available only for dataruns 8 and onwards in DR1. The TS image data did not have timestamps recorded, so a rough estimate was used based on filename and last saved time. Uncertainty in time is tolerable because conditions were identical to datarun shots for a few minutes before and after the dataruns. Fits were taken of the average over the entire datarun; each shot in a datarun has the same recorded TS temperature and density. Dataruns were removed if the error on the density, measured by the square root of the covariance of the fit amplitude, was greater than 50%. Fits above that error threshold were largely unusable. A couple of dataruns looked like pure noise even when averaged over several hundred shots, but were not caught by this broad criterion. 24 dataruns remained out of 30. In some runs there was high-pixel-frequency noise at 128 and 256 (every 4th and 2nd pixel, respectively). The fitting




Figure 4.4: An example of the interferometer skip (blue) and the attempted unwrap (orange).

routine is typically insensitive to this cleaning process, but significant differences can be seen in particularly low-density plasmas where the photon counts are low. An example of this process can be seen in fig. 4.5.

4.5 Data bias

Data bias and imbalance in the training set can be exacerbated by the train-test split. For the nominal test set, 8 out of the 67 dataruns were hand-picked for diversity and held out from the training set. Leaving out entire dataruns – not just shots – is important in order to estimate model performance on new, unseen discharges in new configurations. Four dataruns from each run set were left out: for DR1 08, 15, 23, and 33; for DR2, 02, 10, 19, and 31. As will be demonstrated in chapter 5, this test set appears to characterize the model performance on held out data fairly well.

The dataset predominantly contains gas puff durations of 38 ms. Only six runs in the training set have gas puff durations less than 38 ms: three have 5 ms and three have 10 ms, each having mirror ratios 1, 3, and 6 but otherwise identical configurations in an attempt to see mirror-related




Figure 4.5: The Gaussian fit to the Thomson scattering spectrum before and after crude FFT filtering. The shaded gray region is excluded from the fitting process because they contain the region of the notch filter (the region about 532 nm) and a He-II (ion) line (≈ 541 nm).

interchange instabilities in higher-temperature, lower-collisionality regimes. The 20 ms gas puff duration case is in the test set (DR2 31). This sampling bias towards the 38 ms gas puff duration suggests poor model performance is to be expected in shorter gas puff regimes. The top gas puff valve was operational for only the first nine runs of DR2.

Despite the best efforts to randomize the machine configuration, imbalance in the dataset will be present because of the relatively small amount of samples for the given actuator space. The distribution of I_{sat} signals, averaged from 10 to 20 ms, can be seen in Fig. 4.6. The I_{sat} distribution is clearly different for DR1 and DR2, with DR1 having a much flatter distribution. These distributions imply that if the model is constrained to sample from DR2 via the run set flag, then the model is expected to predict a lower I_{sat} value in general. When predicting from the model in general, performance will likely be worse for I_{sat} values $\gtrsim 11$ mA/mm². The time-averaged I_{sat} distribution is dissimilar between DR1 and DR2: DR1 appears to have a more uniform distribution. Combining the two datasets results in many I_{sat} examples less than 2 mA/mm² and a sharp decrease in number




Figure 4.6: Distribution of I_{sat} signals when averaged from 10 to 20 ms. DR1 appears to have a more uniform distribution than DR2 does. Combining the two datasets results in many I_{sat} examples near 0 mA/mm² and a sharp decrease in number of examples above 10 mA/mm². From these histograms we expect our model to be biased towards fitting lower I_{sat} values better, and to perform poorly in cases with very high I_{sat} values.




Figure 4.7: Distribution of the x-coordinate in the profiles. The increase in data points between roughly $x \approx 0$ to 30 cm is from planes instead of lines. Based on this distribution, the performance of the model is expected to be biased towards this central area.

of examples above 10 mA/mm^2 . Thus, we expect the model to perform better for smaller I_{sat} values than larger ones.

The mix of different probe movements also leads to some imbalance in the dataset. The distribution of probe positions can be seen in fig. 4.7. Notably, samples appear to drop off beyond +25 cm and -15 cm. Measurements over an x-y plane, constituting $\approx 64\%$ of all shots, are predominantly acquired overnight for maximal machine utilization. These longer dataruns lead to particular machine configurations being overrepresented in the dataset.

The distribution of the selected machine settings for all the dataruns is enumerated in Table 4.1. Despite the randomization of the settings of 44 dataruns, the distribution is often uneven. The remaining 23 non-random dataruns also contribute to the imbalance. For example, a source field of 1 kG and discharge voltage of 112 show up disproportionately in the dataset because data were collected at those settings while other equipment was being adjusted or calibrated.

Table 4.1: Data breakdown by class and dataset (percent)

| B source (G) | | | | B mirror (G) | | | | B midplane (G) | | | |
|--------------|-------|-------|-------|--------------|-------|-------|-------|----------------|-------|-------|-------|
| | Train | Test | All | | Train | Test | All | | Train | Test | All |
| 500 | 4.77 | 0 | 4.29 | 250 | 4.30 | 8.41 | 4.72 | 250 | 8.25 | 21.01 | 9.55 |
| 750 | 3.34 | 12.61 | 4.29 | 500 | 30.49 | 8.41 | 28.23 | 500 | 43.80 | 8.41 | 40.19 |
| 1000 | 43.13 | 78.99 | 46.78 | 750 | 6.68 | 16.81 | 7.72 | 750 | 6.62 | 52.19 | 11.27 |
| 1250 | 12.59 | 0 | 11.30 | 1000 | 28.85 | 57.97 | 31.82 | 1000 | 26.36 | 5.78 | 24.26 |
| 1500 | 19.23 | 0 | 17.27 | 1250 | 3.34 | 4.20 | 3.43 | 1250 | 9.24 | 0 | 8.30 |
| 1750 | 1.91 | 0 | 1.71 | 1500 | 26.34 | 4.20 | 24.08 | 1500 | 5.73 | 12.61 | 6.43 |
| 2000 | 15.03 | 8.41 | 14.35 | | | | | | | | |

| Gas puff voltage (V) | | | | Discharge voltage (V) | | | | Axial probe position (cm) | | | |
|----------------------|-------|-------|-------|-----------------------|-------|-------|-------|---------------------------|-------|-------|-------|
| 70 | 12.11 | 16.81 | 12.59 | 70 | 12.22 | 8.41 | 11.83 | 639 | 12.48 | 8.41 | 12.06 |
| 75 | 6.68 | 0 | 6.00 | 80 | 5.25 | 0 | 4.72 | 828 | 17.07 | 36.28 | 19.03 |
| 80 | 11.46 | 8.41 | 11.15 | 90 | 2.86 | 8.41 | 3.43 | 859 | 12.48 | 8.41 | 12.06 |
| 82 | 41.49 | 57.97 | 43.17 | 100 | 3.34 | 8.41 | 3.86 | 895 | 33.01 | 30.10 | 32.71 |
| 85 | 14.13 | 0 | 12.69 | 110 | 8.77 | 0 | 7.87 | 1017 | 12.48 | 8.41 | 12.06 |
| 90 | 14.13 | 16.81 | 14.40 | 112 | 20.62 | 0 | 18.52 | 1145 | 12.48 | 8.41 | 12.06 |
| | | | | 120 | 3.82 | 8.41 | 4.29 | | | | |
| | | | | 130 | 0.95 | 0 | 0.86 | | | | |
| | | | | 140 | 2.86 | 8.41 | 3.43 | | | | |
| | | | | 150 | 39.30 | 57.97 | 41.20 | | | | |

| Gas puff duration (ms) | | | | Vertical probe position (cm) | | | |
|------------------------|-------|-------|-------|------------------------------|-------|-------|-------|
| 38 | 94.27 | 91.59 | 94.00 | ≈ 0 | 36.26 | 46.08 | 37.26 |
| < 38 | 5.73 | 8.41 | 6.00 | $\neq 0$ | 63.74 | 53.92 | 62.74 |




Figure 4.8: An example of the y-axis profile before and after shifting the y-coordinate. The “before” plot (top) is obviously asymmetrical about $y=0$. The shift needed to center was eyeballed from the plot. Each line represents a different x position, from closest to the core (upper lines) to the edge (lower lines).

4.6 Azimuthal asymmetry of probe data

Examining the data, it appears that the y coordinate is not centered properly, possibly because the telescope used to align the probes is set incorrectly. Using profiles from planar data (see the “before” plot in fig. 4.8), the y -coordinate was adjusted. The probes in DR1 were adjusted upward by 2 cm. For DR2, the y -coordinate was adjusted separately for each probe. Port 17 was adjusted 6 cm up, port 21 was adjusted 4 cm up, port 26 was adjusted 4.5 cm up, and port 33 was adjusted 3.35 cm up. This degree of error is consistent with a centering scope crosshair angle error, which would cause a larger absolute y -axis error closer to the cathode. An example of this y -axis error and the profile after shifting the coordinates can be seen in fig. 4.8. It is likely that this y -axis offset applies to other probes in the plasma, not just probes with I_{sat} tips.

4.7 Applying and improving the dataset

The two chapters following this detail machine learning studies utilizing this dataset, though only using a subset of the diagnostics available. Significant opportunities remain for ML-based analysis of the dataset, such as including additional diagnostics, in-situ diagnostic calibration (e.g., I_{sat} or Thomson scattering). Even though the diversity of the dataset is relatively high, many imbalances in machine inputs remain. More data with additional (pseudo-)random samples from broader parameter ranges would be very beneficial in improving downstream ML tasks. Pushing the boundaries of the machine parameters could also lead to discovery and exploitation of new operational modes of the LAPD which could prove beneficial.

CHAPTER 5

Optimizing mirror configurations in the LAPD using machine learning

This study demonstrates the efficacy of machine learning (ML)-based trend inference using data from the Large Plasma Device (LAPD). Neural network (NN) ensembles with uncertainty quantification are trained to predict time-averaged ion saturation current (I_{sat} — proportional to density and the square root of electron temperature) at any position within the dataset domain. Model-inferred trends, such as the effects of introducing mirrors or changing the discharge voltage, are consistent with current understanding. In addition, axial variation is optimized via comprehensive search over I_{sat} predictions. Experimental validation of these optimized machine parameters demonstrate qualitative agreement, with quantitative differences attributable to Langmuir probe variation and cathode conditions. This investigation demonstrates, using ML techniques, a new way of extracting insight from experiments and novel optimization of plasmas. The code and data used in this study are made freely available.

The primary goals of this work are to provide an example of a solid, validated machine learning study and demonstrate how ML can be useful in understanding operating plasma devices.

5.1 Introduction

Understanding and controlling plasma behavior in fusion devices is necessary for developing efficient fusion reactors for energy production. Because of the complex, high-dimensional parameter space, traditional experimental approaches are often time-consuming and require careful planning.

This work explores how machine learning (ML) techniques can accelerate this understanding by elucidating the effect of machine parameters in a basic magnetized plasma device. Trend inference is this process of relationship discovery. While ML methods, particularly neural networks (NNs), have become increasingly prevalent in fusion research for control and stabilization, their application to systematic trend discovery remains largely unexplored.

Many studies have used ML for profile prediction on a variety of tokamaks, particularly for real-time prediction and control. For example, NNs were used to predict electron density, temperature, and other quantities in DIII-D [4], and reservoir NNs have demonstrated the ability to quickly adapt to new scenarios or devices [79]. Temporal evolution of parameters has been successfully modeled using recurrent neural networks (RNNs)[24] for multiple devices, including the EAST[162] and KSTAR tokamaks[146, 144]. These predictions enabled training of a reinforcement learning-based controller[146, 144]. In addition, a decision tree-based controller was trained to maximize β_N while avoiding tearing instabilities[57] on DIII-D. Electron temperature profiles have also been predicted using dense NNs on the J-TEXT tokamak [36].

A parallel focus has been on instability prediction and mitigation in tokamaks, particularly of disruptions. Notable achievements in disruption prediction include RNN-based disruption prediction [86] and random forest approaches[130], with a comprehensive review available by Vega et al [158]. Recent work has extended to active control, such as the mitigation of tearing instabilities in DIII-D using reinforcement learning [145].

While ML has proven effective for prediction and control tasks, inferring trends using data-driven methods has been relatively uncommon. Notable exceptions include finding scaling laws on the JET tokamak[110] via classical ML techniques and the development of the Maris density limit[105] which outperforms other common scalings (including the Greenwald density limit) in predictive capability.

The use of machine learning and Bayesian inference in fusion research has been recently reviewed by Pavone et al.[118]

Outside of magnetized plasmas, the laser plasma community has embraced ML techniques for various applications, enumerated in a review by Dopp et al.[43]. Data-driven plasma science in general has been reviewed by Anirudh et al.[6] Notably, a similar quasi-random method (Sobol sequences) was used to collect a spectroscopy dataset on a plasma processing device over diverse machine settings [30]. This process is similar to what is performed in our work here, but a generative variational autoencoder was instead trained to be used as an empirical surrogate model.

This work advances data-driven methods in plasma physics by taking these methods one step further: instead of learning a model for particular task (e.g., disruption prediction or profile prediction), we infer learned trends directly from the model itself.

The goal of this study is to develop a data-driven model that can provide insight into the effect of machine parameters on plasmas produced in Large Plasma Device (LAPD) in lieu of a theoretical model. In contrast with tokamaks and other fusion devices, the LAPD is particularly well-suited for ML data collection because of its flexibility and high repetition rate. We demonstrate the capability to infer trends in a particular diagnostic signal, the time-averaged ion saturation current (I_{sat}), for any mirror (or anti-mirror) field geometry in a variety of machine configurations. Langmuir probes are commonly used to measure density, temperature, and potential in virtually all plasma devices in low-temperature (less than 10s of an eV) regimes. The I_{sat} signal in particular is almost always used in the LAPD for calculating local plasma density.

This study marks two firsts in magnetized plasma research: using NNs to directly infer trends and collecting data efficiently with partially-randomized machine parameters. We also demonstrate optimizing LAPD plasmas given any cost function by minimizing axial variation in I_{sat} . This global optimization is only possible using ML techniques. This work demonstrates the usefulness of a pure ML approach to modeling device operation and shows how this model can be exploited. We encourage existing ML projects and experiments to consider this approach if possible. Acquiring sufficiently diverse datasets may require assuming some risk because diverse data, such as discharges from randomly sampled machine settings, may not be amenable to conventional analysis techniques.

All the processed data used for training the models in this study are freely available[152] (see

Large Plasma Device

Top view




Figure 5.1: A cartoon of the Large Plasma Device, the coordinate system used, examples of a mirror and anti-mirror magnetic field configuration, and probe locations used in this study. The source, mirror, and midplane regions are labeled; the three fields were programmed independently.

section 5.11). Other devices have also made data publicly available. In particular, data for H-mode confinement scaling has been available since 2008[131], and more recently some MAST[78] and all LHD[2] data are now publicly available.

5.2 Processing of I_{sat} signals

The ion saturation current, denoted as I_{sat} , is obtained by applying a sufficiently negative bias to a Langmuir probe to ensure the exclusive collection of ions. This collected current is proportional to $Sn_e\sqrt{T_e}$, where n_e and T_e are the electron density and temperature, and S is the effective probe collection area. To account for differences in probe tip geometry, the I_{sat} values are normalized to area.

I_{sat} measurements were averaged over 10 to 20 ms to exclude plasma ramp-up and fluctuations. Example I_{sat} probe data can be seen in fig. 5.3 along side gas puff timings. For the probe tip that




Figure 5.2: I_{sat} traces from the swept probe (port 26) from DR2 datarun 03, shot 1 of 6. The orange curve is excluding times when a sweep is active on an opposing tip.

was on the same shaft as the swept probe (in DR2), the signal was instead averaged over when the bias voltage on the swept tip was held constant at the lowest value. A $40 \mu\text{s}$ (250 sample) buffer was used after the sweep was turned off to minimize the impact of transient conditions. A comparison of the full trace and the trace with the swept portion excluded can be seen in fig. 5.2. Notably, the measured I_{sat} value does not attain a steady state before the discharge shuts off.

Profile evolution is not studied to minimize computational requirements. I_{sat} characteristics vary significantly between axial (z) position machine parameters. For I_{sat} measurements on the same probe as a Langmuir sweep (DR2 port 26, $z=863 \text{ cm}$), the averaging process excludes the sweep period with an additional $40 \mu\text{s}$ buffer.

While I_{sat} exhibits a small degree of shot-to-shot variation, the present model only learns the expected value, leaving distributional modeling to future generative approaches. An example of these I_{sat} profiles and the six-shot variance can be seen in fig. 5.4.




Figure 5.3: Gas puff timings and example I_{sat} time series at three different z -axis locations from three different dataruns. Note that some discharges do not achieve steady state in I_{sat} .




Figure 5.4: Examples of I_{sat} profiles from DR2 run 15. The bars represent the minimum and maximum of the six I_{sat} measurements taken at that position.

5.3 Model development and training

5.3.1 Model inputs

Neural network inputs comprise 12 variables: source field, mirror field, midplane field, gas puff voltage, discharge voltage, gas puff duration, probe coordinates (x , y , z), probe rotation, run set identifier, and top gas puff flag. These variables can be interpreted as six control parameters, four probe coordinates, and two flags. These inputs are mean-centered and normalized to the peak-to-peak value with no outliers in the dataset. The baseline models trained in section 5.3.4 did not contain the run set identifier or top gas puff flag.

5.3.2 Training details

For initial experiments in training the model, a mean-squared error (MSE) loss is used:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{m} \sum_{i=1}^m (f(x_i) - y_i)^2 \quad (5.1)$$

where x_i represents the input vector for the i th example, y_i the target measurement, m the batch size, and f the NN. During training, overfitting was assessed via the validation set MSE with a traditional 80-20 train-validation random split. Unless stated otherwise, a dense neural network, 4 hidden layers deep and 256 units wide (201,218 parameters for β -NLL loss, 200,962 parameters for MSE loss), was trained with AdamW using a learning rate of 3×10^{-4} . Leaky ReLU activations (the nonlinearities in the NN) and adaptive gradient clipping[141] (cutting gradients norms above the 90th percentile of recent norms) were used to mitigate vanishing gradients and mitigate exploding gradients, respectively. The models were evaluated after training concluded at 500 epochs.

5.3.3 Validating the training pipeline

ML training processes are relatively simple but bugs, particularly in the data pipeline, can be insidious and can affect final model performance even though training looks fine. Here we validate the data pipeline (which should be performed in every ML study) to verify that the model is training and expected and that there is no accidental data leakage between the train and test sets. Andrej Karpathy’s advice for training neural networks [85] was used as a template for verifying the training procedure used in this project. The data fed into the model immediately before the forward pass (and subsequent backpropagation) was stored and verified: the data are correctly randomly shuffled in each batch. Each epoch contains the same random shot order. To validate the data pipeline, a simple dense model (4 layers, 512 wide with one output; 794113 parameters, tanh activations) was trained. The model is also intentionally overfit on a single batch (128 examples) of training data to ensure that training progresses as expected. A deep double descent is observed as expected [111, 136]. Training on a batch of 8 examples reaches ≈ 0 training loss after 50 steps. Plots of the train and validation losses can be seen in Fig. 5.5.




Figure 5.5: Training and validation losses when overfitting the model. A deep double descent in the validation losses is observed when fitting a single batch of 128 examples. The 8-example batch hits near-zero loss after 50 steps. This process verifies our training process is functioning as expected. The spikes are from exploding gradients which can be mitigated by clipping the gradients. A model trained on blank data is also shown as the black dotted line.

Multiple models were trained with varying depths and widths to verify that training loss decreases with increased model capacity. Doubling the layer width from 512 to 1024 moderately decreases the training loss; doubling the depth of the network from 4 to 8 layers has a larger impact. Increasing the width further to 2048 and depth to 12 layers has a dramatic impact on training loss, so this model and dataset are behaving nominally. The model pipeline is training and performing as expected, so we proceed.

5.3.4 Baselines for mean-squared error

A model was first trained with zeroed-out inputs as a baseline and to validate the data pipeline. This model effectively has only a single, learnable bias parameter at the input. This process yields a

validation loss (simply MSE in this case) of 0.036.

A linear model obviously cannot fit the dataset (see the nonlinear shape of the profiles in fig. 5.4). However, a simple (and mostly linear) model can provide a performance baseline to help spot bugs when training more complex models. Since the x- and y- profiles have a approximate tanh shape, a feature is added at the linear model input stage for the x and y coordinates: $x_{\text{tanh}} = c \cdot \tanh(|x + s| \cdot a + b)$ where s, a, b, c are trainable parameters (independent for each coordinate; c is superfluous). This function was chosen to give the linear model the capability of expressing tanh-like curves. The performance of the linear model on DR2 data, with and without the tanh features, can be seen in fig. 5.6. This baseline linear-like model reaches a training and validation loss of around 0.011, with the $\text{RMSE} = \sqrt{\text{loss}} \sim 0.1$. The linear-only model is marginally worse with losses at around 0.014.

This feature engineering-like approach can continue. For example, the width of the profile is largely controlled by magnetic field configuration of the device, particularly by $\sqrt{B_{\text{midplane}}/B_{\text{source}}}$. This behavior can be added to this model, either as a new feature or as a custom relationship in the model. Note that, as seen in fig. 5.6, the width of the profile also depends on the axial coordinate. Combined with other coordinates and actuators, like discharge voltage and gas pressure, the number of possible features or function space grows combinatorially, making this custom fitting process difficult and tedious to design and test by hand. The obvious solution would be to use symbolic regression or fitting to a function library which may be ideal methods if simple profile prediction were the final goal. However, we are ultimately interested inferring trends in a much more complex input space where neural networks are more flexible and accurate. If NNs do face generalization issues, symbolic regression or a SINDy-like approach can used instead, albeit with limited applications. Symbolic methods are appealing because the fits are simple. However, even though a simple equation may fit the data well, it does not necessarily provide insight or relate to the underlying physics; using a freeform fitting function like a neural network is more appropriate in this use case.

A summary of these baselines is seen at the top of Table 5.1.




Figure 5.6: I_{sat} profiles and predictions for ports 17 and 33 based on inputs from DR2 run 08 using a liner and linear-with-tanh models. DR2 run 08 is in the training set. The “data” points are averaged over six shots. Run 08 was chosen for its representative performance; ports 17 and 33 were chosen to demonstrate the maximal axial variation (across 511 cm). These models fail to describe the data accurately.

Table 5.1: Summary of test set losses for different training data and ensembles

| Model | MSE $\times 10^{-3}$ |
|---------------------------|----------------------|
| Zeroed-input | 36 (validation) |
| Linear model | 14 (validation) |
| Linear with tanh features | 11 (validation) |
| 9 dataruns | 7.0 |
| 19 dataruns | 6.9 |
| 29 dataruns | 4.2 |
| 39 dataruns | 4.1 |
| 49 dataruns | 3.4 |
| DR1 only | 6.4 |
| DR2 only | 5.4 |
| Full set, large model | 2.8 |
| Full set average | 3.6 ± 0.56 |
| Full set ensemble | 2.9 ± 1.1 |
| “Run set” flag average | 2.1 ± 0.15 |
| “Run set” flag ensemble | 1.9 ± 0.64 |
| “Top gas puff” flag | 1.8 |

5.3.5 Effects of training set and model sizes

To study the effects of reduced diversity, the number of unique dataruns in the training set was systematically reduced while evaluating on a fixed test set. The test set loss monotonically increased with this decrease in datarun count. Part of this decrease in performance may be caused by a simple reduction in training set size. In addition, models were individually trained and evaluated on DR1 only or DR2 only. When evaluated on the left-out run set, the test set losses were high, near or above the zero-input baseline of 3.6×10^{-2} . This result suggests that both run sets contain significant information missing in the other, and training on both provides beneficial information on the structure of the I_{sat} measurement despite different probe calibrations and cathode state.

A larger model, consisting of a 12-deep 2048-wide dense network, was trained on the full training dataset, evaluated at 30 epochs. This larger model yielded a test MSE of 2.8×10^{-3} , indicating that these NNs are behaving as expected. Longer training or larger models may yield better test set results, but will likely not come close to the training and validation losses which are on the order of 10^{-5} . Combined models with differing initializations (an ensemble), were trained to measure the MSE variance over model parameters which was about 16%. When the I_{sat} predictions were averaged, the test set MSE was $2.9 \pm 1.1 \times 10^{-3}$, achieving the best performance for that model size. These test set losses are also seen in Table 5.1.

5.3.6 Improving performance with machine state flags

Data from DR1 and DR2 were collected 14 months apart leading to differing machine states. In DR1, only one turbo pump was operating leading to much higher neutral pressures than in the DR2 run set. A new parameter (mean-centered and scaled) was added to the inputs to distinguish between these two run sets. All the predictions in this study use the DR2 run set flag (a value of 1.0) because turning off the turbopumps is not a commonly desired mode of operating the LAPD. The inclusion of this parameter also provides the model the ability to distinguish between the probe calibration differences between DR1 and DR2. An ensemble prediction with this run set flag brings the test set

MSE down to 1.9×10^{-3} .

A flag indicating when the top gas puff valve was enabled in DR2 was also added to all training data, allowing the model to further distinguish between different fueling cases. The addition of this flag incrementally improved test set MSE to 1.8×10^{-3} . The effect on MSE on the inclusion of these new parameters is compared to the performance of other models in Table 5.1.

5.3.7 Learning rate scheduling

Modifying the learning rate over time (scheduling) is known to improve model learning. The following schedules were compared: constant learning rate ($\gamma = 3 \times 10^{-4}$), $\gamma \propto \text{epoch}^{-1}$, $\gamma \propto \exp(-\text{epoch})$, and $\gamma \propto \text{epoch}^{-1/2}$. The epoch is the training step divided by the number of batches in one epoch, so “epoch” in this case takes on a floating-point value. $\gamma \propto \text{epoch}^{-1}$ appears to give the best test set loss by a test MSE difference of 1×10^{-4} , and any schedule beats a constant learning rate by a difference of $2 - 4 \times 10^{-4}$.

5.4 Uncertainty quantification

5.4.1 β -NLL loss

Instead of predicting a single point, the model can predict a mean μ and variance σ^2 using the negative-log likelihood (NLL) loss [114, 95] by assuming a Gaussian likelihood. An adaptive scaling factor $\text{StopGrad}(\sigma_i^{2\beta})$ is introduced that can be interpreted as an interpolation between an MSE loss and Gaussian NLL loss, yielding the β -NLL loss:

$$\mathcal{L}_{\beta\text{-NLL}} = \frac{1}{2} \left(\log \sigma_i^2(\mathbf{x}_n) + \frac{(\mu_i(\mathbf{x}_n) - y_n)^2}{\sigma_i^2(\mathbf{x}_n)} \right) \text{StopGrad} \left(\sigma_i^{2\beta} \right) \quad (5.2)$$

for example n and model i , with an implicit expectation over training examples. $\beta = 0$ yields the original Gaussian NLL loss function and $\beta = 1$ yields the MSE loss function. This factor improves

MSE performance by scaling via an effective learning rate for each example (which necessitates the StopGrad operation) [142], and improves both aleatoric and epistemic uncertainty quantification [155]. $\beta = 0.5$ was used by default in this study. This β -NLL loss function also improved training stability.

This NLL-like loss assumes the prediction – the likelihood of y given input \mathbf{x} : $p(y|\mathbf{x})$ – follows a Gaussian distribution. Treating each prediction as an independent random variable (considering each model in the ensemble is sampled from some weight distribution $\theta \sim p(\theta|\mathbf{x}, y)$) and finding the mean of the random variables yields a normal distribution with mean $\mu_*(\mathbf{x}) = \langle \mu_i(\mathbf{x}) \rangle$ and variance $\sigma_*^2 = \langle \sigma_i^2(\mathbf{x}) + \mu_i^2(\mathbf{x}) \rangle - \mu_*^2(\mathbf{x})$ where $\langle \rangle$ indicates an average over the ensemble.

The loss function for one of the NNs in an ensemble is seen in fig. 5.7. The MSE decreases monotonically for the training and validation set, but does not for the test set. The loss function can no longer be interpreted as a log-likelihood because of the effective per-example learning rate set by the β term in the loss (eq. 5.2). Note that early stopping (at around 8 epochs) would improve the test set loss, but the MSE would still be several factors higher than after 500 epochs. Early stopping was not explored in this study.

The ensemble predictive uncertainty can be decomposed into the aleatoric and epistemic components [155]: the aleatoric uncertainty is $\langle \sigma_i^2(\mathbf{x}) \rangle$ and the epistemic uncertainty is $\langle \mu_i^2(\mathbf{x}) \rangle - \mu_*^2(\mathbf{x}) = \text{Var}[\mu_i(\mathbf{x})]$. The intuition behind these uncertainties is that the random fluctuations in the recorded data are captured in the variance of a single network, σ_i^2 . If the choice of model parameters were significant, we would expect the predicted mean for a single model, μ_i , to fluctuate as captured by $\text{Var}[\mu_i(\mathbf{x})]$.

5.4.2 Cross-validation MSE

For cross-validation, multiple train-test set pairs were created. Test set 0 comprises deliberately chosen dataruns to encompass a diverse set of machine settings and probe movements. The other six datasets were compiled with randomly chosen dataruns (without replacement) while keeping

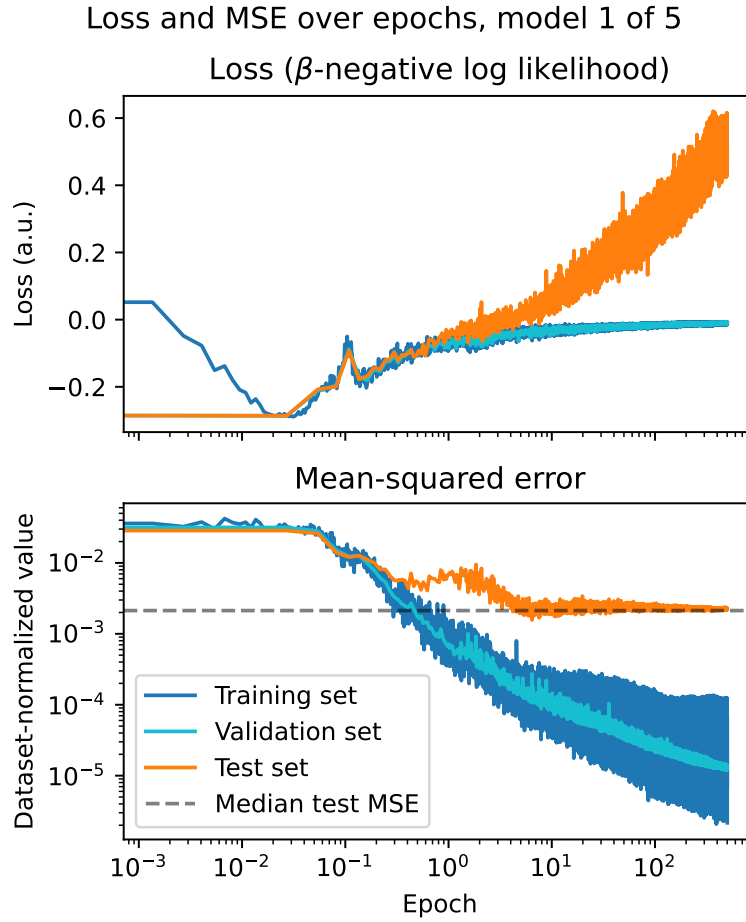


Figure 5.7: The loss and MSE for the training, validation, and test sets over the entire training duration of 500 epochs. The inclusion of the β term in the loss function – interpreted as a per-example learning rate – makes the loss function no longer interpretable in simple terms. The mean-squared error benefits from longer training for all sets.

the number of dataruns from DR1 and DR2 equal. Seven model ensembles (5 NNs per ensemble – 35 NNs total) were trained to evaluate the effect of test set choice on model MSE. The test set MSE performance can be seen in fig. 5.8, and the training MSE performance in fig. 5.9. The median ensemble test set MSE for these seven sets was 2.13×10^{-3} with a mean of 3.6×10^{-3} . The handpicked dataset had an ensemble test set MSE of 1.85×10^{-3} , indicating that the choice of dataruns was adequately representative. This median MSE will be used to estimate model prediction error in addition to uncertainty quantification. This cross-validation also provides an error estimate if the models were to be trained on *all* dataruns. Ensembles always out-performed the average error of single-model predictions.

All validation set MSEs fall between 1 and 6×10^{-5} , with the average training MSE falling within that range as well. These MSEs indicate that the model is able to fit the training data to a high degree of accuracy regardless of which dataruns are held out. The loss and MSE curves over training epochs can be seen in the appendix in fig. 5.7.

5.4.3 Model calibration via weight decay

The predicted uncertainty may not provide an accurate range of I_{sat} values when compared to the measured value. Calibrating the model means changing the predicted uncertainty range so that the measured values fall within that range according to some distribution, such as a Gaussian in this case. One of the ways assessing this calibration is by the z-score of predictions, where $z_n = (x_n - \mu_n) / \sigma_n(x_n)$ for example x_n , predicted mean μ_n , and standard deviation σ_n . Perfect model calibration would lead to identical z-score distribution $\mathcal{N}(\mu = 0, \sigma = 1)$ for the training and test sets. When evaluated on the training set, the distribution should be a Gaussian with a standard deviation of 1. The z-score distributions for the train and test sets with a model weight decay of 0 can be seen in fig. 5.10.

Increased weight decay can lead to better model calibration [67]. Weight decay penalizes large parameter values by adding the L2 norm of model weights to the loss. Model ensembles were




Figure 5.8: Model performance as measured by MSE over test sets with different dataruns. Test set 0 is the hand-picked dataset, and the rest were randomly compiled without replacement (though separate for DR1 and DR2). The number at the bottom of the bar chart is the number of shots in the testing set. The median test set performance is very close to the hand picked (set 0) performance. Ensembles always out-perform the average single-model prediction.




Figure 5.9: See caption for fig. 5.8. Note that the training loss is dramatically less than the testing loss, but otherwise there is no discernible trend.




Figure 5.10: Z-scores for the training and testing set for the model with a weight decay coefficient λ of zero. The magnitude of counts for the test set is scaled up by a factor of 8.8 (the train-to-test example ratio). The histograms are clipped between z of -8 and 8 with a bin width of 0.25; the spike at the negative side of the test set histogram is from the long tail.




Figure 5.11: Model performance and calibration for different weight decays. Highly biased models are better calibrated, but come at great expense of mean prediction error. At the weight decay value of 50, the model has worse error than a linear model. Note the linear scale below 10^{-2} .

trained with weight decay coefficients between 0 and 50 to determine the best calibrated model determined by the distribution of z-scores of the training and test sets. The results of this weight decay scan are seen in Fig. 5.11. Increasing the weight decay increases the test MSE and decreases its z-score standard deviation. This large standard deviation is caused by outliers. Excluding z-scores magnitudes above 10, or 4.4% of the test set, yields a standard deviation of 2.53. This long tail indicates that the distribution of predictions on the test set is not Gaussian. Nonetheless, the trend remains that increasing weight decay leads to smaller test set z-score standard deviations. However, the test set MSE increases after a weight decay of 1. This increase in test MSE implies that the model is making less accurate predictions but is better calibrated. Highly biased models are better calibrated, but come at great expense of mean prediction error. At the weight decay value of 50, the model has worse error than a linear model. Despite the attempts using weight decay, the model never becomes well-calibrated: the predicted uncertainty is always too low by a factor of 2 to 5.




Figure 5.12: Model extrapolation performance (top plots) with uncertainty (bottom plots) for a model ensemble trained on a β -NLL loss function. DR2 run 10 was chosen as an illustrative example. The *relative* uncertainty appears to be more useful when zero weight decay ($\lambda = 0$, left) is used: the uncertainty increases when the model is predicting outside its training data along the x -axis.

Despite the better calibration, the uncertainty predicted by a model with a large weight decay is decidedly worse: the uncertainty is similar across an entire profile, and when projected beyond the training data, the total uncertainty remains largely constant as seen in Fig. 5.12. The zero weight decay model exhibits relatively increasing uncertainty beyond the bounds of the training data. Although not well-calibrated, this uncertainty can provide a hint of where the model lacks confidence relative to other predictions, even though the uncertainty is much less than it should be.

5.5 Evaluating model performance

Model performance is evaluated in three ways by comparing against intuition from geometry, an absolute measurement, and extrapolated machine conditions.

5.5.1 Checking geometrical intuition

Assuming magnetic flux conservation, we know that modifying the mirror geometry can control the effective width of the plasma. One way to check that the model is learning appropriate trends is to check with this intuition. If the magnetic field at the source is not equal to the field at the probe, the probe will see the plasma expanded (or contracted) by roughly a factor of $\sqrt{B_{\text{probe}}/B_{\text{source}}}$. The cathode is about 35 cm in diameter, so a magnetic field ratio of 3 would give produce a plasma approximately 60 cm in diameter. All the probes used in this study are in or very close to the zero-curvature midplane region of a mirror.

To check this intuition, the model is given the following inputs: $B_{\text{source}}=500$ G, $B_{\text{mirror}}=1500$ G, $B_{\text{midplane}}=500$ G, discharge voltage=110 V, gas puff voltage=70 V, gas puff duration=38 ms, run set flag=DR2 and top gas puff=off. The discharge voltage and gas puffing parameters were arbitrarily chosen. The x coordinate is scanned from 0 to 30 cm, and the z coordinate from 640 to 1140 cm. This discharge is then modified by separately changing B_{source} to 1500 G and B_{midplane} to 750 G ($M=1.5$). The x profiles at the midplane ($z=790$ cm) of the reference $M=3$ prediction, source field change, and midplane field change, all scaled to cathode radius, can be seen in Fig. 5.13. Changing the source field to 1500 G increases the I_{sat} towards the edge of the plasma, as expected. When the midplane field is increased, the I_{sat} values decrease at the edge and increase at the core ($x=0$ cm), implying a thinner plasma column and is consistent with previously measured behavior. When only the mirror field is modified (not shown), the strongest effect on I_{sat} is on or near $x=0$ cm, and the plasma column width does not appear to appreciably change.




Figure 5.13: Plot of various mirror configurations scaled to the cathode radius $x_c = 17.5$ cm at the midplane ($z=790$ cm). When scaled according to the expected magnetic expansion, the profiles generally agree. The smaller the plasma diameter (and thus smaller volume), the higher the peak in I_{sat} at the core, as expected.

5.5.2 Directly comparing prediction to measurement

I_{sat} measurements were taken with the following LAPD machine settings: $B_{\text{source}}=1250$ G, $B_{\text{mirror}}=500$ G, $B_{\text{midplane}}=1500$ G, discharge voltage=90 V, gas puff voltage=90 V, gas puff duration=38 ms, run set flag=DR2 and top gas puff=off. These settings were from a previous discharge optimization attempt. The probes utilized were the permanently-mounted 45° probe drives. These probes were known to have identical effective areas relative to each other from the previous experiment and from analyzing the discharge rampup.

Because of data acquisition issues, only a single useful shot was collected at a nominal -45° angle (relative to the x-axis) 10 cm past the center ($x=0$ cm, $y=0$ cm) of the plasma on three probes at z-positions of 990, 767, and 607 cm (ports 22, 29, and 34, respectively). The probe drives were slightly uncentered, leading to the real coordinates of the probes to be around $x = 9.75$ cm and $y = -8.4$ cm. Note that the model can predict anywhere in LAPD bounded by the training data, so off-axis measurements are not an issue. The resulting predictions using these coordinates and machine conditions can be seen in Fig. 5.14. The model reproduces the axial trend well, but slightly underestimates I_{sat} on an absolute basis. However, given the lack of absolute I_{sat} calibration and variable machine state, the agreement of the absolute I_{sat} values may be coincidental. Nevertheless, the trend exhibited by this validation study matches the predicted trend and increases our trust in model predictions.

An additional validation datarun was performed. For this run, the discharge voltage was increased to 160 V, and the source field changed to 822 G. The training data contains discharge voltages up to 150 V, so this case tests the extrapolation capabilities of the model. The comparison of model predictions and the measured data can be seen in Fig. 5.14. As stated earlier, the absolute uncertainty provided of the model is not calibrated. However, note that the level of uncertainty provided by the model, as well as the large spread in model predictions, are much greater than seen in the interpolation regime (Fig. 5.14) and eclipses the cross-validated test set root mean squared error (RMSE). We conclude that this model has good interpolation capabilities, but extrapolation –




Figure 5.14: Top: data collected at off-axis positions around $x = 9.75$ cm and $y = -8.4$ cm are compared with predictions from the machine learning model at the same points in addition to two interpolating predictions. The model predicts the trend well, but underestimates I_{sat} in general. The shaded orange region is the total model uncertainty ($\sigma = \sqrt{\text{Var}}$). Bottom: Measured vs predicted I_{sat} values for an odd machine configuration with $B_{\text{source}}=822$ G and discharge voltage=160 V. The training data only covers discharge voltages up to 150 V. The machine was also in an odd discharge state, so it's no surprise that the predicted uncertainty bounds are very large (even greater than the test set RMSE value) and that accuracy suffers.

as with any model – remains difficult.

5.5.3 Comparison with Thomson scattering

The z-axis interpolation for dataruns in the training and test sets can be evaluated using the Thomson scattering (TS) diagnostic. The TS measurement is taken 8 ms into the discharge for DR1 12 ms into the discharge for DR2, but in this study the measured and predicted I_{sat} are instead averaged over 10 to 20 ms. The Thomson scattering measurement is compared in with DR1 and DR2 in figs. 5.16a and 5.16b, respectively. The linear slope fits do not take model error into account. In DR1, I_{sat} predictions disagree with the I_{sat} derived from TS. Measurements from probes, when nearby the TS beam, can also have very different values from the TS-derived measurement. The TS density measurement may suffer from misalignment, and has not been calibrated since January 2022, roughly a year (DR1) or two (DR2) before these data were taken. The density measurement is photon counting and requires absolute calibration. This disagreement likely comes from this error in density because fitting the temperature is robust to absolute calibration errors. In addition, the I_{sat} is time-varying; the average may differ substantially from single points in time earlier in the discharge.

All these issues considered, the model predictions roughly agree with TS on average in DR1, which is encouraging because the TS beam at port 32 (671 cm) is substantially further from the closest probe at port 27 (831). We should expect rough agreement or a slight underestimate on average based on the skewed test-set z-score seen in fig. 5.10). DR2 has a probe past the TS beam at port 33 (639 cm), but the I_{sat} measurement rarely agrees with Thomson. Because of this density error and measurement time discrepancy, we conclude that the TS diagnostic may not be a good way to verify the predictions of the model. Note that, when calibrated, TS agrees with I_{sat} measurements quite closely as seen in the LAPD Thomson scattering paper [62].






Figure 5.15: Comparison of original DR2 profiles with the profiles from the optimization dataset (DROpt) for the same machine configuration. The I_{sat} values in the DROpt dataset are not calibrated in this plot, indicating significant variation in probe calibration in this DROpt dataset.

5.6 Effect of I_{sat} calibration

The Langmuir probes did not seem to be behaving correctly when the optimization validation data were taken. The probes showed an *increasing* I_{sat} profile when moving further from the cathode in the lowest gas puff condition, which is at odds with previous measurements and intuition. An example of this discrepancy can be seen in Fig. 5.15, where a run from the original testing set (specifically DR2 run 10) is duplicated. The probes for the validation run can be corrected for by either assuming the 5 ms gas puff run has a flat axial profile or by normalizing the probes to the DR2 run 10 axial profile. Calibrating the probes using the DR2 run 10 reference was the best way to go because it corrects for both probe discrepancies as well as changes in the condition (or emissivity) of the main cathode.



(a) Thomson scattering (TS) 8 ms into the discharge compared to the model predictions (10 to 20 ms averaged). Broadly speaking, the TS measurement roughly agrees with the model estimate on average.



(b) Thomson scattering (TS) 12 ms into the discharge compared to the model predictions (10 to 20 ms averaged) and I_{sat} measurements one port away. The TS underestimates I_{sat} in general.

Figure 5.16

5.7 Inferring trends

A systematic study of the impact of discharge voltage on I_{sat} profiles has not been conducted using conventional techniques. Collecting both z- and x-axis profiles over a wide range of discharge voltages would take a considerable amount of time, mostly from the requirement to dismount and reattach the probes and probe drives along the length of the LAPD. This study has now been performed using the learned model, circumventing these time-consuming challenges. Model input parameters were chosen to be common, reasonable values: 1 kG flat field, 80 V gas puff, 38 ms gas puff duration, run set=DR2, and top gas puff off. The 38 ms puff is used in these predictions because it is the most common gas puff duration in the training set – the model is biased in favor of this gas puff setting. The results of changing the discharge voltage only can be seen in fig 5.17. Notably, the I_{sat} increases across both axes. Steeper axial gradients are seen with lower discharge voltages, but peaked x-profiles are seen at higher discharge voltages. The area closer to the source region (+z direction) appears to have a steep drop but flatter profiles down the length of the machine.

Unfortunately the discharge current was not included as an output in the training set. Otherwise the effect of changes in discharge power, rather than simply voltage, could be computed. The discharge current – and thus discharge power – is set by cathode condition, cathode heater settings, and the downstream machine configuration, and thus cannot be set to a desired value easily before the discharge. Discharge voltage, however, can remain fixed.

Of particular interest for some LAPD users is achieving the most uniform axial profile possible. We explore this problem in the context of mirrors. The gas puff duration is known to be a large actuator for controlling density and temperature and so is explored as a way of shaping the axial profile. We predict discharges with a flat 1 kG field with the probe in the center. The discharge voltage was set at 110 (a reasonable, middle value) with the run set flag=DR2 and top gas puffing=off. The inferred effect of gas puff duration on the axial gradient and axial gradient scale length can be seen in Fig. 5.18. Care was taken to handle the aleatoric (independent) uncertainty separately from the axially-dependent epistemic uncertainty. As seen in the figure, the mean axial gradient




Figure 5.17: The z profile at x=0 cm and x profile at z=1140 cm for different discharge voltages. The I_{sat} decreases with increasing voltage, and the error (filled regions) stays roughly the same, but in general increase slightly towards the cathode and at higher discharge voltages.




Figure 5.18: ML prediction: mean axial gradients decrease with decreased gas puff duration. Five durations are plotted between 5 and 38 ms (which are the bounds of the training data), evenly spaced. The gradient scale length also increases, indicating that the gradient change was not solely caused by a decrease in the bulk I_{sat} .

decreases when the gas puff duration is shortened. The gradient scale length also increases, so the mean gradient is not decreasing simply because the bulk I_{sat} is decreasing. This result suggests that the gas puff duration may be a useful actuator to consider when planning future experiments.

These applicability of these results are somewhat muted because the gas puff duration was not chosen randomly in the training discharges. Given this lack of data diversity, the accuracy and applicability of this study must be interpreted cautiously. When a model is trained on *all* data available (using the cross-validated test set MSE as a guide for error), which includes the 20 ms gas puff case, the predicted gradient scale length decreases uniformly across the duration scan by 1 meter. The fact that the trend remains intact when an additional, randomized intermediate gas puff case is added gives confidence in the predictions of the model despite the lack of data diversity.

5.8 Optimizing profiles

One particular issue seen in LAPD plasmas is sharp axial density and temperature gradients. Some experiments require flat gradients, such as Alfvén wave propagation studies. We explore optimizing the axial I_{sat} variation as an approximation to this problem. In addition, in this case the optimization problem is used as a tool to evaluate the quality of the learned model. This optimization is a very demanding task because the inferred trends along all model inputs must be accurate simultaneously. Constraints on this optimization further increase the difficulty of the problem. Success in optimization provides strong evidence that the model has inferred relevant trends in predicting I_{sat} . We quantify the uniformity of the axial profile by taking the standard deviation of I_{sat} of 11 points along the z-axis ($x, y = 0$). The required LAPD state for attaining the most uniform axial profile can be found by finding the minimum of this standard deviation with respect to the LAPD control parameters and flags:

$$\text{Inputs} = \arg \min_{\text{Inputs} \neq z} \text{sd}(I_{\text{sat}}|_{x=0}) \quad (5.3)$$

The largest axial variation can likewise be found by finding the maximum. The model inputs used for this optimization can be found in Table 5.2.

For this optimization we use an ensemble of five β -NLL-loss models with weight decay $\lambda = 0$. The $\lambda = 0$ model is used because it appears to give the most useful uncertainty estimate (seen in Fig. 5.12). The optimal machine actuator states are found by feeding a grid of inputs into the neural network. This variance estimate is not well-calibrated: the error of the predictions on the test set falls far outside the predicted uncertainty. However, this uncertainty can be used in a relative way: when the model is predicting far outside its training range, the predicted variance is much larger. The ranges of inputs into this model are seen in Table 5.2. These inputs yield 127,234,800 different machine states (times five models) which takes 151 seconds to process on an RTX 3090 (≈ 4.2 million forward passes per second) when implemented in a naive way. The number of forward passes can be reduced by a factor of 51 if the x value is set to 0 cm. Note that gradient-based methods can be used for search because the network is differentiable everywhere but this network

Table 5.2: Machine inputs and actuators for model inference

| Input or actuator | Range | Step | Count |
|-------------------|-------------------|---------|-------|
| Source field | 500 G to 2000 G | 250 G | 7 |
| Mirror field | 250 G to 1500 G | 250 G | 6 |
| Midplane field | 250 G to 1500 G | 250 G | 6 |
| Gas puff voltage | 70 V to 90 V | 5 V | 5 |
| Discharge voltage | 70 V to 150 V | 10 V | 9 |
| Gas puff duration | 5 ms to 38 ms | 8.25 ms | 5 |
| Probe x positions | -50 cm to 50 cm | 2 cm | 51 |
| Probe y positions | 0 cm | – | – |
| Probe z positions | 640 cm to 1140 cm | 50 cm | 11 |
| Probe angle | 0 rad | – | – |
| Run set flag | off and on | 1 | 2 |
| Top gas puff flag | off and on | 1 | 2 |

and parameter space is sufficiently small that a comprehensive search is computationally tractable.

Like any optimization method, the results may be pathologically optimal. In this scenario, the unconstrained minimal axial variation is found when the I_{sat} is only around 1 mA/mm², which is quite small and corresponds to $1\text{-}2 \times 10^{12}$ cm⁻³ depending on Te. The inputs corresponding to this optimum are given in the second column of Table 5.3. This density range is below what is required or useful for many studies in the LAPD.

Since many physics studies require higher densities, we constrain the mean axial I_{sat} value to be greater than 7.5 mA/mm² (roughly $0.5\text{-}2 \times 10^{13}$ cm⁻³). The “run set flag” is set to “on” for cases to be validated (bolded in Table 5.3) because we wish to keep the turbopumps on to represent typical LAPD operating conditions. In addition the “top gas puff flag” was set to ‘off’ to minimize the complexity of operating the fueling system on followup dataruns and experiments. Turning the top gas puff valve on is predicted to decrease the average I_{sat} by -2 mA/mm² for strongly varying profiles, but otherwise the shapes are very similar. The inputs corresponding to the maximum and minimum axial variation under these constraints can be seen in columns 3 and 4 of Table 5.3. For contrast we also consider what machine settings would lead to the greatest axial variation. The results of both of these optimizations can be seen in Fig. 5.19. The optimizations yield profiles that have the largest I_{sat} values closest to the cathode, which is expected.

The prediction for an intermediate axial variation case is also seen in Fig. 5.19. The intermediate case was chosen as somewhere around half way between the strongest and weakest case with a round index number (15000, in this case). The parameters for intermediate case are also enumerated in Table 5.3.

The predicted configurations with the run set flag on and top gas puff flag off (bolded in Table 5.3) were then applied on the LAPD. The data collected, compared with the predictions can be seen in Fig. 5.19.

For the optimized axial profiles, the absolute value of the I_{sat} predictions compared to measurement do not agree. All of the predicted profiles have overlapping predictions (within the predicted




Figure 5.19: Axial profiles, predicted and measured, for the optimized weakest (blue), intermediate (green), and strongest (orange) cases. a. The shaded region covers the mean prediction \pm one standard deviation, and the dashed lines are \pm the median cross-validation RMSE values. b. The measured I_{sat} values are calibrated to DR2 run 10 (solid lines), or using triple probe Te measurements on the probe and linearly extrapolating the interferometer measurements (dotted lines). The absolute values disagree between the predicted and measured values, but axial trends are consistent with the optimization.

Table 5.3: Machine inputs and actuators for optimized axial profiles

| Input or actuator | Weakest | Weakest | Strongest | Intermediate |
|---|-------------------------------|------------------------|------------------------|------------------------|
| I_{sat} constraint (mA/mm ²) | $I_{\text{sat}} = \text{any}$ | $I_{\text{sat}} > 7.5$ | $I_{\text{sat}} > 7.5$ | $I_{\text{sat}} > 7.5$ |
| Source field | 750 G | 1000 G | 500 G | 2000 G |
| Mirror field | 1000 G | 750 G | 500 G | 1250 G |
| Midplane field | 250 G | 250 G | 1500 G | 750 G |
| Gas puff voltage | 70 V | 75 V | 90 V | 90 V |
| Discharge voltage | 130 V | 150 V | 150 V | 120 V |
| Gas puff duration | 5 ms | 5 ms | 38 ms | 38 ms |
| Run set flag | on | on | on | on |
| Top gas puff flag | on | off | off | off |

error) at the region furthest from the cathode, but the measured values do not show that behavior. Although the mean I_{sat} value was constrained to be greater than 7.5 mA/mm², the measured mean was 2 mA/mm² for the weakest case.

The important result is that the optimized LAPD settings, when implemented on the machine, do yield profiles with strong, intermediate, and weak axial variation. Although the minimum- I_{sat} constraint was violated for the case of weakest axial variation case, this optimization would nonetheless be very useful for creating axial profiles of the desired shape.

There are three contributing factors to the mismatch of the ML-predicted values and the real measured values. First, the condition of the machine, such as the cathode emissivity or temperature or the downstream neutral pressure, are unquantified and cannot be compensated for in data preprocessing or in the model itself. Second, the calibration of the Langmuir probes could differ substantially between runs. The probes in the training data run sets (DR1 and DR2) were well-calibrated to each other within the run set, but were not absolutely calibrated. The probes used for verifying the optimization were not calibrated. A rough calibration was performed by linearly extrapolating interferometer measurements and using triple probes (dotted lines on the right panel

in Fig. 5.19). A configuration identical to DR2 run 10 was also measured to simultaneously correct cathode condition and probe calibration (solid lines on the right panel in Fig. 5.19). Langmuir probe calibration is discussed further in Appendix 5.6. Third, the original dataset may not have sufficient diversity to make accurate predictions on such a constrained optimization problem.

If this optimization were performed using the dataset instead of the model, the constrained search would encompass just 10084 shots out of the 131550 shots total in the training dataset, or around 7.7%. Including the on-axis constraint reduces the number of shots down to 303 (270 in the training set), or 0.23% of all shots in the dataset. We conclude that this optimization of an arbitrary objective function, as done here, would be intractable using traditional, non-machine learning techniques because orders of magnitude more dataruns would need to be collected.

Optimization requires correctly learning the trends of all inputs and how they interact. In addition, as seen from the shot statistics, the model was trained on very few shots in the constrained input and output space. These two factors – the need for the model to learn all trends and the constrained search space – combine to make an incredibly difficult task that functions as a benchmark for the model. These factors considered, it is not surprising that the model incorrectly predicts the absolute value. The uncertainty predicted by the model, though not well-calibrated, was nonetheless very large compared to the median test set RMSE. The model did predict the trends correctly, however; the optimized, measured profiles were strong, intermediate, or weak.

We did not check to see if the predicted optima were actually true optima: an approximation of the local derivative using a finite-difference technique would require much more run time on the LAPD than was available.

5.9 Discussion

5.9.1 Key achievements

To my knowledge, this work is one of the first instances in which machine learning has been used to infer specific trends and optimize profiles in magnetized plasmas. Three examples of trend inference were shown in this study: influence of magnetic geometry on plasma width, changes in the axial and radial profiles with changing discharge voltage, and the relationship of gas puff duration with axial gradient scale length. In addition, the axial profile was optimized by minimizing (or maximizing) the axial standard deviation. There is no other way of simultaneously uncovering many trends or finding optima without using an ML model trained on a diverse dataset. Traditionally, such studies would require extensive scans over grids to map the parameter space, but here it was accomplished with a relatively small amount of data.

The trends inferred in this work, such as changing discharge voltages, gas puff durations, or mirror fields, would traditionally require a grid scan (varying one parameter at a time) in LAPD settings space. Here instead we are able to extract any trend covered by the training set with only a minimal amount of machine configurations sampled. Both data collection runs lacked absolute I_{sat} calibration and had potential differences in cathode condition. Despite these issues the model learned trends that were exploited via optimization.

In addition, this work demonstrates uncertainty quantification broken down into epistemic and aleatoric components by using ensembles and a negative-log likelihood loss. This uncertainty estimate is useful in gauging relative certainty between different predictions of the model which increases confidence in the predictions of the model. In general, the total uncertainty predicted by the model increases when predictions are made outside the bounds of the training data.

Fundamentally, this model can predict I_{sat} with uncertainty at any point in space covered by the training data. No other model currently exists that can perform this prediction. Traditionally, this capability would be possible only with a detailed theoretical study.

5.9.2 Current limitations

This study would be dramatically improved by collecting more, diverse data. Only 44 of the 67 dataruns in this dataset had randomly sampled LAPD machine settings which is very small compared to the over 60,000 possible combinations. In addition, there are many other settings or parameters that were not changed in this study, such as gas puff timings, gas puff valve asymmetries, wall/limiter biasing, cathode heater settings, operation of the north cathode, and so on. The bounds of the inputs were also conservative; all settings in this study could be pushed higher or lower with a small amount of risk to LAPD operations. In addition, the placement of the probes could be further varied and placed outside the mirror cell, which would provide a more complete picture of LAPD plasmas, particularly axial effects.

Probe calibrations differed between the two training run sets (DR1 and DR2) – and a flag was introduced to distinguish between them – but despite this deficiency combining the two run sets was shown to be advantageous for model performance. The condition of the cathode (e.g., electron emissivity and uniformity) also has a large impact on the measured I_{sat} . The improved model performance with the flag suggests that inconsistencies between dataruns could be compensated for using latent variables if a generative modeling approach is to be taken. At the very least, this model provides a way to benchmark these differences in machine state.

Concerning the model, hyperparameter tuning could be performed. In this study a few extra percent in MSE performance is not meaningful considering the limited dataset. Instead, we focused on the trends and insights that can be extracted from this model rather than simple predictive accuracy. There may also be regimes in hyperparameter space where the uncertainty is better calibrated (perhaps using early stopping). Uncertainty estimation is important, even if the absolute uncertainty is not well-calibrated because it can provide a useful relative estimate as shown in this study.

Trends such as the dependence of axial gradient on the gas puff duration (fig. 5.18) or the effect discharge voltage on x-z profiles (fig. 5.17), although intuitive, remain unverified. Verification of

these trends would increase confidence of model predictions when setting LAPD parameters in future experiments.

5.9.3 Future directions

The neural network architecture used here can readily scale to additional inputs and outputs; including time-series signals is the obvious next step. Integration of multiple diagnostics – perhaps starting with individual models before combining them – could enable inference of plasma parameters throughout the device volume. For example, combining triple probe electron temperature measurements with existing I_{sat} data would allow density predictions anywhere in the plasma. This capability could enable in-situ diagnostic cross-calibration (e.g., the Thomson scattering density measurement) and prediction of higher-order distribution moments like particle flux. The model could be further enhanced by incorporating physics constraints such as boundary conditions (e.g., zero I_{sat} at the machine wall) or symmetries.

The problem presented here – learning time-averaged I_{sat} trends – is fairly simple and required a relatively simple model. As demonstrated in this work, ML provides a way to explore regions of parameter space quickly and efficiently. Most physics studies on plasma devices (and fusion devices) are dedicated to a single particular problem, use grid scans, and are not useful to other experiments or campaigns. This work shows a way of using data and trends uncovered from other experimental studies. This work also demonstrates that random exploration can be a useful tool: the increased diversity of the aggregated data will generally benefit an ML model whether the experimenter discovers something new or not.

5.10 Conclusion

We demonstrate the first randomized experiments in a magnetized plasma experiment to train a neural network model. This learned model was then used to infer trends when changing field configuration, discharge voltage, or gas puff duration. This model was also used to optimize axial

variation of I_{sat} as measured by the standard deviation, which was validated in later experiments despite poor absolute error.

We strongly advocate that all ML-based analyses in plasma and fusion research should be validated and used to gain insight by inferring trends, as demonstrated here. This validation step is crucial for ensuring that ML models capture physically meaningful relationships and the insights provided may provide direction for future research. We hope this is the first step towards automating plasma science.

5.11 The open dataset and repository

All the code to perform the ML portion of this study is available at <https://doi.org/10.5281/zenodo.15007853>. The training datasets are also available in that repository in the datasets directory. Additional data are available upon request. The repository also contains additional training details and the notebooks for generating the plots in this document. The plots used in this study were made in jupyter notebooks, which are also uploaded. The final training code can be found in `train_dense_beta_NLL.py`. Trained models are found in the `code/training_runs` directory. The history of many training runs can be found on Weights and Biases: <https://wandb.ai/phil/profile-predict> and the accompanying notes on these trained models are found in the associated pdf on github. The code and dataset are licensed under Creative Commons Attribution Share Alike 4.0 International License.

CHAPTER 6

Energy-based models for diagnostic reconstruction and distribution modeling

Energy-based models (EBMs) provide a powerful and flexible way of learning relationships in data by constructing an energy surface. We extend EBMs to laboratory plasma physics, a domain characterized by highly nonlinear phenomena studied using incomplete diagnostic information. These diagnostics can be unreliable or difficult to analyze. In addition, the possible configuration space of a plasma device is sufficiently large that it cannot be efficiently searched using conventional analysis techniques. EBMs provide a way to address these issues. At the Large Plasma Device (LAPD), a CNN- and attention-based EBM is trained on a set of randomly generated machine conditions and their corresponding diagnostic time series. We demonstrate diagnostic reconstruction using this EBM and also show that including additional diagnostics improves reconstruction error and generation quality. Symmetries in the data can be found by directly evaluating the energy surface, potentially leading to a new line of inquiry using learned models. In addition, this multimodal EBM is able to unconditionally reproduce all distributional modes, suggesting future potential in anomaly detection on the LAPD. Fundamentally, this work demonstrates the flexibility and efficacy of EBM-based generative modeling of laboratory plasma data, and showcases practical use of EBMs in the physical sciences.

6.1 Goals and prior work

6.1.1 Goals

We seek to use machine learning – particularly generative models – to alleviate some of the challenges facing fusion-related plasma science and to accelerate the advent of fusion power. In particular, we want to be able to reconstruct missing diagnostic signals from any other set of signals. In addition, the reconstruction should be a probability distribution, not just a single estimate. This distribution should also be combinable with other sources of information such as simulations or constraints placed on it via prior knowledge.

6.1.2 Generative modeling of plasmas

The use of generative models in plasma physics is not without precedent, but remains relatively uncommon. Variational autoencoders [89] have been used to generate new, realistic output from stellarator transport simulations for inferring trends and uncovering physics [98] and to discover relationships between inputs and outputs [160]. These VAEs have also been used in the COMPASS to identify a rare instability characterized by fluctuations in magnetic probes [169] by pretraining on unlabeled data and combining the model with a classifier over a smaller, labeled dataset.

On the Joint European Torus experiment, a generative topographic map was used to create a 2d representation of a 7d input space (information from 1d profiles) [117]. This mapping clearly shows a disruptive-nondisruptive boundary and the relative stability of locations in this 2d representation can be evaluated by visualizing cluster distances. Discharge trajectories can be visualized in this reduced 2d space. Likewise on the HBT-EP tokamak, a VAE was trained on coil currents, equilibrium properties, and MHD information to learn a 2d latent representation [163]. This model was run in real time to identify threshold crossing events which then triggered preprogrammed countermeasures.

Generative-adversarial networks (GANs) [64] have also been used to generate synthetic training

data (time series of plasma current) for use in training a disruption predictor [31]. GANs have also been used to generate posterior distributions of surface temperature and emissivity given the measured radiance for the thermal surface-pointed cameras on the WEST tokamak [83].

Diagnostic reconstruction is useful in the event a diagnostic goes offline, but it can also be used to supplement existing diagnostics. One such example is the upsampling of the Thomson scattering signal on DIII-D based on information from other diagnostics using a neural network-based solution [80]. In addition, Bayesian and ML work in fusion has been recently reviewed [118].

Outside of magnetized plasmas, random experiments were performed in an inductively coupled plasma (ICP) similar to the process used to collect for this work. These data were used to train a VAE as a surrogate collisional radiative model [30] to construct an interpolatable latent representation.

Concerning EBMs, they have yet to be applied to plasma physics problems. One notable application in the physical sciences has been in the high-energy physics community. EBMs were used for modeling event patterns in the Large Hadron Collider (LHC) for anomaly detection and to augment a classifier [26] with success.

6.2 Introduction to energy-based models (EBMs)

Energy-based models interpret a probability distribution through the lens of the Boltzmann distribution [73, 5, 97]. In the EBM formulation, the unnormalized probability density is parameterized by an energy function, that is $\tilde{p}(\mathbf{x}) = \exp(-E_{\theta}(\mathbf{x}))$, where θ are the parameters of this energy function. In this work and the works cited, this energy function is parameterized by a neural network. EBMs have been historically difficult to train, but recent work has demonstrated high-quality sampling using MCMC techniques in high-dimensional spaces [40, 41, 39, 37, 38, 112, 113, 33, 59]. These MCMC techniques are fundamentally based in contrastive divergence [69, 135, 150]. The nature of MCMC-based sampling of EBMs, detailing the convergence of and expansion and contraction of learned models (a paradigm particularly helpful for training EBMs in this work) has been studied [112, 113]. Training via contrastive divergence can be improved by implementing a term typically

left out which approximates the KL-divergence [39].

EBMs can attain GAN-like performance while generating all modes of a distribution [41]. These models have also been used for text generation [33] and model-based planning [40]. In addition, EBMs can be composed by combining the energy functions in various ways which has been demonstrated in image generation tasks [37, 38].

Overview of EBMs, how they are trained, and their place among generative models [21]. Another helpful review of generative modeling that helped in the selection by detailing advantages and disadvantages of various methods [17]. EBMs in particular have some desirable properties. The implicit generation scheme does not require balancing an explicit generator with a discriminator or encoder. In addition, the energy-based representation is amenable to modification of an auxiliary function, allowing the generated samples to be easily steered, enabling the composition of multiple energy functions. Lastly, EBMs do not easily suffer from mode collapse or otherwise ignore parts of the distribution. For these reasons, we use EBMs for building a distribution of our LAPD data.

6.3 Data preparation

Around 130 thousand shots were taken on the Large Plasma Device for 67 different dataruns, varying machine conditions and probe configurations as detailed in chapter 4. The train-test split is identical to what was performed in earlier chapters. Namely, four dataruns each from DR1 and DR2 were held out as representative configurations. In order to reduce computational requirements, the time series data was downsampled. The original sampling rates for the machine state information and auxiliary diagnostics were 25 kHz. The sampling rate for the I_{sat} probe was 6.25 MHz. All time-series data were downsampled to a common sampling rate of 2.5 kHz. This downsample leads to the time series having a length of 76 points long, and the MSI were truncated to be identical in duration and start time to the I_{sat} time series.

The dataset includes nine time series in total: discharge current, discharge voltage, 5 diode signals (the last of which having a He-II filter), interferometer (line-integrated density), and I_{sat}

from a Langmuir probe. Additionally, several one-dimensional inputs and flags were included: probe position (x, y, z), magnetic field (source, mirror, and midplane), gas puff duration and voltage, total gas pressure, first 4 amu RGA readings, run set flag, and top gas puff flag. In total, the input length is 699 features when concatenated.

When fed into the model, each input is mean-subtracted and scaled independently by the peak-to-peak value. Scaling each value was found to work much better than scaling a whole time series by the same value. Often, the beginning and end of the time series would have very little variance (such as in discharge current always starting at 1 kA). Learning distributions of very small width appears to be difficult for EBMs.

6.4 Training the EBM

The following loss function is used to train the EBM:

$$\mathcal{L} = \mathcal{L}_{\text{CD}} + \mathcal{L}_{\text{KL}} + \alpha \mathcal{L}_{\text{reg}} \quad (6.1)$$

where \mathcal{L}_{CD} is the contrastive divergence (CD) loss, \mathcal{L}_{KL} is the KL-divergence loss, and \mathcal{L}_{reg} is the energy regularization loss, listed in order of importance. The contrastive divergence loss is defined as:

$$\mathcal{L}_{\text{CD}} = \frac{1}{M} \sum_i E_{\theta}(\tilde{x}_i^+) - E_{\theta}(\tilde{x}_i^L) \quad (6.2)$$

This loss calculates the difference in energies between sample from the energy surface \tilde{x}_i^L via Langevin dynamics – the “negative” samples – and the samples from the data distribution \tilde{x}_i^+ , i.e., the training data. This loss places the energy surface in tension, with the surface being pulled down by the data and pushed up on the negative samples. Both terms must be included: absence of negative samples would lead to a vacuous solution such as a flat energy surface. The KL-divergence loss is:

$$\mathcal{L}_{\text{KL}} = \frac{1}{M} \sum_i E_{\Omega(\theta)}(E_{\theta}(\hat{x}_i^K) - \text{NN}(X, \hat{x}_i^K)) \quad (6.3)$$

This loss, suggested in [39] includes portions of the CD loss typically left out. Specifically, this loss minimizes the sampler energy by propagating gradients through the MCMC steps themselves instead of the final samples. In addition, the nearest-neighbor (NN) samples are used to estimate the energy of the negative samples so that the entropy of the distribution is maximized. This loss, though not necessary, significantly improves training stability. The regularization loss

$$\mathcal{L}_{\text{reg}} = \frac{1}{M} \sum_i E_{\theta}(\tilde{x}_i^+)^2 + E_{\theta}(\tilde{x}_i^L)^2 \quad (6.4)$$

keeps the energy values centered near 0. The absolute value of the energy does not matter – only the gradients and scale – but this is included to keep the results easily interpretable and to avoid extreme energy values as to not run into floating-point representation boundaries. In this work, the scale of this regularization is very small with $\alpha = 1 \times 10^{-6}$. This loss also functions as a great diagnostic for when the sampler fails: the regularization loss will reach $1/\alpha$. The EBM training process is presented in algorithm 1.

6.4.1 The sampler

The sampler is fundamental to how the model is trained and how samples are generated; the configuration of the sampler is one of the most important considerations when building the EBM. The samples are initialized from random noise between -1 and 1 . We use Langevin dynamics to move the samples as formulated in Nijkamp et al. [112], as seen in algorithm 2. A variety of step sizes and number of steps were used. The number of sampler steps per batch was critical for training stability: 30 appears to be close to the minimum number of steps for stable training, similar to what has been observed in other studies [26].

The step size is typically chosen to match the standard deviation of the smallest feature, but that was found to be too large for this use case. Given that our data are highly multimodal, some modes of the input distribution had a deviation of 0 (such as the flags). For this reason, a smaller step size of $\epsilon = 0.01$ was used.

The model was trained for 27 hours over 172 epochs. The batch size was 128 and was trained

Algorithm 1 EBM training algorithm

Require: Training samples x_i^+ , training data distribution p_D , energy function E_θ , replay buffer \mathcal{B} , step size ε , MCMC steps L , KL MCMC steps K , energy regularization strength α , stop gradient operator $\Omega(\cdot)$, replay fraction $f_{\mathcal{B}}$, batch size M

$\mathcal{B} \leftarrow \mathcal{U}(-1, 1)$

▷ Fill buffer from uniform distribution

while not converged **do**

$x_i^+ \sim p_D$

$\tilde{x}_i^0 \sim \mathcal{B}$ sample $Mf_{\mathcal{B}}$ negative examples, $\mathcal{U}(-1, 1)$ otherwise

$X \sim \mathcal{B}$ nearest-neighbor samples such that $X \cap \tilde{x}_i^0 = \emptyset$

for sample step $\ell = 1$ to L **do**

▷ Run Langevin dynamics

$\tilde{x}_i^\ell \leftarrow \tilde{x}_i^{\ell-1} - \frac{\varepsilon^2}{2} \nabla_x E_\theta(\tilde{x}_i^{\ell-1}) + \varepsilon \mathcal{N}(0, 1)$

end for

$\tilde{x}_i^L = \Omega(\tilde{x}_i^L)$

$\hat{x}_i^0 = \tilde{x}_i^\ell$ where $\ell = L - K$

▷ Run Langevin dynamics for KL loss

for KL sample step $k = 1$ to K **do**

$\hat{x}_i^k \leftarrow \hat{x}_i^{k-1} - \frac{\varepsilon^2}{2} \nabla_x E_\theta(\hat{x}_i^{k-1}) + \varepsilon \mathcal{N}(0, 1)$

end for

$\mathcal{L}_{\text{CD}} = \frac{1}{M} \sum_i E_\theta(\tilde{x}_i^+) - E_\theta(\tilde{x}_i^L)$

$\mathcal{L}_{\text{KL}} = \frac{1}{M} \sum_i E_{\Omega(\theta)}(E_\theta(\hat{x}_i^K) - \text{NN}(X, \hat{x}_i^K))$

▷ Has gradients through MCMC

$\mathcal{L}_{\text{reg}} = \frac{1}{M} \sum_i E_\theta(\tilde{x}_i^+)^2 + E_\theta(\tilde{x}_i^L)^2$

$\mathcal{L} = \mathcal{L}_{\text{CD}} + \mathcal{L}_{\text{KL}} + \alpha \mathcal{L}_{\text{reg}}$

Apply $\nabla_\theta \mathcal{L}$ to θ via the Adam optimizer

$\mathcal{B} \leftarrow \mathcal{B} \cup \mathcal{U}(-1, 1)$ and remove samples to maintain buffer size

end while

Algorithm 2 EBM sampling

Require: Energy function E_θ , auxiliary energy function F , step size ε , MCMC sampling steps L

$\tilde{x}_i^0 \sim \mathcal{U}(-1, 1)$ ▷ Initialize on uniform distribution
for sample step $\ell = 1$ to L **do** ▷ Run Langevin dynamics
 $\tilde{x}_i^\ell \leftarrow \tilde{x}_i^{\ell-1} - \frac{\varepsilon^2}{2} \nabla_x \left(E_\theta(\tilde{x}_i^{\ell-1}) + F(\tilde{x}_i^{\ell-1}) \right) + \varepsilon \mathcal{N}(0, 1)$
end for

Return \tilde{x}_i^L

with the AdamW optimizer with a learning rate of 10^{-4} , decaying to 10^{-5} after 20 epochs. Training using stochastic gradient descent required very large learning rates (10^{-2}) and could not generate realistic samples even with a learning rate decay schedule. The first 100 epochs were trained in a single run; the next 72 were resumed from the final checkpoint of the previous run, leading to reinitialization of the replay buffer. Given the low learning rate, this did not have a significant impact on model training. This model, taken from the 39th epoch of the second run, was used for inference; after which, training diverged.

The training curves of loss, relative energy, and energy gradient can be seen in fig. 6.1. The loss curve is very hard to interpret given all the components of the loss function. Maintaining a relative energy near 0 between negative and positive samples is crucial for model performance. A value near 0 indicates that the generated samples are close to the data distribution. When the sampler diverges, samples are no longer accurately representing the data and thus training fails. The energy function learns the scale of the gradients required to produce good samples. When training is proceeding normally, the ratio of the gradient step to the noise scale (step size) ε is near 1.

6.4.2 Replay buffer configuration

A replay buffer is used to provide a warm start for the sampler chains. Every batch iteration, 5% of the samples from the buffer are replaced with random noise (a replay fraction of 0.95). This

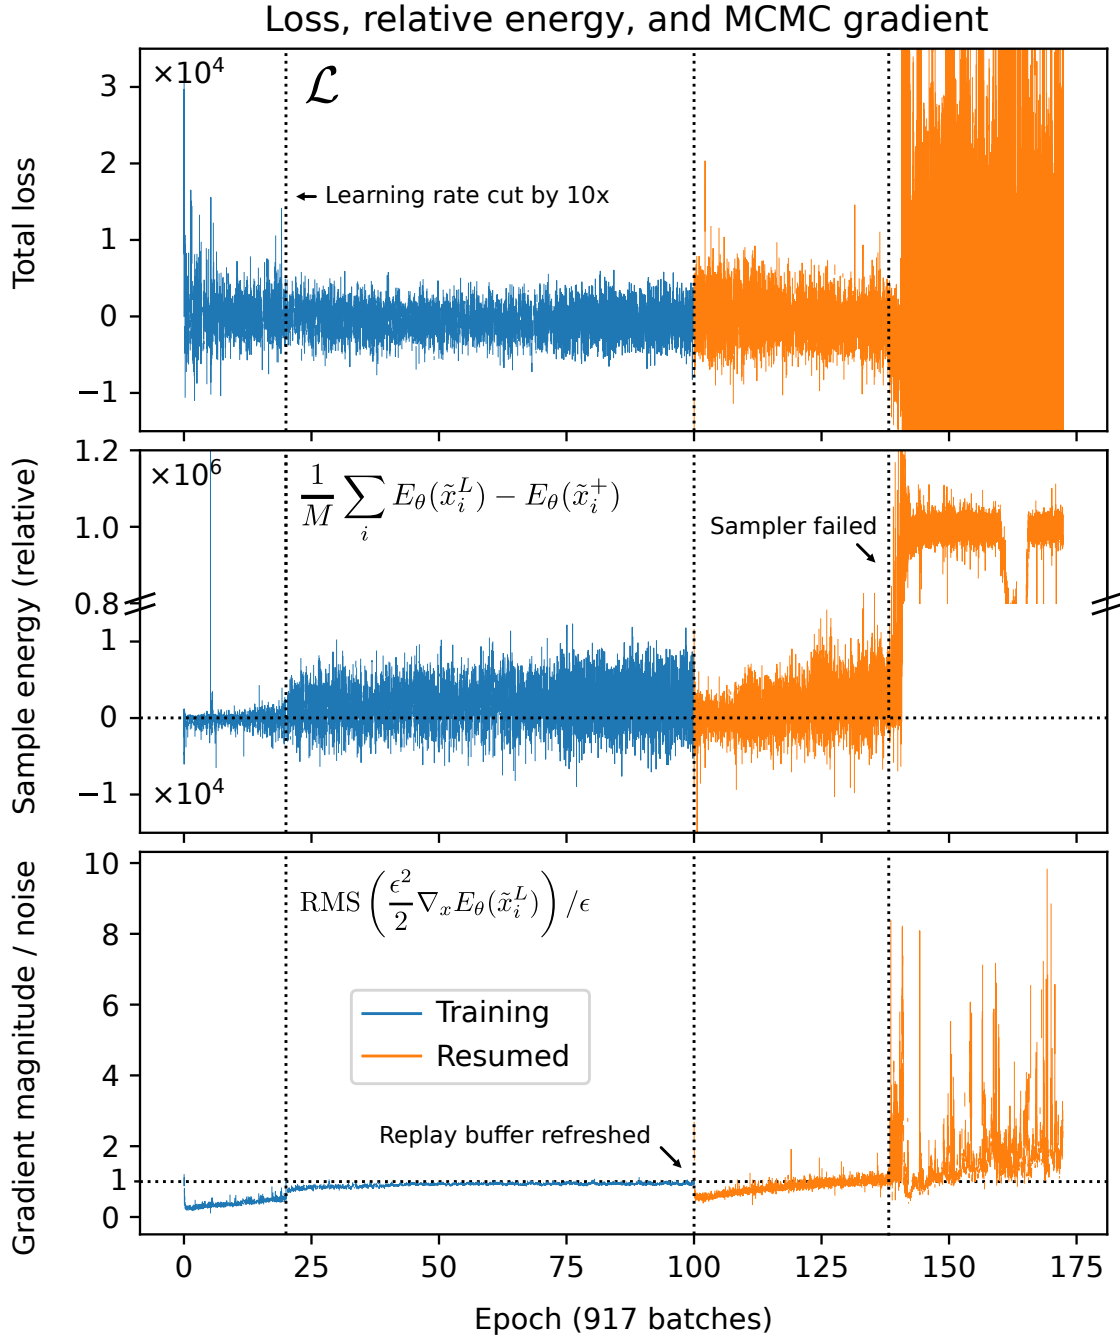


Figure 6.1: Training curves of the model. Top: the total loss, middle: the relative energies of the negative samples and the training examples (positive samples), bottom: the gradient of the energy function normalized to the noise scale (step size ϵ).




Figure 6.2: Distribution of batches in replay buffer. When training is starting (epoch 0, left), the number of batches each sample experiences is low and somewhat uniform. After 10 training epochs (right), the number of batches experienced by a sample converged to an exponential distribution.

replay fraction leads to a mean of $1/0.05=20$ batch iterations for each chain, with half the chains experiencing $\ln(2)/0.05 \approx 14$ batches. The replay buffer requires 4 to 5 epochs to converge to an exponential distribution in the number of steps experienced by each chain in the buffer. This diversity of chain lengths likely encourages quick convergence of the chain but good long-term samples (on average, each chain experience 600 MCMC updates/steps). The distribution of the batches among the replay buffer can be seen in fig. 6.2. One downside of this long buffer is that if a sample ends up in a region far outside the domain (-1 to 1) it can cause exploding gradients; samples have a $\approx 50\%$ chance of lasting 194 batches if one goes awry. Implementing a reject step for these prodigal samples may improve training stability.

6.4.3 Sampling behavior

In general, if the model is not learning, increase the number of sample steps, decreasing the step size, and decreasing the learning rate are beneficial steps to take. A lower learning rate decreases the change per batch in the energy surface, so stale sample chains from the replay buffer may find themselves in more familiar territory than if the energy function is rapidly changing. More sample

steps relaxes the requirement for the samples to rapidly find realistic (lower energy) locations, perhaps leading to shallower gradients and greater stability. Much work needs to be done to really understand the dynamics of the sampler in the training process. The samples from the replay buffer sampling may also benefit from a non-memoryless distribution so that the chains are guaranteed a certain number of sample steps over their lifetime, with a hard limit so that chains do not persist for too long.

6.5 Architecture

The model is intrinsically multi-modal: time-series data from diagnostics is mixed in with machine settings, state, and probe position. In this model we preprocess separately each time series and the LAPD configuration. Convolutional NNs were used for the time series input, and transformer-like multi-head attention blocks were used for the settings, state, and probe position. The time series convolutions were merged in another convolutional pass, and the two branches were combined using multi-head attention. A visual representation of the architecture can be seen in fig. 6.3, with the layer blocks defined in fig. 6.4. Convolutions were chosen for the time series input because they are relatively parameter efficient. The depth of the network guarantees that the receptive field of the downstream neurons covers the entire input space so the positional dependence of the signal is maintained. No pooling layers were used. In general, fully-connected networks were found to be difficult to train (which has been observed in other studies [26]) and are only used when projecting representations to higher or lower dimensions. The cause of this training difficulty for dense networks is unknown. In general, including residual connections appeared to help improve training stability.

During experimentation, it was found that preprocessing the signals before combining into the energy function worked much better than combining all the signals near the inputs. This may indicate that feature learning is very important for the EBM to learn a representative energy surface, at least for this dataset. The energy layers – responsible for combining the learned representations

Energy model architecture

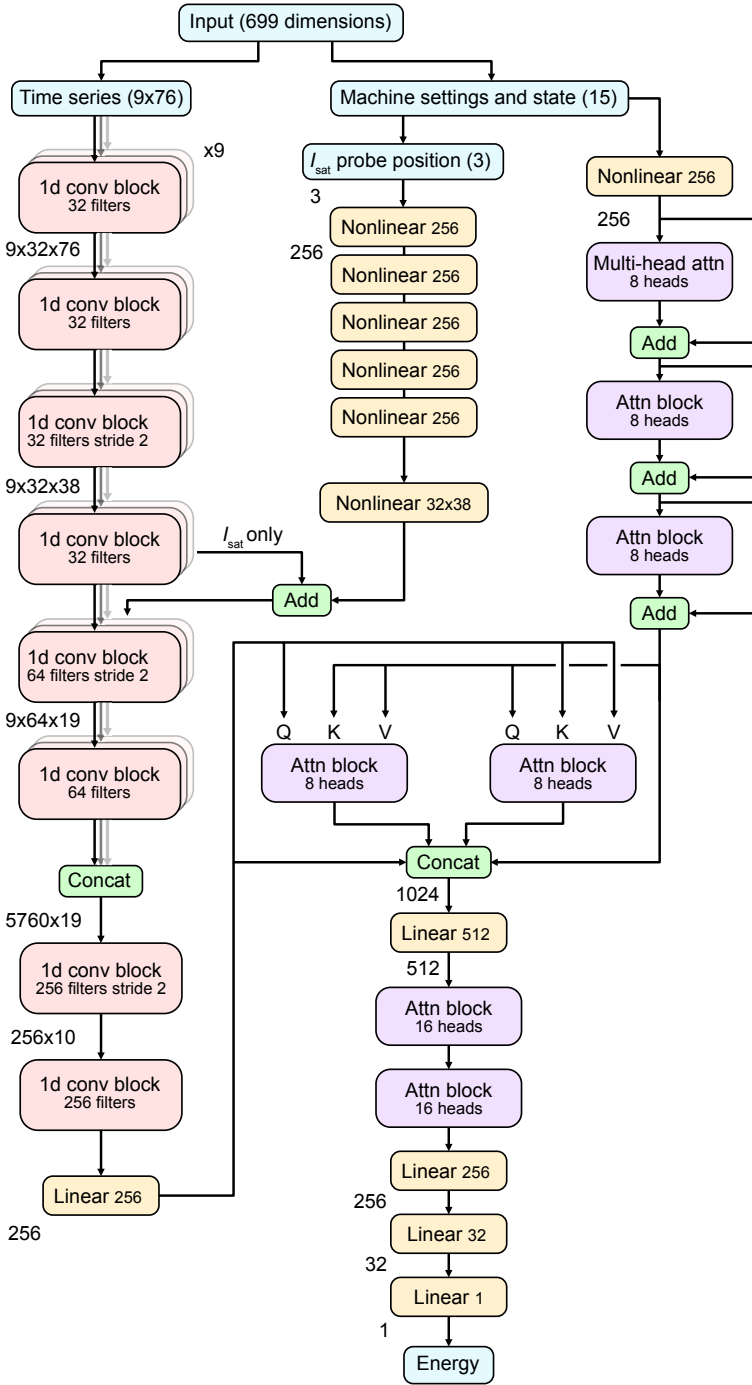


Figure 6.3: Architecture for learning the energy function. Two main branches were used: processing the time series inputs and the machine settings. The layer blocks are defined in fig. 6.4.

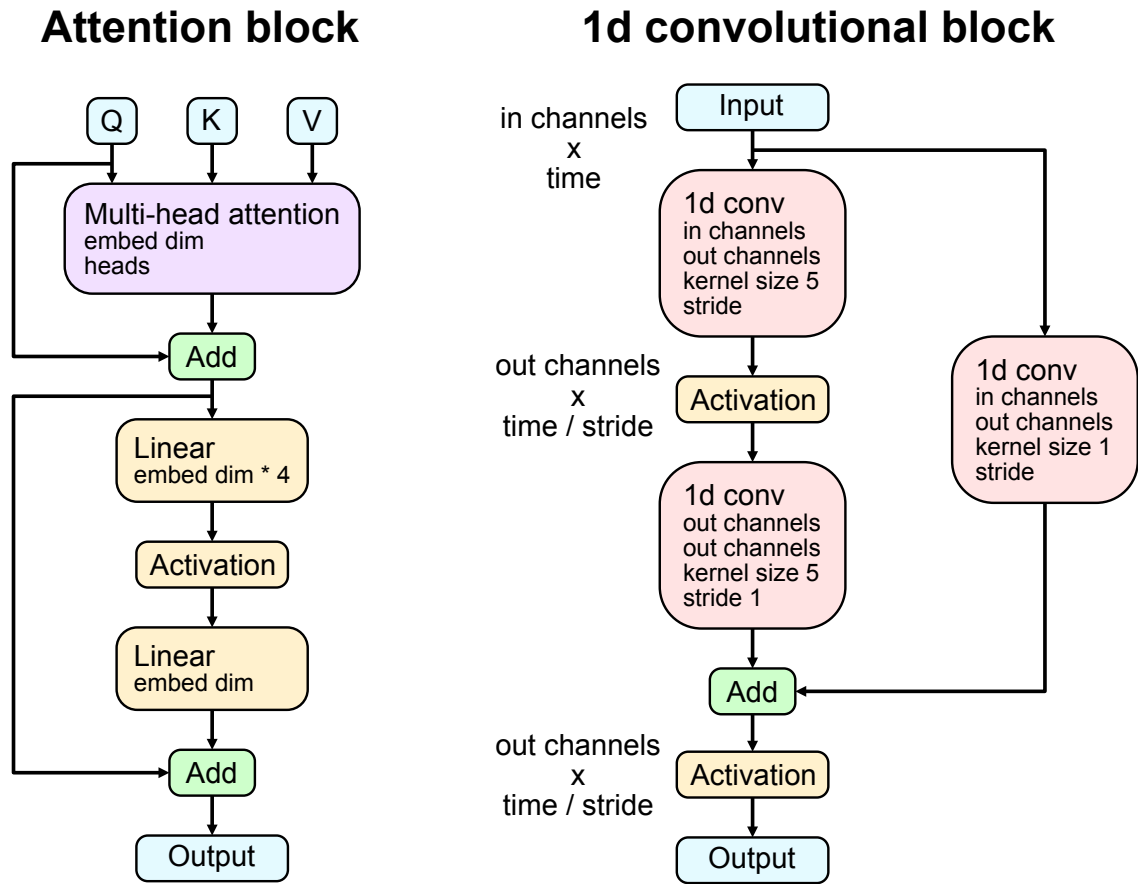


Figure 6.4: Layer blocks used in the EBM architecture. Residual layers work well.

– were not as sensitive to parameter count. Despite this importance of feature learning, attempts to train the EBM with a pretrained encoder was difficult and could not converge to a good energy surface.

The EBM struggled to model the positional dependence of the I_{sat} signal, so a branch was added from the probe positions to add onto the intermediate representation of the I_{sat} signal. This branch improved performance of the positional mapping, but is not as accurate as the simple feed forward model from the previous chapter. Performance may be improved by improving model capacity, or reducing the learning rate in conjunction with the step size.

The architecture developed for this study is not optimal – it is simply one that works and is mostly stable. The architecture likely has ample room for improvement, particularly regarding the learning of the probe position. Additional, intermediate connections between the time series and 1d inputs should be explored. The current architecture can be thought of as a medium-term multimodal fusion design (early fusion did not work).

6.6 Unconditional sampling

Unconditional samples can tell us how well the network is modeling the data distribution. 5000 samples were generated with the inputs initialized from a uniform distribution between -1 and 1. These samples ran for 120 steps of Langevin dynamics with the default step size of 0.01. On an RTX 3090, this process took 64 seconds. The MCMC trajectories of unconditional samples steps these samples can be seen in fig. 6.5. Notably after a small number of sampling steps – around 30 – the energy surface gradients appear to flatten out and thus the energies of the samples level off. This leveling point may be determined by the number of samples steps used while training the model: each randomly-initialized sample runs for 30 steps before being added to the replay buffer. Despite converging relatively quickly, long-run MCMC chains look just as realistic as shorter-duration chains and do not exhibit the burn-in or high-saturation that has been observed in other work [41].

Given that the data distributions are highly multi-modal, it is important that the EBM captures




Figure 6.5: MCMC energies, gradients, and integrated trajectory length for unconditional samples. Left: the model converges after approximately 50 sampling steps. Middle: the gradients approach an asymptote; long-term samples are realistic. Right: integrated trajectory length show that individual MCMC trajectories vary in total distance traveled along the energy surface.

many, if not all, modes of the distribution. Representative examples of these distributions, namely of diode 3 at 16 ms (chosen arbitrarily) and the mirror field, can be seen in fig. 6.6. The full unconditional distribution for each input (or at 16 ms for time-series data) can be seen in fig. 6.7. Notably, although most – if not all – modes of the distribution are covered, the mass associated with each mode may not agree. This behavior is evident in the aggregate energy distribution seen in fig. 6.8. On average, the unconditional samples have higher energy than the data. In terms of the scaled values of all of the inputs, the model appears to struggle with extreme values. This could hint towards the need for data augmentation so that chains can properly mix, or a need for a different random initialization before commencing Langevin dynamics.

6.7 Diagnostic reconstruction via conditional sampling

Conditional sampling of these models is typically performed by initializing a portion of the inputs on real (or otherwise desired) data and only sampling the other inputs. In practice, this can be done by freezing the gradients of the inputs to be conditioned on. In this work, we use a different




Figure 6.6: Unconditional samples of diode 3 at 16 ms and the mirror coil magnetic field inputs, chosen as representative examples. The EBM learns all modes of the distributions, though the probability mass is not well-aligned.

approach of instead modifying the energy function based on the data to be conditioned on. Given the energy model $E(x)$, where x is an input into the model, we add an auxiliary energy function $F(x)$ that is added to $E(x)$. This creates a new aggregate energy function $E(x) + F(x)$. By definition of energy $p_E(x) \propto \exp(-E(x))$, this is equivalent to multiplying the two distributions $p_E(x) \cdot p_F(x)$. In other words, adding this auxiliary energy function $F(x)$ to $E(x)$ implies we are sampling over the distribution $p_F \cap p_F$.

The choice of auxiliary energy function $F(x)$ can make a large difference on the conditional samples. We choose $F(x)$ to be a quadratic energy function centered on the data:

$$F(x) = \left(\frac{x - x_{\text{fixed}}}{\sigma_F} \right)^2 \quad (6.5)$$

where σ_F controls the horizontal scale of the quadratic function. Interpreted as a probability distribution, this is a Gaussian with standard deviation σ_F . The minimal width for stable sampling appears to be $\approx (2\epsilon)^2$. This makes sense from a sampling point of view: if the width of $F(x)$ approach the step size, a Langevin update step may place the particle at a much higher energy with much larger gradients. This behavior also ties the conditional sampling directly to the training process because the conditional samples are naturally limited to the step sized used while training,

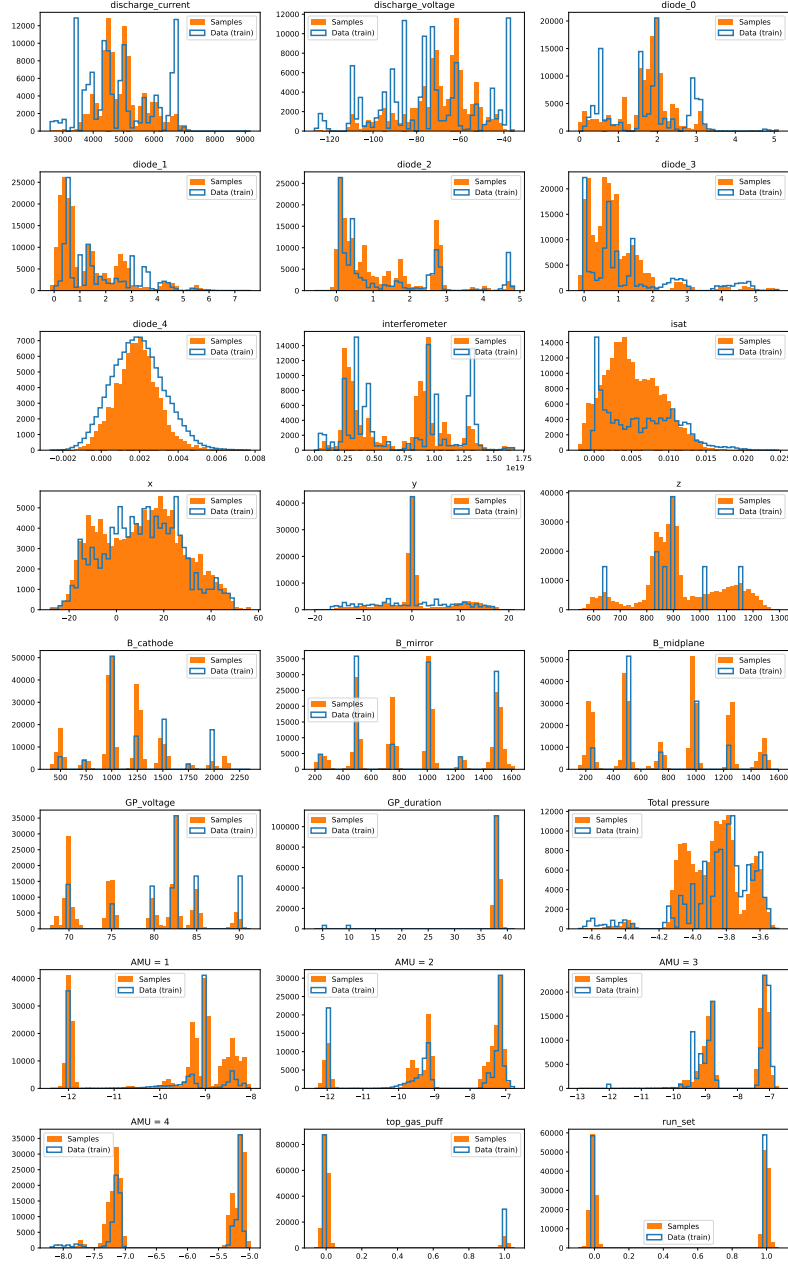


Figure 6.7: Unconditional samples of all inputs, or at 16 ms – chosen arbitrarily – for time series. The model appears to learn most, if not all, modes of the distributions, but can perform poorly when modeling the probability mass, such as in the I_{sat} distribution.

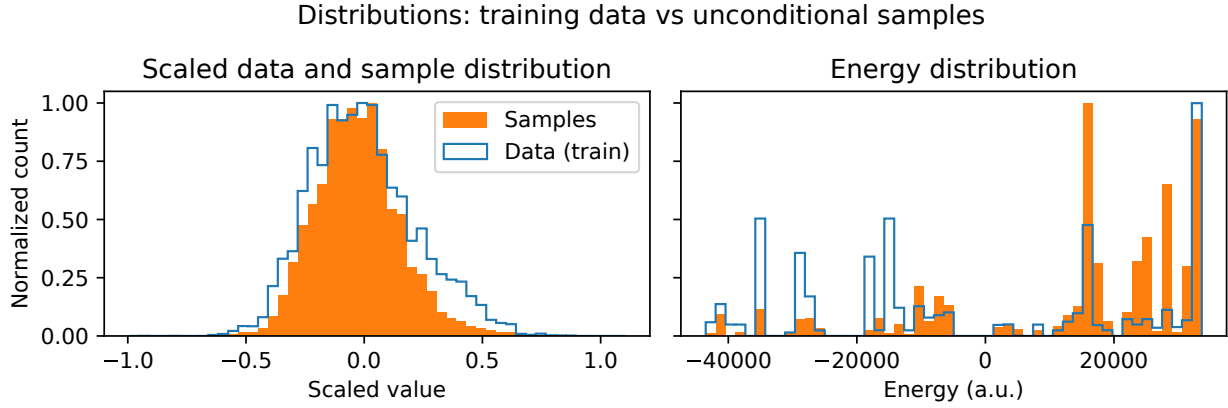


Figure 6.8: Left: all scaled inputs from the training set vs samples inputs. The distributions are similar, but the EBM does not appear to learn more extreme values. Right: corresponding energy distribution. The EBM learns all the modes, but the probability mass is not properly distributed.

and thus are also tied to the fidelity of the model. Other distributions were used, such as a Laplace distribution via $F(x) = |x - x_{\text{fixed}}|$, but the samples produced lacked the diversity seen when using a quadratic $F(x)$. Note that the normalizing constants of the probability distribution $p_F(x)$ do not matter because they vanish after $-\log(\cdot)$ is applied and the energy gradients $\nabla_x F(x)$ are taken.

6.7.1 Sampling interferometer signals

Using this conditional sampling method, we now reconstruct interferometer signals. We choose to reconstruct interferometer signals because the results are more easily interpreted physically than the diodes and the discharge current. Using conditional sampling, we compare the samples when only the LAPD control inputs are given and compare with sampling when other diagnostics are given – namely the discharge current, diodes, and I_{sat} . The machine inputs are the discharge voltage, gas puff duration and voltage, gas pressures, and magnetic field configuration. We also compare the traditional method of initializing on data and freezing gradients. A summary of the results and standard deviation of the distributions can be seen in table. 6.1. The full time series of the samples can be seen in fig. 6.9. Note that the training set RMSE is an order of magnitude better than the test

Table 6.1: RMSE and 2σ of conditional interferometer samples for test set and DR2_02

| Given: | LAPD settings only | All signals | Frozen gradients |
|--------------------|-----------------------|-----------------------|-----------------------|
| RMSE (test set) | 4.12×10^{18} | 2.91×10^{18} | 3.13×10^{18} |
| RMSE (DR2_02) | 3.77×10^{18} | 3.54×10^{18} | 2.51×10^{18} |
| 2σ (DR2_02) | 6.93×10^{18} | 8.37×10^{18} | 3.38×10^{17} |
| Training RMSE | 4.40×10^{17} | | |

set, which is in line with expectations given the limited dataset diversity.

For sampling, 90 samples steps were taken with the training step size of $\varepsilon = 0.01$. The auxiliary energy function used a width of $(2\varepsilon)^2$. Interferometer traces from a single shot from each of the eight test set dataruns were sampled. 128 samples of the interferometer signal were taken from each datarun, taking approximately 9 seconds on an RTX 3090.

Sampling with only machine inputs given led to a large variety of potential interferometer signals, seen on the left of fig. 6.9. Comparing to the case where all data is given (middle of the figure) yields some interesting insights. First, the variety of the sample distribution is dramatically decreased (seen in the bottom row). The number of modes in the distribution dramatically decreases markedly from over ten to four. Second, the shapes of the curves when diagnostics are given matches better. Third, the RMSE improves slightly for this particular case, but reduces significantly – $\gtrsim 25\%$ – when computed over the entire test set (tab. 6.1). The 2σ range of the samples increases, however, perhaps indicating increased uncertainty. This uncertainty increase is counterintuitive given the increased context. The particular I_{sat} time trace did not make much of a difference on sample quality, indicating that the model was also using information in the discharge current and diode signals when reconstructing the interferometer signal.

Sampling while freezing gradients (right side of fig. 6.9) can lead to good RMSE values relative to the other samples, but the actual samples look terrible and are unphysical: negative density is impossible. Constraining the samples to be greater than 0 using an auxiliary energy function ensures positive interferometer values, but the sample quality remains very poor and ironically increases the

RMSE. The distribution of signals is also very narrow, leading to a very overconfident prediction.

Note that there is nothing special about sampling the interferometer in particular. Any diagnostic or feature of the dataset can be conditionally sampled given other data. In other words, this model allows you to predict any feature of the data given any other feature, which is incredibly powerful, and may enable opportunities such as diagnostic calibration after-the-fact.

6.8 Symmetries and trends in the energy function

Structure in the energy function itself can also be examined to extract insights from the data. One example of this can be seen in scans along the energy function for the x-axis probe position. Since LAPD plasmas are approximately azimuthally symmetric (because of the cylindrical geometry), we expect the energy function to exhibit similar behavior. Given a certain I_{sat} time series, we expect the position of the probe to be equally likely at a constant radius if the azimuthal symmetry is perfect. Energy scans along the x axis for a given shot (and I_{sat} time series) for a training and a test set datarun can be seen in fig. 6.10. When a shot near or on the magnetic axis is provided, the energy function is largely symmetric about that point. When a shot is provided further out, the energy function takes on a shape with two minima, indicating that two positions of the energy surface are likely. This behavior is obvious in the test set case of DR2_19 (yellow curve) – either side of the y axis on the x axis are nearly equally likely. For a shot near the machine wall (red curve), the energy distribution has two minima, but the minimum on the opposite side is closer than ideal, but the true symmetrical position is beyond the limits of the training data. Not all shots yield symmetric energy functions, as seen in the red curve of the training set (left side of fig. 6.10).

This direct analysis of the energy function is uncommon, if not novel, owing to the intrinsic structure of the data. Analyzing energy functions in this way requires a dataset that has strong dependence on a low dimensional subset – sufficiently low that a comprehensive grid scan is tractable. The probe coordinates and machine settings satisfy this condition: changing one of those values dramatically affects likely values of other inputs. For most datasets, and the time series

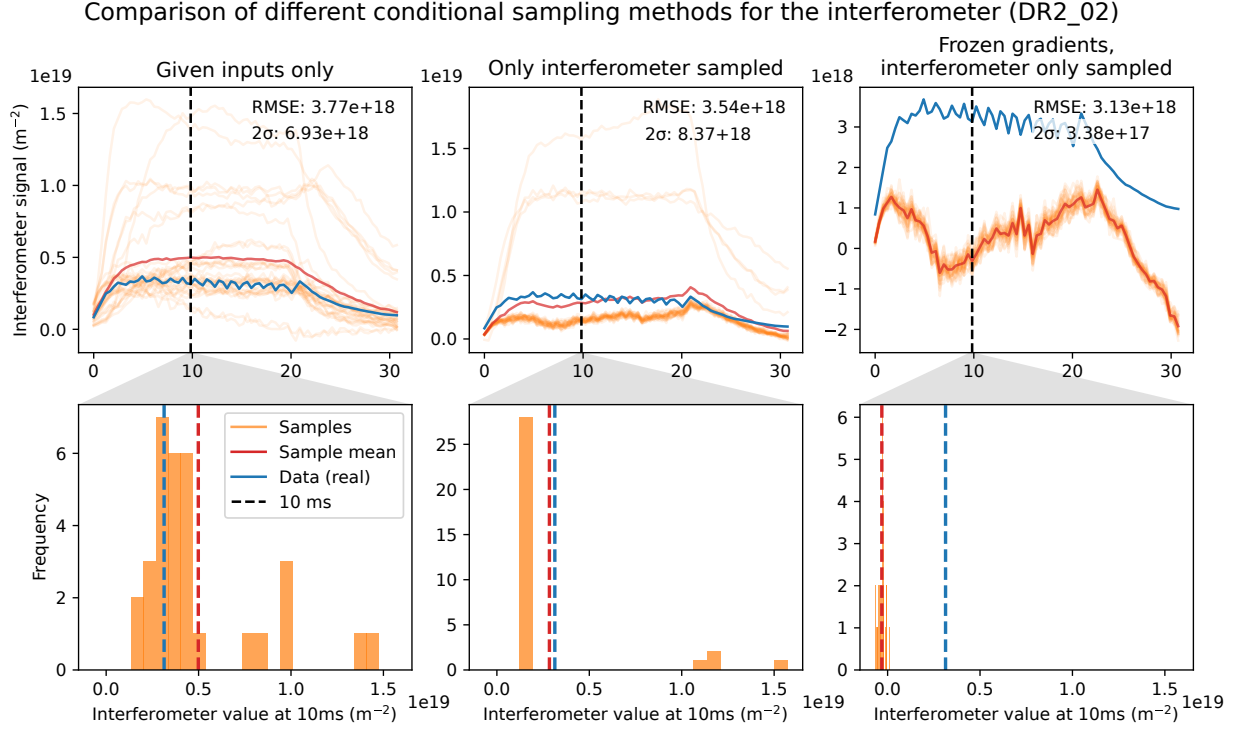


Figure 6.9: Reconstructing the interferometer signal for a test-set datarun, showing only 32 samples for clarity. Given only the inputs (left), the interferometer signal reconstruction uncertainty is quite large with many possible modes. When given other diagnostics signals, the RMSE improves by $2 \times 10^{17} \text{ m}^{-2}$, but the uncertainty increases. If the model is sampled by instead initializing all inputs on real data and freezing the gradients (right), the model produces unphysical results and is poorly calibrated. The datarun chosen (DR2_02) is representative of performance across the test set.




Figure 6.10: Scans along the x-axis input for the energy function of a real shot. When off-axis shots are provided, the energy function may encode the symmetry about the y axis.

inputs in the dataset used here, changing a single input away from the data value does not yield any particular insight. For example, changing a single pixel in an image does not significantly change the image, though changing a group of pixels could. In the case of this dataset, changing the probe position could change the I_{sat} time series significantly, so scanning along probe positions could yield insight into the dependence of the I_{sat} time series.

Certainly other symmetries in the data could be uncovered by directly analyzing this energy function (future work includes the relationship of gas puff duration and voltage to discharge power). Two- or three-dimensional scans along on particular inputs, such as gas puff or magnetic field settings, could illuminate regions of LAPD parameter space where expected behavior is identical. Although this positional symmetry presented here is obvious in retrospect, this capability highlights an important feature of EBMs. This feature is that if the inversion is not well posed, or has many possible solutions – a multivalued function – the EBM will predict many of them and not just the mean. In other words, using the energy function, anything can be predicted from anything else regardless of the invertibility of the problem. This invertibility is very important if multiple machine configurations lead to similar plasmas. Simply predicting the mean of these possibilities could lead

to wildly different plasmas.

6.9 Conclusion

In this work, we demonstrate the usefulness of energy-based models (EBMs) for modeling the data distribution. Using a multi-branch, medium-term fusion architecture, signals from multiple time-series diagnostics, machine state information, and LAPD inputs are mapped onto an energy surface. This dataset is highly multimodal in both definitions of the term: the distribution has many modes and the inputs are different modalities. This energy surface is related to a probability distribution via $p(x) = \exp(-E(x))$ and can be easily sampled using Langevin dynamics. The distribution captures all modes of the data, though the represented probability mass has some room for improvement. Using this learned distribution, it is possible to detect anomalies, reconstruct diagnostic signals given other data, and uncover trends or symmetries. When reconstructing diagnostics, additional data is very beneficial, even when the signals are from an uncalibrated diagnostic and are unanalyzed. This result implies that collecting many different diagnostics signals – even if they are basic – can be very useful and contain exploitable information. In other words, this model can exploit information we humans cannot, either because of lack of time, manpower, or the complexity of analysis.

These energy-based distributions are amenable to outside modification, as demonstrated by the novel conditional sampling technique in this work. This direct access to the distribution representations enables the scientist to shape the sampling process however desired, so that trends along any input or combination of inputs can be performed.

6.10 Future work

Given the flexibility of EBMs, there are many potential ways of improving on this work. The most obvious improvement would be to collect additional data from random LAPD configurations. Expanded probe positions would also be useful. This EBM can be extended to include

arbitrary amounts of diagnostics, enabling after-the-fact diagnostic calibration or inference of plasma parameters.

The compositional capabilities of EBMs could enable the model to expand to various machine states (e.g., cathode condition, heater temperature). An EBM could be trained on a 29-million-shot dataset which covers various machine states over 3 years. Composing this model with the one trained in this work could extend its applicability to a broader range of scenarios.

EBMs may also provide a way of joining theory and experiment. One could train an EBM on a diverse array of simulations with a variety of included physics or strengths of various effects. Jointly sampling from the sum of the theory and experiment models will effectively select the appropriate physics model for the desired scenario, enabling better predictive performance of the simulations.

In summary, EBMs provide a powerful new way of representing and analyzing data from a plasma device, and may become more useful and easier to train as compute becomes cheaper. These EBMs may enable the future automation of fusion science and device optimization.

CHAPTER 7

Conclusions

7.1 Mirror turbulence and transport

Several studies were undertaken, exploring (i) turbulence and transport in mirror configurations, (ii) optimization and trend inference in mirrors using machine learning, and (iii) generative modeling of machine state and diagnostics using energy-based models.

The primary goal of the study of turbulence and transport was to examine the coupling of the curvature-induced interchange instability with drift waves (and other modes) in the LAPD. The interchange instability was not observed: all modes seen were also present in the flat $M = 1$ case, ruling out curvature drive. The instability was likely not visible because of the many stabilization mechanisms present, such as FLR stabilization for azimuthal mode numbers $m > 4$, line-tying to the cathode which reduces the growth rate, a stabilization mechanism caused by ionization of neutral gas, and shear flow which disrupts the mode structure. Although no evidence for interchange mode was seen, other interesting behavior was present.

Surprisingly, the cross-field transport as measured by the $\tilde{E} \times B$ particle flux decreased as the mirror ratio was increased from $M = 1$ to $M = 2.68$. This decrease in particle flux was largely caused by a decrease in power of the density fluctuations for mirror ratios up to $M = 1.9$. The decrease in particle flux and fluctuation power is, in part, likely caused by the increasing density gradient scale length L_n when the mirror ratio is increased (the plasma expands radially).

Coincident peaks in density and magnetic fluctuations were observed at 12 kHz and up, with the frequency increasing with mirror ratio. These fluctuations were identified as drift-Alfvén waves,

with the frequency set by the maximum Alfvén half-wavelength using the magnetic field averaged over the length of the plasma. Additional modes were seen at 3 kHz in density and 6 kHz in magnetics. The origin of these fluctuations could not be determined, but could potentially be a combination drift waves, a rotational interchange instability, a nonlinear instability, or a conducting wall mode.

This work indicates that the edge of a mirror fusion reactor may be fairly stable, particularly if it has larger gradient scale length.

7.2 Machine learning on the LAPD

Two machine learning studies were performed on the LAPD: predicting a time-averaged ion saturation current (I_{sat} signal) and generative modeling (creating a joint distribution) of time-series diagnostics and machine state. Both of these studies utilized a diverse dataset collected on the LAPD. Machine configurations were sampled – using Latin hypercube sampling (LHS) – by varying gas puff settings, magnetic field configurations, and discharge parameters. This is the first time on the LAPD – and possibly in magnetized plasmas – that machine settings have been set using a sampling technique instead of a conventional grid search. This diversity in machine configurations leads to a large diversity in plasma discharge behavior. Although only 44 machine configurations were sampled out of 67 dataruns in total, this dataset was sufficient to train models to predict trends in LAPD behavior.

By training a relatively simple neural network (NN) model, machine settings were mapped to I_{sat} signals averaged from 10 to 20 ms. This simple model was able to learn trends over the data despite the limited diversity of the dataset. Trends inferred by varying discharge voltage or magnetic field configuration agreed with intuition and expectations. When validated on other LAPD data, the model predictions matched reality fairly well. By performing a comprehensive search over machine inputs, any function of I_{sat} can be optimized. This model was used to find optima in axial variation as measured by the standard deviation of I_{sat} signals. The predicted machine configurations for

minimal, maximal, and intermediate axial variation were applied to the LAPD. The predicted I_{sat} , although fairly inaccurate on an absolute scale, accurately reflected the trends in axial variation very well. This model could be used to identify ideal conditions for experimental studies requiring a particular I_{sat} profile.

This simple model used a β -NLL loss and ensembles to measure uncertainty. The uncertainty was not calibrated on an absolute scale, but is useful when comparing two different predictions of the model: greater model uncertainty corresponds to increased error of the prediction. The thorough uncertainty quantification used in this study, and the distinction between aleatoric and epistemic uncertainty, is unique among machine learning studies in plasma physics.

An energy-based model (EBM) was also trained on this diverse dataset. EBMs are an uncommon generative model and this work is the first time EBMs have been applied to plasmas. Instead of predicting a single average I_{sat} value, additional time-series signals were included. These time-series signals included discharge current and voltages, five diodes, the interferometer, and I_{sat} . The EBM learned an energy surface that represented the joint distribution of all these signals along with machine configuration. This multimodal energy surface was parameterized by a hybrid fusion architecture, separately learning input representations and combining them later, with some early crossover for probe positions. This energy surface was sampled using Langevin dynamics. Generating conditional samples by freezing gradients was demonstrated. A novel method of conditional sampling via energy surface modification was also demonstrated, leading to higher quality, physically realistic, samples. Unconditional samples indicated that the EBM learned all modes of the data distribution, but struggled to learn the mass associated with each mode.

In terms of scope, the flexibility of EBMs makes their analysis comparable to traditional data analysis. A couple ways of using these EBMs were demonstrated. Interferometer signals were reconstructed, both with and without additional time-series signals. Including other diagnostics, such as diodes, increased the accuracy of the predicted time traces. This use of diode signals is notable because they are uncalibrated and not analyzed, yet the ML model was able to exploit this information that would otherwise require extensive manual analysis. Given that the EBM learned a

joint probability distribution, any diagnostic signal can be reconstructed using any combination of machine settings or other diagnostics signals. This flexibility is unparalleled in machine learning in plasma physics. In addition to conditional sampling, the energy surface was directly examined. By scanning the I_{sat} probe x position for a given data point, several dips in the energy surface were occasionally seen, indicating the x-coordinate symmetry intrinsic in the plasma. The inclusion of single-dimensional parameters in the joint distribution allows for this novel analysis of the energy function. Typical EBMs only have one type of input, such as an image, which intrinsically has a high-dimensional surface. This energy surface scan highlights the ability of the EBM to learn symmetries and the ability for it to learn multivalued functions.

7.3 Future work

The interchange instability may be visible in the LAPD by increasing the plasma β and magnetic curvature. Increased β would increase the growth rate and also increase the curvature by opposing the applied field, and a larger difference in mirror and midplane fields could further encourage growth of the interchange mode. Once this mode is excited, a study can be performed to analyze the coupling of interchange to other modes. Achieving a β of 1% – an order of magnitude higher than present in this study – is possible in the LAPD using the new LaB₆ cathode.

Disentangling the different modes present in LAPD mirrors proved to be quite difficult. Further experiments with additional probes are needed to accurately measure parallel wavelength of different modes; the work presented here used only two probes which gave an effective average over all modes present. In particular, differentiating between flute like $k_{\parallel} = 0$ and finite- k modes would prove beneficial in mode identification. Additional diagnostics and probes outside of the mirror cell would also be useful in determining if modes exist that are confined to the mirror cell.

Although different length mirrors were briefly examined, the dependence of the modes present on mirror length could be more thoroughly studied. This length dependence (and thus average field) could prove beneficial when attempting to differentiate between the instabilities present.

Additional diagnostics and a diversity of machine conditions are helpful, but simulations would be the most useful in understanding which modes are possible or present. The computational approach could range from an eigenvalue solver to gyrokinetics, the latter of which would be capable of modeling the turbulence present in the device (nearly all fluctuation power was far below the cyclotron frequencies so the gyrokinetic approach is appropriate). As performed in previous work [55], a two-fluid edge transport code (BOUT++ [42]) could also be used as an intermediate approach.

If this mirror turbulence project were to be taken up again on the current LaB₆ cathode, machine learning could have a very prominent role to play in finding the stability boundary. Like the work done on optimizing I_{sat} profiles, the optimization function could instead be strength of an instability instead of axial variation. The current dataset collected – although diverse – likely does not contain sufficient information to find the stability boundary over the multidimensional LAPD actuator or settings space. More data would need to be taken for this approach to work.

The machine learning work performed here could be substantially improved by collecting a greater diversity of LAPD configurations. This greater diversity could be enabled by automated adjustment of LAPD settings (given some safe operating range). Automated setting changes and randomized probe positions could enable much faster data acquisition and sampling of LAPD parameter space since grid-based scans are no longer necessary with an ML approach. This increased diversity would dramatically improve ML model performance. The cathode – and thus plasma conditions – also change over hours, days, and weeks – inclusion of these effects would improve the generalizability of learned plasma behavior. A dataset of 29 million discharges spanning four years has already been collected and may allow categorization and accounting of intrinsic changes in LAPD plasmas.

Additional diagnostics, such as Langmuir sweeps, triple probes, floating potential, Thomson scattering, or the diamagnetic loop, could also be included. The inclusion of additional diagnostics enables diagnostic calibration after data collection. For example, I_{sat} and triple probe T_e measurements could be sampled along an interferometer chord to determine the effective area of the

Langmuir probe, assuming the Langmuir probes are relatively calibrated to each other. Probe T_e measurements could be calibrated by sampling at the Thomson scattering measurement volume.

The flexibility of EBM allows joint sampling with other distributions. As long as there is some overlap with experimental data, EBMs trained on simulation data can be jointly sampled to generate discharges that are both experimentally and theoretically plausible. This joint sampling would allow interrogation of plasma physics only available to simulations (such as transport coefficients or the entire multi-species plasma distribution function) but grounded in experimental measurements. I see this combination as potent and a promising next step for deepening our understanding of experimental plasmas.

7.4 Accelerating fusion science

The machine learning work in this paper partially addressed the questions laid out in the introduction, in short: how can we accelerate fusion science? Evidently, an optimization can be performed by collecting randomized data, to a degree of qualitative success. Despite the poor quantitative performance, qualitative guidance may be sufficient, as many incremental steps in the correct direction can achieve a desired outcome. In that way machine learning could help accelerate progress – I hope this work encourages other plasma and fusion groups to consider exploring beyond conventional experimental campaigns and increase the diversity of experiments. The ability to explore and infer trends using ML is also very useful – this trend discovery is often the goal of many experimental campaigns in fusion and plasma physics.

The energy-based modeling work also indicates another path to accelerating fusion science. The learned energy function was evidently useful and flexible for reconstructing diagnostics (or theoretically performing any inverse problem) and extracting insight. This model, as demonstrated for conditional sampling, can be combined with any other energy-based model. This model could be trained on another experimental dataset or simulation results, and the joint distribution sampled from. The performance improvement by including uncalibrated and otherwise unused

diagnostics while sampling hints that much more information is available to be autonomously exploited from our devices, even if this information is not directly interpretable. Should an ML-first approach be adopted to study fusion plasmas, dramatic gains could be achieved in optimization and understanding.

In conclusion, when done properly, machine learning can be a very useful tool for extracting insight from data and has the potential to dramatically accelerate fusion science. If fusion science is accelerated, then our iteration on experiments can likewise be faster; and mirror machines, in particular, may be the device best suited for this rapid iteration.

APPENDIX A

0D mirror optimization

This appendix describes a 0d mirror reactor optimization project. The equations and assumptions used for this optimization are described. This project provided good intuition on the limitations of simple mirrors. In addition, it gave practical experience using SymPy and Jax. The work has yet to be extended to tandem mirrors.

This work was built off an Excel spreadsheet created by Cary Forest, with contributions from Kunal Sanwalka.

A.1 List of assumptions / conditions

There are many issues and assumptions with this analysis (in no particular order):

1. Powers and particles are not strictly balanced in tandem mirrors
2. Thermal barriers are ignored which may be very important for a practical reactor
3. A fudge factor is used for electron temperatures when plug electrons are heating the central cell
4. Macrostability is not considered
5. Microstability is not considered
6. Plug confinement time is not self-consistent with plug temperatures
7. Effects of field ripple are not calculated

8. T-T and T-He3 reaction rates are not considered
9. Radial transport is assumed to be only classical
10. Cat-DD assumes instant burnup of products which is unreasonable, particularly at the high ion energies needed in mirror reactors
11. Impurities are assumed to be zero
12. Heating and magnet costs are not justified
13. All fusion power exits the plasma immediately (charged particles are collected by the direct-energy converter, neutrons absorbed by the blanket)
14. When using the temperature model from Egedal 2022 [45], we assume that the auxiliary power is much less than the beam power ($P_{aux}/P_{NBI} \ll 1$) or else the model may be inaccurate. Auxiliary power (say, to compensate for classical diffusion losses or additional ECH) can be included in this model but it would require iterative solving.
15. Burnup fraction is sufficiently small that fusion reactions are not a significant loss of fuel (ironically).
16. The DEC's collect all ion losses at a fixed efficiency

A.2 User specified parameters

A.2.1 Simple mirror endplug

1. Mirror field, plug (T): $B_{p,m}$
2. Plug (i.e., midplane) cell field (T): B_p
3. Magnet bore/throat radius (m): r_b

4. Plug length (m): L_p
5. Neutral beam energy (keV): E_{inj} or E_b
6. Beta limit (critical stability): β_{limit} (set to 0.8 [91])
7. Effective species mass (amu): μ
8. Effective atomic number (account for He and other impurities): Z_{eff}

The β_{limit} (discussed in Kotelnikov 2021 [91]) assumes a stationary plasma (no rotation, no flow out the ends), ignores finite-Larmor-radius (FLR) effects (which stabilize $m > a^2/L\rho_i$ modes), and uses the paraxial approximation ($L_m \gg a$). It also assumes $\beta \ll 1$ but this paper shows that these results match up with GDT experiments. The β_{limit} depends on the radial pressure profile falloff; the faster the falloff, the lower the β_{limit} . L is the length from midplane to throat, and L_m is the length of the mirror (highest field to lowest, I think). Profile calculations will not be included in a 0D optimization. The relevant assumptions for FLR effects and the paraxial approximation should be calculated and shown in the output to make sure they are not dramatically violated.

A.2.2 Tandem mirror

Central cell parameters defined below. Simple mirror endplugs are used on either end.

1. Central cell field (T): B_{cc}
2. Central cell to plug density ratio: n_{cc}/n_p
3. Central cell ion to plug electron temperature ratio: $T_{cc,i}/T_{p,e}$ (assumes Maxwellian)
4. Central cell to plug electron temperature ratio: $T_{cc,e}/T_{p,e}$
5. Central cell length (m): L_{cc}
6. Electron temperature fudge factor: $T_{ep, fudge}$ if electron are heating the central cell. Default value is 0.5

A.2.3 Engineering parameters

1. Vessel wall radius (with respect to plasma radius): $a_{\text{wall}} = a_{\text{wall, ratio}} \cdot a_{\text{plasma}}$
2. Blanket thickness: d_{blanket}
3. Vacuum vessel thickness: d_{vv}
4. Direct Converter Efficiency (used in the mirror exhaust): η_{DEC}
5. Thermal to electric conversion efficiency: η_{TE}
6. ECH heating efficiency: η_{ECH}
7. NBI heating efficiency: η_{NBI}
8. RF heating efficiency: η_{RF}

Optimizing the blanket and vacuum vessel thickness would probably require some neutronics calculations which would probably depend on the fuel mix, so we're just going to leave those constant in our optimization.

A.3 Fusion

DT fusion helium energy (keV): $E_{\alpha} = 3500$

A.3.1 Reactivity

DD and DT fusion cross-section parameterizations can be found in Bosch 1992. [19]. What we care most about is the fusion reaction rate per unit volume (eq. 10 in the paper):

$$\frac{dR}{dV} = \frac{n_i n_j}{1 + \delta_{ij}} \langle \sigma v \rangle \quad (\text{A.1})$$




Figure A.1: DT reactivities

This parameterization accepts ion temperature in keV and gives reactivity in cm³/s:

$$\langle \sigma v \rangle = C1 \cdot \theta \sqrt{\xi / (m_r c^2 T^3)} e^{-3\xi} \quad (\text{A.2})$$

$$\theta = T / \left[1 - \frac{T(C2 + T(C4 + TC6))}{1 + T(C3 + T(C5 + TC7))} \right] \quad (\text{A.3})$$

$$\xi = (B_G^2 / 4\theta)^{1/3} \quad (\text{A.4})$$

$$B_G = \pi \alpha Z_1 Z_2 \sqrt{2m_r c^2} \quad (\text{A.5})$$

where m_r is the reduced mass and α is the fine structure constant. The coefficients ($C1$, $C2$, and so on) are in the paper cited above. This parameterization is valid for T_i between 0.2 to 100 keV. Max error is 0.25% for DT and 0.35% and 0.3% for $DD \Rightarrow p \text{ T}$ and $DD \Rightarrow n \text{ He3}$, respectively.

A.3.2 Fusion power

DT fusion reaction rate (#/s):

$$R_{x,DT} = V n_D n_T \langle \sigma v \rangle_{DT} \quad (\text{A.6})$$




Figure A.2: DD reactivities




Figure A.3: DT, DD, and D-He3 reactivity comparison

If $n_D = n_T = n/2$, then this becomes $V \frac{n^2}{4} \langle \sigma v \rangle_{DT}$

DD fusion reaction rate (#/s):

$$R_{x,DD} = V \frac{n_{\text{plug},D}^2}{2} \langle \sigma v \rangle_{DD} \quad (\text{A.7})$$

The $\frac{1}{2}$ factor is to avoid double counting DD reactions.

Fusion power (MW):

$$P_{DT,n} = 14.1 |e| R_{x,DT} \quad (\text{A.8})$$

$$P_{DT,+} = 3.5 |e| R_{x,DT} \quad (\text{A.9})$$

$$P_{DD,n} = 2.45 |e| R_{x,DD} \cdot \frac{1}{2} \quad (\text{A.10})$$

$$P_{DD,+} = (4.02 + 0.82) |e| R_{x,DD} \cdot \frac{1}{2} \quad (\text{A.11})$$

$$(\text{A.12})$$

It's useful to split the power into charged and neutrons because energy is extract from them in different ways. Neutrons provide thermal power, charged particles heat the plasma and/or are directly captured by the DECs. The $\frac{1}{2}$ coefficient on the DD reactions assumes a 50-50 split on the DD branching ratio which actually varies with energy and may be significant above around 100 keV. If we assume the tritium produced from a DD reaction is burned instantly, then the additional power produced ("catalyed DD") is:

$$P_{\text{cat DD},n} = 14.1 |e| R_{x,DD} \cdot \frac{1}{2} \quad (\text{A.13})$$

$$P_{\text{cat DD},+} = (3.5 + 18.3) |e| R_{x,DD} \cdot \frac{1}{2} \quad (\text{A.14})$$

$$(\text{A.15})$$

We assume the tritium is burned instantly because the DT reaction rate is much higher than DD and D-He3 fusion up to around 200 keV, after which it's only slightly higher up to around 1 MeV. A more accurate estimate of fusion power would require estimates of D-He3, TT, and T-He reaction rates and density evolution of each species. A plot of reactivities can be found in Fig. A.3.

A.4 General formulae

Electron cyclotron frequency (GHz):

$$f_{\text{ECH}} = \frac{eB}{2\pi m_e c} = 28B \quad (\text{A.16})$$

Ion cyclotron frequency (MHz):

$$f_{ci} = \frac{ZeB}{2\pi m_i c} \quad (\text{A.17})$$

$$f_{ci,D} = 7.63B \quad (\text{A.18})$$

$$f_{ci,T} = 5.09B \quad (\text{A.19})$$

Here, Z is the charge state of the ion.

Electron plasma frequency (Hz):

$$f_{pe} = \frac{1}{2\pi} \sqrt{\frac{4\pi n_e e^2}{m_e}} = 8.98 \cdot 10^3 \sqrt{n_{\text{plug}}} \quad (\text{A.20})$$

Ion plasma frequency (Hz):

$$f_{pi} = \frac{1}{2\pi} \sqrt{\frac{4\pi n_i Z^2 e^2}{\mu m_p}} \quad (\text{A.21})$$

$$f_{pi,D} = 2100 \sqrt{\frac{n_{\text{plug}}}{2}} \quad (\text{A.22})$$

$$f_{pi,T} = 2100 \sqrt{\frac{n_{\text{plug}}}{3}} \quad (\text{A.23})$$

Here, μ is the mass of the ion in proton mass units (e.g. $\mu_{\text{Deuterium}} = 2$ and $\mu_{\text{Tritium}} = 3$).

Lorentz factor (γ):

$$\gamma = \sqrt{1 + \frac{T_e}{m_e c^2}} = \sqrt{1 + \frac{T_e}{511 \text{keV}}} \quad (\text{A.24})$$

Ion thermal velocity:

$$v_{Ti} = 97900 \sqrt{\frac{10^3 E_{\text{ion}}}{\mu}} \quad (\text{A.25})$$

Ion gyroradius:

$$\rho_i = \frac{mv_{\perp}}{qB} = 3.22 \cdot 10^{-3} \frac{\sqrt{\mu E_{\text{ion}}}}{B_p} \quad (\text{A.26})$$

Whistler wavelength:

$$\lambda_{\text{whistler}} = \sqrt{\frac{2\pi\Omega_e c^2}{\Pi_e^2 f}} \quad (\text{A.27})$$

$$\lambda_{\text{whistler}} = \sqrt{\frac{90 f_{\text{ECH}}}{f_{pe}^2 f_{D,2\text{nd Harmonic}}}} \quad (\text{A.28})$$

The 2nd formula is what appears on the spreadsheet and is used to estimate the size of the RF wave used for HHFW as compared to the size of the plasma. It takes into account the various constants and units used in the spreadsheet.

Collision rates (from NRL):

$$v_e = 2.91 \cdot 10^{-6} \frac{n_e \ln \Lambda}{T_e^{3/2}} \quad (\text{A.29})$$

$$v_i = 4.80 \cdot 10^{-8} \frac{Z_{\text{eff}}^4 n_i \ln \Lambda}{\mu^{1/2} T_i^{3/2}} \quad (\text{A.30})$$

These can be rearranged to give the following collision times (s):

$$\tau_{ee} = 10^{-4} \frac{T_e^{3/2}}{n_{20} \lambda_{ee}} \quad (\text{A.31})$$

$$\tau_{ii} = 1.25 \cdot 10^{-4} \frac{\mu^{1/2} E_{\text{ion}}^{3/2}}{n_{20} Z_{\text{eff}}^4} \quad (\text{A.32})$$

Slowing down times [34]

$$\tau_{i,\text{slow}} = 0.1 \frac{\mu T_e^{3/2}}{n_{20} Z^2 \lambda_{ei}} \quad (\text{A.33})$$

$$\tau_{i,\text{fast}} = \left(\left(\tau_{ii} 0.4 \log \frac{R_m}{\sqrt{1-\beta}} \right)^{-1} + \frac{1}{\tau_{\text{slow}}} \right)^{-1} \quad (\text{A.34})$$

Here, we can write $Z = 2$ for alpha particles. The 2nd equation comes from substituting the expression for T_e in a purely NBI heated case seen above.

Coulomb logarithms:

$$\lambda_{ee} = 23.5 - 0.5 \ln n_e + 1.25 \ln T_e - \left(10^{-5} + \frac{(\ln T_e - 2)^2}{16} \right)^{1/2} \quad (\text{A.35})$$

$$\lambda_{ei} = 24 - 0.5 \ln n_e + \ln T_e \quad (\text{A.36})$$

$$\lambda_{ii,\text{Cary}} = 31 - 0.5 \ln n_e + \ln T_e \quad (\text{A.37})$$

$$\lambda_{ii,\text{NRL}} = 23 - 0.5 \ln n_e + 1.5 \ln T_i \quad (\text{A.38})$$

The formula for λ_{ee} is from NRL. The formula for λ_{ei} is from NRL. However, the plasma does not fit into any of the 3 limiting cases described in the formulary. We have picked the formula based on the condition that is violated the least severely. There are 2 formulas for λ_{ii} . They do not have a large disagreement in the ranges of T_e and T_i of interest.

A.5 Radial particle transport

As of the time of writing, diffusive radial transport in mirror reactors appears to be an open question. The goal here is to provide reasonable estimates of radial particle loss and how each scale, not necessarily going for high-accuracy predictions (though being close would be nice!)

A.5.1 Classical diffusion

Assuming Fick's law and a linear density gradient from $3n$ to 0 (from Chen 5.8):

$$\tau_{\text{classical}} = \frac{nV}{A \cdot \Gamma} = \frac{na}{-2D_{\perp} \nabla n} \quad (\text{A.39})$$

where D_{\perp} is defined as

$$D_{\perp} = \eta_{\perp} n \sum T / B^2 \quad (\text{A.40})$$

and the parallel and perpendicular conductivities are

$$\eta_{\parallel} = 5.2 \cdot 10^{-5} \frac{Z \ln \Lambda}{T^{3/2}} \quad (\text{A.41})$$

$$\eta_{\perp} = 2 \cdot \eta_{\parallel} \quad (\text{A.42})$$

Combining all of this together gives

$$\tau_{\text{classical}} = \frac{a^2 B^2 T_e^{\frac{3}{2}}}{3.12 \cdot 10^{-4} \cdot n Z \sqrt{\mu} \ln \Lambda \Sigma T} \quad (\text{A.43})$$

A.5.2 Bohm diffusion

Bohm diffusivity:

$$D_{\text{Bohm}} = \frac{1}{16} \cdot \frac{T_{i,cc} \cdot 10^3}{B_{cc}} \quad (\text{A.44})$$

Normalized gyroradius (assuming deuterium):

$$r_{\text{Larmor}} = \frac{\sqrt{2mE_{\perp}}}{2eB_{cc}} = \frac{0.00791 \sqrt{T_{i,cc}}}{B_{cc}} \text{cm} \quad (\text{A.45})$$

$$\rho^* = \frac{r_{\text{Larmor}}}{a_{cc}} \quad (\text{A.46})$$

Again, using Fick's law and assuming a linear density gradient from $3n$ to 0 (so that total particle number remains $n \cdot V$), cross-field particle flux is:

$$\Gamma = -D_{\text{Bohm}} \cdot \nabla n_e \quad (\text{A.47})$$

$$= \frac{1}{16} \cdot \frac{T_{i,cc} \cdot 10^3}{B_{cc}} \cdot 3n_i \quad (\text{A.48})$$

which implies a characteristic confinement time of

$$\tau_{\text{Bohm}} = N_{\text{tot}} \left/ \frac{dN}{dt} \right. = n_i \cdot V / (\Gamma \cdot A) \quad (\text{A.49})$$

$$= n_i \cdot \pi a^2 L \left/ \left(\frac{1}{16} \cdot \frac{T_{i,cc} \cdot 10^3}{B_{cc}} \cdot 3n_i \cdot 2\pi a_{cc} L \right) \right. \quad (\text{A.50})$$

$$= \frac{8aB_{cc}}{3T_{i,cc}10^3} \quad (\text{A.51})$$

A.5.3 Gyro-Bohm diffusion

The gyro-Bohm scaling assumes cross-field transport is dominated by small ion-gyroscale turbulence. Though commonly used for tokamak scaling laws, we should be able to get some rough estimates for

mirrors. Right now it can be estimated by just diving the Bohm confinement time by the normalized gyroradius ρ^* . The gyro-Bohm estimate is then:

$$\tau_{\text{gyro-Bohm}} = \frac{8aB_{cc}}{3T_{i,cc}10^3} \cdot \frac{1}{\rho^*} \quad (\text{A.52})$$

The $\frac{1}{\rho^*}$ factor can boost the confinement time estimate by a factor of 50-100.

A.5.4 ETG-driven transport

$$\chi_{\text{ETG}} = 0.1 \frac{T_{e,cc}^{3/2}}{B_{cc}} \quad (\text{A.53})$$

$$\tau_{\text{ETG}} = \frac{a_{cc}^2}{\chi_{\text{ETG}}} \quad (\text{A.54})$$

A.6 Mirror-specific derived quantities

A.6.1 Temperatures and confinement time in a beam-heated mirror from Egedal et al 2022 [45]

Electron and ion temperature (keV) via pure beam heating [45]: we must solve a system of equations which considers the power balance of the machine. The ion temperature, given by eq. 22 in [45] is:

$$\frac{3}{2} \frac{T_i}{E_{beam}} = \frac{\exp(-\alpha) - \alpha\Gamma(0, \alpha)}{\Gamma(0, \alpha)} \quad (\text{A.55})$$

Note that eq. 22 in [45] is missing a factor of α in the numerator in front of the Γ function. The electron temperature can be found in terms of T_i and α by rearranging the definition of α (eq 21):

$$\frac{T_e}{E_{beam}} = \left(\frac{T_i}{E_{beam}} \frac{2}{3} \frac{\alpha^2 \ln R_m^2}{(22.4)^2} \right)^{1/3} \quad (\text{A.56})$$

These can be solved for with the help of a power (really energy-per-particle) balance equation (eq. 24 in [45]):

$$E_{beam} + p_{aux} = T_i + 6T_e \quad (\text{A.57})$$

where p_{aux} is the combined sources and losses, such as alpha-particle heating, plasma heat losses, RF heating, radial transport, and so on. By balancing the power lost with auxiliary heating power we can keep $p_{aux} = 0$ to avoid iteratively solving this equation. Each ion loses $T_i + e\Phi \approx T_i + 5T_e$ units of energy and each electron loses $\approx T_e$ units because only hotter electrons can surmount the ambipolar potential. p_{aux} isn't so much a power as it is the energy gained/lost per particle – an actual power would require evaluation of the confinement time (eq 29 in [45]):

$$\tau_p = \tau_{Ti}^{90} \frac{1}{\alpha_1 \lambda_1} \frac{\mathcal{H}}{T_i/E_{beam}} \int_0^1 M_1(\xi) d\xi \quad (\text{A.58})$$

where τ_{Ti}^{90} is the "scattering reactivity", α_1 , λ_1 , and M_1 are the normalization value, eigenvalue, and eigenfunction of the Lorentz scattering operator (eq 4 in [45]). By particle conservation and because NBI will be the dominant fueling mechanism, confinement time relates to density and beam current by:

$$\tau_p = eVn_b/I_{\text{NBI}} \quad (\text{A.59})$$

For ion temperature,

$$T_i = \frac{2}{3} E_{inj} \quad (\text{A.60})$$

This emerges from the relation that $E = \frac{1}{2} k_B T$ for every degree of freedom. For single particles, we assume 3 degrees of freedom to get $E = \frac{3}{2} T$ where T is expressed in eV.

For electron temperature,

$$T_e = 0.089 E_b \log_{10}(R_p)^{0.4} \quad (\text{A.61})$$

$$= 0.089 E_b \log_{10} \left(\frac{R_p}{1 - \beta} \right)^{0.4} \quad (\text{A.62})$$

which seems to give a roughly 2x higher electron temperature than the reduced model in Egedal 2022 [45], which means that our estimate will be more optimistic.

Particle confinement time (Convention: $R_p = R_m$) found in Baldwin's end-loss paper [11] equations 4.14 and 4.13. The same equation can be applied to tandem mirrors with thermal barriers

and plug cells [12]. This may be pessimistic. Equation 4.14 [11] states:

$$n\tau_{\text{Fowler/Baldwin}} = \kappa \times 10^{10} E_b^{3/2} \log R_{\text{eff}} / \log 10 \quad (\text{A.63})$$

where $[n]$ is cm^{-3} and $[E_b]$ is keV, and

$$R_{\text{eff}} = R_m / (1 + (q\phi / mE_i)) \quad (\text{A.64})$$

For 90° NBI, κ falls between 2.4 and 2.8 according to Fokker-Planck calculations [11]; it would be ~ 1.7 if the ion distribution did not have a loss-cone hole because the average energy is higher. Converting to $[n]$ in m^{-3} :

$$\tau_{\text{Fowler/Baldwin}} = 2.8 \cdot 10^{16} \frac{E_b^{3/2}}{n_e} \log R_m / \log 10 \quad (\text{A.65})$$

We may also need finite- β corrections to the mirror ratio.

A.6.2 Confinement time given by classical transport

Classical confinement time estimates assumes that transport is dominated by diffusion of gyrocenters via Coulomb collisions (from Chen section 5.8[25]). The diffusivity is:

$$D_{\text{classical}} = \eta_{\perp} n \sum T / B^2 \quad (\text{A.66})$$

where the perpendicular conductivity (for hydrogen) η_{\perp} is (temperatures in eV):

$$\eta_{\perp} = 2 \cdot \eta_{\parallel}, \quad (\text{A.67})$$

$$\eta_{\parallel} = 5.2 \cdot 10^{-5} \frac{Z \ln \Lambda_{ei}}{T_e^{3/2}} \sqrt{\mu} \quad (\text{A.68})$$

The confinement time is then (summing over species):

$$\tau_{\text{classical}} = \frac{nV}{A \cdot \Gamma} \quad (\text{A.69})$$

$$= \frac{na}{-2D_{\text{perp}} \nabla n} \quad (\text{A.70})$$

$$\tau_{\text{classical}} = \frac{aB^2}{-2\eta_{\perp} \nabla n \sum T} \quad (\text{A.71})$$

Again assuming a linear radial density profile with a peak of $3n_i$ to keep the total particle number $n_i \cdot V$:

$$\tau_{\text{classical}} = \frac{a^2 B^2 T_e^{\frac{3}{2}}}{3.12 \cdot 10^{-4} \cdot n Z \sqrt{\mu} \ln \Lambda \Sigma T} \quad (\text{A.72})$$

The aggregate confinement time is then:

$$\tau_{\text{tot}} = \frac{1}{\frac{1}{\tau_{\text{classical}}} + \frac{1}{\tau_{\text{Fowler/Baldwin}}}} \quad (\text{A.73})$$

A.6.3 End Cells/Plugs

Mirror ratio:

$$R_{\text{plug}} = \frac{B_{p,m}}{B_p} \quad (\text{A.74})$$

Radius at the midplane (mapped from bore radius):

$$a_{\text{plug}} = r_b \sqrt{\frac{B_{p,m}}{B_p}} \quad (\text{A.75})$$

Volume:

$$V_p = L_p \pi a_p^2 \quad (\text{A.76})$$

Total particle number:

$$N_{\text{tot}} = V_p n_{\text{plug}} \quad (\text{A.77})$$

Particles lost per second:

$$\frac{dN}{dt} = \frac{N_{\text{tot}}}{\tau_{\text{Fowler/Baldwin}}} \quad (\text{A.78})$$

Number of gyroradii in the plasma radius:

$$N_{\text{gyro}} = \frac{a_p}{\rho_i} \quad (\text{A.79})$$

Density (m^{-3}) at the β limit:

$$n_{\text{plug}} = B_p^2 \frac{\beta_{\text{limit}}}{2\mu_0 |e| (T_{\text{ion}} + T_e)} \quad (\text{A.80})$$

Here, T_{ion} and T_e are expressed in eV. This can be found in Wesson page 115. Rolling all the constants together and with T_i and T_e in keV:

$$n_{20} = n_{\text{plug}}/10^{20} = B_p^2 \frac{\beta_{\text{limit}}}{0.04(T_{\text{ion}} + T_e)} \quad (\text{A.81})$$

NBI Current (A):

$$I_{\text{NBI}} = |e| \frac{dN}{dt} \quad (\text{A.82})$$

The neutral beam current is enough to replace the particles lost by the end plugs. In reality, this number will be larger since the beam neutrals are ionized via charge exchange as well as ion/electron impact.

Electron heating by fast ions (MW):

$$P_{\text{e heating by fast ions}} = 10^{-3} \frac{I_{\text{NBI}} E_b}{\tau_{\text{slow}}} \quad (\text{A.83})$$

Synchrotron radiation power loss (MW) [164]:

$$P_{\text{synch}} = 6 \cdot 10^{-3} V_p n_{20} T_e \gamma^2 B_p^2 \quad (\text{A.84})$$

Bremsstrahlung radiation power loss (MW) [164]:

$$P_{\text{brem}} = 5.35 \cdot 10^{-3} n_{20}^2 Z_{\text{eff}} \sqrt{T_e} V_p \quad (\text{A.85})$$

Power loss from escaping electrons (MW):

$$P_{\text{e, endloss}} = 10^{-3} (I_{\text{NBI}} + I_{\text{cooling}}) \cdot 7T_e \quad (\text{A.86})$$

I_{cooling} is non-zero when there is current in the expander/divertor. The $7T_e$ is because only electrons with an energy greater than the ambipolar potential can escape.

Power loss from escaping fast ions (MW):

$$P_{\text{i, endloss}} = 10^{-3} I_{\text{NBI}} (E_b - T_e) \quad (\text{A.87})$$

Injected NBI Power (MW):

$$P_{\text{NBI}} = 10^{-3} I_{\text{NBI}} E_b \quad (\text{A.88})$$

Injected ECH Power (MW):

$$P_{\text{ECH}} = \frac{P_{\text{synch}}}{20} + P_{\text{e, endloss}} - (\text{Electron heating from fast ions}) \quad (\text{A.89})$$

Divide by 20 since the plasma recaptures most of the synchrotron losses are reabsorbed.

Lawson Triple Product ($10^{20} \text{keV} \cdot \text{s}/\text{m}^3$):

$$\tau_{\text{Fowler/Baldwin}} n_{20} T_i \quad (\text{A.90})$$

Neutron Flux (MW/m^2):

$$\frac{14}{17.6} \frac{P_{\text{plug}}}{4\pi a_{\text{wall}}^2} \quad (\text{A.91})$$

Burnup fraction:

$$\frac{R_{\text{x, plug, DT}}}{dN/dt} \quad (\text{A.92})$$

α particle density (10^{20}m^{-3}):

$$n_{\alpha} = \frac{I_{\text{NBI}} Q_{\text{plug}} \tau_{\alpha} E_b}{16 V_p E_{\alpha}} \quad (\text{A.93})$$

but a more intuitive way of putting it may be

$$n_{\alpha} = \frac{\tau_{\alpha} (R_{\text{x, DT}} + \frac{1}{2} R_{\text{x, DD}})}{V} \quad (\text{A.94})$$

Z_{eff} : (from Wesson section 2.16 [164]) assuming no impurities!:

$$Z_{\text{eff}} = \frac{\sum_j n_j Z_j^2}{\sum_j n_j Z_j} = \frac{n + 4n_{\alpha}}{n + 2n_{\alpha}} \quad (\text{A.95})$$

Q_{plug} :

$$Q_{\text{plug}} = \frac{P_{\text{plug}}}{P_{\text{injected}}} \quad (\text{A.96})$$

$P_{\text{electric, in}}$:

$$P_{\text{electric, in}} = P_{\text{total}} \left(\frac{1}{\eta_{\text{HS}}} - \eta_{\text{DC}} \left(1 - \frac{T_e}{E_b} \right) \right) \quad (\text{A.97})$$

$P_{\text{electric,out}}$:

$$P_{\text{electric,out}} = 0.8\eta_{HS}P_{\text{plug}} \quad (\text{A.98})$$

Q^* :

$$Q^* = \frac{Q_{\text{plug}}}{\frac{1}{\eta_{HS}} - \eta_{DC} \left(1 - \frac{T_e}{E_b} + 0.2Q_{\text{plug}} \right)} \quad (\text{A.99})$$

Q_{electric} :

$$Q_{\text{electric}} = Q^* \cdot 0.8 \cdot \eta_{HS} \quad (\text{A.100})$$

A.6.4 Tandem mirror — central cell

Radius at the midplane:

$$a_{cc} = r_b \sqrt{\frac{B_{p,m}}{B_{cc}}} \quad (\text{A.101})$$

Central cell mirror ratio:

$$R_{cc} = \frac{B_{p,m}}{B_{cc}} \quad (\text{A.102})$$

Central cell beta:

$$\beta_{cc} = \frac{2\mu_0 |e| n_{cc} (T_{cc,i} + T_{cc,e})}{B_{cc}^2} \quad (\text{A.103})$$

$\beta_{cc} \geq 1$ will lead to an infinite Pastukhov factor, so the β -enhanced mirror ratio $R_{cc,\text{eff}} = R_{cc} \left(\sqrt{1 - \beta_{cc}} \right)^{-\frac{1}{2}}$ will be limited by keeping $\beta_{cc} \leq 0.9$.

In a tandem mirror (without a thermal barrier), we assume that the central cell electrons and plug cell electrons are Maxwellian and in thermal equilibrium, and that the central cell ions are also at the same temperature (Introduction to Tandem Mirror Physics, eq 1-3 (pg 78)):

$$T_{cc,i} = T_{cc,e} = T_{\text{plug},e} \cdot T_{\text{fudge factor}} \quad (\text{A.104})$$

The plug cell electron temperature is reduced by some fudge factor because they are heating the central cell plasma. Since the electrons follow a Maxwellian distribution along field lines, they

follow the Maxwell-Boltzman relationship, where the potential difference between the plug and central cells are given by:

$$\Phi_i = \Phi_p - \Phi_c = T_{ep} \ln \left(\frac{n_p}{n_c} \right) \quad (\text{A.105})$$

The enhancement in ion confinement time in the central cell is then given by the Pastukhov factor (Pastukhov 1974, eq. 21 [116], Kesner et al. eqs. 1-3 [87]):

$$n_c \tau_i = n_c \tau_{ii} g(R) \frac{\Phi_i}{T_{ic}} \exp \left(\frac{\Phi_i}{T_{ic}} \right) \quad (\text{A.106})$$

where $g(R)$ is a weak function of the mirror ratio. We assume the $g(R)$ is:

$$g(R) = \log \left(2R_{cc} \frac{1}{\sqrt{1 - \beta_{cc}}} + 1 \right) \quad (\text{A.107})$$

The ion confinement time is then:

$$\tau_E = \text{Pastukhov} \cdot \tau_{cc,ii} \quad (\text{A.108})$$

$$= \log \left(2R_{cc} \frac{1}{\sqrt{1 - \beta_{cc}}} + 1 \right) \frac{T_{ep}}{T_{ic}} \ln \left(\frac{n_p}{n_{cc}} \right) \left(\frac{n_p}{n_{cc}} \right)^{T_{p,e}/T_{c,i}} \cdot \tau_{cc,ii} \quad (\text{A.109})$$

Since $T_{p,e} = T_{c,i}$, this reduces to

$$\tau_E = \log \left(2R_{cc} \frac{1}{\sqrt{1 - \beta_{cc}}} + 1 \right) \ln \left(\frac{n_p}{n_{cc}} \right) \left(\frac{n_p}{n_{cc}} \right) \cdot \tau_{cc,ii} \quad (\text{A.110})$$

Thermal barriers are not considered in this analysis, which enhance the central cell confinement by elevating plug electron temperatures instead of only modifying the plug-central cell density ratio (see Post 1987 eq. 10-110[124]). Thermal barriers require additional heating and ion pumpout methods. If estimates of the power requirements of thermal barriers are available, they can be easily included in this analysis and optimization process.

Power lost from the reactor by central cell particles, per meter (MW, T in keV):

$$P_{cc,loss} = 10^{-3} \pi \cdot a_{cc}^2 n_{cc} \cdot e \frac{3}{2} (T_{cc,i} + T_{cc,e}) / \tau_E \quad (\text{A.111})$$

Since this is axial power lost, it's assumed that this power (at least the ion contribution) is recovered by the DEC's.

The power lost can be account for by lowering T_e by some fudge factor, or re-heating the electrons back up to the self-consistent temperature by injecting ECH:

$$P_{aux,ECH} = P_{cc,loss} \quad (A.112)$$

The central cell will be fuelled using cold gas puffing and is ionized and heated by electrons from the plugs. The fueling current is then:

$$I_{cc,fuel} = \frac{dN_{cc}}{dt} = \pi a_{cc}^2 L_{cc} n_{cc} / \tau_E \quad (A.113)$$

Fusion Power per meter (MW/m):

$$P_{fusion} = 17.6 |e| R_x \quad (A.114)$$

Breakeven length:

$$L_{breakeven} = \frac{2P_{plug,injected}}{P_{fusion \text{ per m}}} \quad (A.115)$$

Central cell length:

$$L_{cc} = Q \cdot L_{breakeven} \quad (A.116)$$

Total fusion power (MW):

$$P_{total} = 2P_{plug} + L_{cc} P_{fusion} \quad (A.117)$$

A.6.5 Overall power balance and plant power estimates

Total electric power in:

$$P_{electric,in} = \eta_{ECH} P_{ECH} + \eta_{NBI} P_{NBI} + \eta_{RF} P_{RF} \quad (A.118)$$

Recirculating power:

$$P_{recirculating} = \eta_{DEC} (P_{fusion,charged} + P_{cc,i,endloss} + P_{plug,i,endloss}) \quad (A.119)$$

Thermal power, ignoring power generated by the blanket (the last term is thermal losses caused by DEC inefficiencies):

$$P_{\text{thermal}} = P_{\text{fusion,neutrons}} + (1 - \eta_{\text{DEC}}) \left(\frac{P_{\text{recirculating}}}{\eta_{\text{DEC}}} \right) \quad (\text{A.120})$$

Net electric power:

$$P_{\text{electric,net}} = -P_{\text{electric,in}} + P_{\text{recirculating}} + \eta_{\text{thermal}} P_{\text{thermal}} \quad (\text{A.121})$$

Q electric:

$$Q_{\text{electric}} = \frac{P_{\text{recirculating}} + \eta_{\text{thermal}} P_{\text{thermal}}}{P_{\text{electric,in}}} \quad (\text{A.122})$$

A.6.6 Instabilities

DCLC ratio (need to keep $\sim 1,000$) [93, 123]:

$$\text{DCLC ratio} = \left(\frac{f_{pi}}{f_{ci,D}} \right)^2 \quad (\text{A.123})$$

The DCLC ratio must be kept $\sim 1,000$ as the radial density gradient needed to trigger the DCLC instability is very small ($I_{\text{gradient}} < 0.01 \rho_{g,i}$ for stability). The above condition keeps the plasma radius large enough to prevent radial gradients that are sharper than those needed to trigger the DCLC instability from forming.

Interchange growth rate (s^{-1}):

$$\gamma_{\text{interchange}} = \frac{v_{Ti}}{L_p} \quad (\text{A.124})$$

Electron temperature gradient

$$\chi_{\text{ETG}} = 0.1 \frac{T_{cc,e}^{3/2}}{B_{cc}} \quad (\text{A.125})$$

$$\tau_{\text{ETG}} = \frac{a_{cc}^2}{\chi_{\text{ETG}}} \quad (\text{A.126})$$

A.7 Costs and economics

A.7.1 Heating

ECH: \$10/W

RF: \$1/W

NBI: \$5/W

A.7.2 Magnets

kA-turns of coil needed for a given field and radius:

$$I_{\text{kA-turns}} = \frac{2 \cdot B \cdot a}{1000 \cdot \mu_0} \quad (\text{A.127})$$

kA-m of superconductor needed:

$$S = 2\pi R \cdot I_{\text{kA-turns}} \quad (\text{A.128})$$

Cost per kA·m = 10^{-4} M\$ / kA·m

Cost of magnet = $S \cdot (\text{cost per kA} \cdot \text{m})$

Radii of magnet coils needed:

1. Mirror: $r_{\text{bore}} + d_{\text{vv}} (0.1\text{m}) + d_{\text{blanket}} (0.6\text{m})$
2. Plug midplane: $(a_{\text{wall, ratio}} \cdot a_{\text{plasma}}) + d_{\text{blanket}} + d_{\text{vv}} (0.2\text{m})$
3. Plug divertor: beta limit + 0.2
4. Central cell: $(a_{\text{wall, ratio}} \cdot a_{\text{plasma}}) + d_{\text{blanket}}$

For the central cell solenoid we are assuming a spacing of one coil per meter for diagnostic access. This is an adjustable parameter but will not be optimized because that would require

energetic particle confinement estimates for coil ripple. For reference, the MARS study [101] had 42 central cell magnets spaced 3.16m apart with an inner radius of roughly 2m which led to 6% field ripple, which I assume is tolerable.

1. Solenoid field: $B = \mu_0 \cdot n_{\text{cc,turns}} \cdot I$, where $n_{\text{cc,turns}}$ is number of turns per coil. This becomes

$$B = \mu_0 \cdot I_{\text{kA-turns}}$$

2. kA-m per meter length (or per coil): $S_{\text{cc}} = 2\pi a_{\text{cc}} \cdot \frac{B}{\mu_0} \cdot (1/\text{coil spacing})$

A.8 Optimization constraints

A.8.1 Midplane fields regularization via alpha particle confinement penalties

If we do not regularize field strengths, then the optimizer will bring the central cell (or plug) magnetic fields to 0 or negative. Only the midplane fields of the central cell and plugs will be regularized because the cost functions of interest tend towards higher reactor performance (and/or lower cost), and thus higher mirror ratios (and less HTS tape). The vacuum vessel should be, at minimum, four alpha gyroradii across. If an alpha is produced in the core, it will reach a distance of two gyroradii if all the energy is perpendicular to the field (aside: this is more likely with spin-polarized fuels). Doubling the vacuum vessel radius to four alpha gyroradii is the safer bet. The 3.5 MeV alpha gyroradius is:

$$r_{\text{Larmor}} = \frac{\sqrt{2mE_{\perp}}}{2eB} = \frac{0.2694\text{cm}}{B} \quad (\text{A.129})$$

This regularization is enforced as a penalty coefficient on charged particle fusion power as an exponential function of the vessel wall:

$$\mathcal{C}_{\text{power penalty}} = \begin{cases} e^{a_{\text{diff}}/r_{\text{Larmor}}} & \text{if } a_{\text{diff}} > 0 \\ 1.0 & \text{otherwise} \end{cases} \quad (\text{A.130})$$

where a_{diff} is the difference between the vessel wall and 4 alpha gyroradii: $a_{\text{diff}} = 4r_{\text{Larmor, 3.5 MeV alphas}} - a_{\text{vv}}$. These particle losses depend on the radial plasma profile and should be simulated and imple-

itly affect the optimization instead of the explicit penalty as done here. Only the 3.5 MeV alpha gyroradius is considered because it's the largest of all the usual fusion products but we apply the penalty to all fusion products. This penalty aims to be a conservative estimate.

A.9 Optimizing mirror configurations

A.9.1 Gradient descent using SymPy and JAX

Optimization is performed via gradient descent, that is, taking the gradient of some cost function \mathcal{C} with respect to some input parameter vector \vec{x} :

$$\vec{x} := \vec{x} - \nabla_{\vec{x}} \mathcal{C} \cdot \lambda \quad (\text{A.131})$$

where λ is the step size. Specific input values can be frozen by multiplying the gradient by a mask.

Equations are defined in SymPy, which are then `lambdified` to JAX expressions and then compiled by JAX's just-in-time (JIT) compiler on first run, or when `jax.jit` is called. JAX [1] calculates the gradients of \mathcal{C} with respect to \vec{x} automatically. The step size λ may be tuned; larger step sizes may not be able to be used because propagating gradients through exponential functions in the temperature calculations can be unstable. We also use 64-bit floats so that large values of α (in the reduced temperature model from Egedal 2022 [45]) remain calculable.

A.9.2 Example: optimizing Q in a simple mirror

As an example of a simple optimization task, we optimize to increase the Q of a simple mirror with classical radial transport. In this case, Q is just fusion power over NBI and ECH power. ECH power is only used to replace Bremsstrahlung and electron cyclotron losses to maintain self-consistent temperatures without requiring iterative solving. D-D fusion products are assumed to be burned instantly, though this only increases fusion power by roughly 7%.

Because the optimal solution is to decrease B_p until the mirror ratio explodes, we will add a

$1/B_p$ penalty term to keep values reasonable. The cost function is then:

$$\mathcal{C} = -Q + 1/B_p \quad (\text{A.132})$$

This cost function has no meaningful physical interpretation.

For this optimization case, we hold constant auxiliary heating power ($p_{aux} = 0$ MW), plasma beta ($\beta = 0.8$), mirror bore radius ($r_b = 0.25$ m), length ($L_p = 20$ m), tritium fraction ($T_{\text{frac}} = 0.5$), and beam energy ($E_b = 1000$ keV) and optimize only the mirror field (B_{pm} , T) and central (midplane) field (B_p , T), Z_{eff} is assumed to be 1. B_p is initialized to 6 T, and eight different values of B_{pm} are initialized between 7 and 20 T.

In this optimization, the step size λ is set to 1. The optimization was run for 1000 steps which was chosen arbitrarily—it doesn't converge in that step range (and we don't expect it to in this case).

Plots of the cost function \mathcal{C} and the gradient L2 norm for each different configuration can be seen in fig. A.4. The effects of the optimization on the fields B_{pm} and B_p can be seen in fig. A.5. The optimization favors lowering B_p until the regularization cost becomes significant at around step 60. The dramatic increase in mirror ratio leads to greater axial confinement, which decreases NBI current and power, leading to increased Q and decreased fusion power. Plots of Q and fusion power can be seen in fig. A.6. The effects of this optimization on the temperatures (or average energy in the ion case) can be seen in in fig. A.7. The increased confinement time allows the beam ions more time to slow on the background electrons, decreasing T_i and increase T_e . The decreased T_i decreases D-D reactivity but *increases* D-T reactivity at a faster rate, leading to higher fusion power. However, the lower density caused by the lower midplane field (as mandated by the β limit) causes a net *decrease* in fusion power.

A.10 Conclusions

Two insights can be gleaned from this simplified optimization task. Firstly, given optimistic physics, excessively high beam energies, incredibly high field strengths, and ignored impurity and ash




Figure A.4: The cost function and gradient magnitude for each optimization step.

accumulation, the reactor still only tops out at Q of around 2.3. This low Q implies that simple mirrors will never be a viable source of electricity. Secondly, Q is a shockingly bad optimization target because it maximizes fusion power *and* minimizes heating power simultaneously, thus high Q 's can be obtained at low fusion power as demonstrated here. An expensive, low-power reactor is not useful for power production.




Figure A.5: The mirror and midplane fields for each optimization step. Note the logarithmic x-axis — the rate of increase of the mirror field B_{pm} with respect to optimization step decreases with iteration.




Figure A.6: The total fusion power and Q for each optimization step.




Figure A.7: The ion and electron temperatures for each optimization step

APPENDIX B

Wisdom acquired

1. Simple stuff goes a long way. Try it first
2. Structure your projects so progress is linear with effort
3. Listen to your advisor
4. Write up as you go along. It helps with the thinking too
5. Wrap up projects (or at least get them to a terminal state) as soon as possible
6. Get a dog
7. Nothing really matters
8. Have two simultaneous projects ongoing
9. Good feedback is hard to find. Seek it out
10. Take risks (see item 7)
11. Don't do solo projects
12. Make lots of friends
13. When life gives you lemons, give them to Mel to make lemon bars

REFERENCES

- [1] Jax: High-performance array computing.
- [2] Lhd experiment data repository. doi:10.57451/lhd.analyzed-data.
- [3] Turbulence and transport in mirror geometries in the Large Plasma Device. 91.
- [4] J. Abbate, R. Conlin, and E. Kolemen. Data-driven profile prediction for DIII-D. *Nuclear Fusion*, 61(4):046027, April 2021.
- [5] D Ackley, G Hinton, and T Sejnowski. A learning algorithm for boltzmann machines. *Cognitive Science*, 9(1):147–169, March 1985.
- [6] Rushil Anirudh, Rick Archibald, M. Salman Asif, Markus M. Becker, Sadruddin Benkadda, Peer-Timo Bremer, Rick H. S. Budé, C. S. Chang, Lei Chen, R. M. Churchill, Jonathan Citrin, Jim A. Gaffney, Ana Gainaru, Walter Gekelman, Tom Gibbs, Satoshi Hamaguchi, Christian Hill, Kelli Humbird, Sören Jalas, Satoru Kawaguchi, Gon-Ho Kim, Manuel Kirchen, Scott Klasky, John L. Kline, Karl Krushelnick, Bogdan Kustowski, Giovanni Lapenta, Wenting Li, Tammy Ma, Nigel J. Mason, Ali Mesbah, Craig Michoski, Todd Munson, Izumi Murakami, Habib N. Najm, K. Erik J. Olofsson, Seolhye Park, J. Luc Peterson, Michael Probst, David Pugmire, Brian Sammulu, Kapil Sawlani, Alexander Scheinker, David P. Schissel, Rob J. Shalloo, Jun Shinagawa, Jaegu Seong, Brian K. Spears, Jonathan Tennyson, Jayaraman Thiagarajan, Catalin M. Ticoş, Jan Trieschmann, Jan Van Dijk, Brian Van Essen, Peter Ventzek, Haimin Wang, Jason T. L. Wang, Zhehui Wang, Kristian Wende, Xueqiao Xu, Hiroshi Yamada, Tatsuya Yokoyama, and Xinhua Zhang. 2022 Review of Data-Driven Plasma Science. *IEEE Transactions on Plasma Science*, 51(7):1750–1838, July 2023.
- [7] Rushil Anirudh, Jayaraman J. Thiagarajan, Peer-Timo Bremer, and Brian K. Spears. Improved Surrogates in Inertial Confinement Fusion with Manifold and Cycle Consistencies. *arXiv:1912.08113 [physics, stat]*, December 2019. arXiv: 1912.08113.
- [8] E. Aymerich, F. Pisano, B. Cannas, G. Sias, A. Fanni, Y. Gao, D. Böckenhoff, and M. Jakubowski. Physics Informed Neural Networks towards the real-time calculation of heat fluxes at W7-X. *Nuclear Materials and Energy*, 34:101401, March 2023.
- [9] P.A. Bagryansky, A.D. Beklemishev, and E.I. Soldatkina. Influence of Radial Electric Field on High-Beta Plasma Confinement in the Gas Dynamic Trap. *Fusion Science and Technology*, 51(2T):340–342, February 2007.
- [10] P.A. Bagryansky, A.A. Lizunov, A.A. Zuev, E. Yu. Kolesnikov, and A.L. Solomachin. Experiments with Controllable Application of Radial Electric Fields in GDT Central Cell. *Fusion Science and Technology*, 43(1T):152–156, January 2003.

- [11] D. E. Baldwin. End-loss processes from mirror machines. *Reviews of Modern Physics*, 49(2):317–339, April 1977.
- [12] D. E. Baldwin and B. G. Logan. Improved tandem mirror fusion reactor. *Phys. Rev. Lett.*, 43:1318–1321, 10 1979.
- [13] David E. Baldwin and Raynard A. Jong. Potential-confined cold plasma stabilization of loss-cone modes in mirror machines. *The Physics of Fluids*, 22(1):119–121, January 1979.
- [14] J. M. Beall, Y. C. Kim, and E. J. Powers. Estimation of wavenumber and frequency spectra using fixed probe pairs. *Journal of Applied Physics*, 53(6):3933–3940, 06 1982.
- [15] Alexei D. Beklemishev, Peter A. Bagryansky, Maxim S. Chaschin, and Elena I. Soldatkina. Vortex Confinement of Plasmas in Symmetric Mirror Traps. *Fusion Science and Technology*, 57(4):351–360, May 2010.
- [16] H. L. Berk, D. D. Ryutov, and Yu. A. Tsidulko. Temperature-gradient instability induced by conducting end walls. *Physics of Fluids B: Plasma Physics*, 3(6):1346–1354, June 1991.
- [17] Sam Bond-Taylor, Adam Leach, Yang Long, and Chris G. Willcocks. Deep Generative Modelling: A Comparative Review of VAEs, GANs, Normalizing Flows, Energy-Based and Autoregressive Models. *arXiv:2103.04922 [cs, stat]*, March 2021. arXiv: 2103.04922.
- [18] Allen H. Boozer. What is a stellarator? *Physics of Plasmas*, 5(5):1647–1655, May 1998.
- [19] H.-S Bosch and G.M Hale. Improved formulas for fusion cross-sections and thermal reactivities. *Nuclear Fusion*, 32(4):611–631, 4 1992.
- [20] F. Brochard, E. Gravier, and G. Bonhomme. Transition from flute modes to drift waves in a magnetized plasma column. *Physics of Plasmas*, 12(6):062104, June 2005.
- [21] Davide Carbone. Hitchhiker’s guide on Energy-Based Models: a comprehensive review on the relation with other generative models, sampling and statistical physics, June 2024. arXiv:2406.13661 [cs].
- [22] T. A. Carter and J. E. Maggs. Modifications of turbulence and turbulent transport associated with a bias-induced confinement transition in the Large Plasma Device. *Physics of Plasmas*, 16(1):012304, January 2009.
- [23] T. A. Casper and Gary R. Smith. Observation of alfvén ion-cyclotron fluctuations in the end-cell plasma in the tandem mirror experiment. *Phys. Rev. Lett.*, 48:1015–1018, Apr 1982.
- [24] Ian Char, Youngseog Chung, Joseph Abbate, Egemen Kolemen, and Jeff Schneider. Full Shot Predictions for the DIII-D Tokamak via Deep Recurrent Networks, April 2024. arXiv:2404.12416 [physics].

- [25] F. Chen. *Introduction to Plasma Physics and Controlled Fusion*. Springer, third edition, 2016.
- [26] Taoli Cheng and Aaron Courville. Versatile Energy-Based Probabilistic Models for High Energy Physics, January 2024. arXiv:2302.00695 [cs].
- [27] T. Cho, M. Yoshida, J. Kohagura, M. Hirata, T. Numakura, H. Higaki, H. Hojo, M. Ichimura, K. Ishii, K. Md. Islam, A. Itakura, I. Katanuma, Y. Nakashima, T. Saito, Y. Tatematsu, M. Yoshikawa, Y. Kojima, S. Tokioka, N. Yokoyama, Y. Tomii, T. Imai, V. P. Pastukhov, and S. Miyoshi. Observation of the Effects of Radially Sheared Electric Fields on the Suppression of Turbulent Vortex Structures and the Associated Transverse Loss in GAMMA 10. *Physical Review Letters*, 94(8):085002, March 2005.
- [28] J W Connor. A review of models for ELMs. *Plasma Physics and Controlled Fusion*, 40(2):191–213, February 1998.
- [29] A. J. Creely, M. J. Greenwald, S. B. Ballinger, D. Brunner, J. Canik, J. Doody, T. Fülöp, D. T. Garnier, R. Granetz, T. K. Gray, C. Holland, N. T. Howard, J. W. Hughes, J. H. Irby, V. A. Izzo, G. J. Kramer, A. Q. Kuang, B. LaBombard, Y. Lin, B. Lipschultz, N. C. Logan, J. D. Lore, E. S. Marmor, K. Montes, R. T. Mumgaard, C. Paz-Soldan, C. Rea, M. L. Reinke, P. Rodriguez-Fernandez, K. Särkimäki, F. Sciortino, S. D. Scott, A. Snicker, P. B. Snyder, B. N. Sorbom, R. Sweeney, R. A. Tinguely, E. A. Tolman, M. Umansky, O. Vallhagen, J. Varje, D. G. Whyte, J. C. Wright, S. J. Wukitch, J. Zhu, and the SPARC Team. Overview of the SPARC tokamak. *Journal of Plasma Physics*, 86(5):865860502, October 2020.
- [30] G A Daly, J E Fieldsend, G Hassall, and G R Tabor. Data-driven plasma modelling: surrogate collisional radiative models of fluorocarbon plasmas from deep generative autoencoders. *Machine Learning: Science and Technology*, 4(3):035035, September 2023.
- [31] Bhruvu Dave, Sarthak Patel, Rishi Shivani, Shishir Purohit, and Bhaskar Chaudhury. Synthetic data generation using generative adversarial network for tokamak plasma current quench experiments. *Contributions to Plasma Physics*, 63(5-6):e202200051, June 2023.
- [32] Jonas Degraeve, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego De Las Casas, Craig Donner, Leslie Fritz, Cristian Galperti, Andrea Huber, James Keeling, Maria Tsimpoukelli, Jackie Kay, Antoine Merle, Jean-Marc Moret, Seb Noury, Federico Pesamosca, David Pfau, Olivier Sauter, Cristian Sommariva, Stefano Coda, Basil Duval, Ambrogio Fasoli, Pushmeet Kohli, Koray Kavukcuoglu, Demis Hassabis, and Martin Riedmiller. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–419, February 2022.
- [33] Yuntian Deng, Anton Bakhtin, Myle Ott, Arthur Szlam, and Marc’Aurelio Ranzato. Residual Energy-Based Models for Text Generation, April 2020. arXiv:2004.11714 [cs].

- [34] Thomas James Dolan. *Fusion research*. Pergamon Press, 1982.
- [35] G. Dong, X. Wei, J. Bao, G. Brochard, Z. Lin, and W. Tang. Deep learning based surrogate models for first-principles global simulations of fusion plasmas. *Nuclear Fusion*, 61(12):126061, December 2021.
- [36] Jiaolong Dong, Jianchao Li, Yonghua Ding, Xiaoqing Zhang, Nengchao Wang, Da Li, Wei Yan, Chengshuo Shen, Ying He, Xiehang Ren, and Donghui Xia. Machine learning application to predict the electron temperature on the J-TEXT tokamak. *Plasma Science and Technology*, 23(8):085101, August 2021.
- [37] Yilun Du, Shuang Li, and Igor Mordatch. Compositional Visual Generation and Inference with Energy Based Models, December 2020. arXiv:2004.06030 [cs].
- [38] Yilun Du, Shuang Li, Yash Sharma, Joshua B. Tenenbaum, and Igor Mordatch. Unsupervised Learning of Compositional Energy Concepts, November 2021. arXiv:2111.03042 [cs].
- [39] Yilun Du, Shuang Li, Joshua Tenenbaum, and Igor Mordatch. Improved Contrastive Divergence Training of Energy Based Models. *arXiv:2012.01316 [cs]*, June 2021. arXiv: 2012.01316.
- [40] Yilun Du, Toru Lin, and Igor Mordatch. Model Based Planning with Energy Based Models. page 10, 2019.
- [41] Yilun Du and Igor Mordatch. Implicit Generation and Generalization in Energy-Based Models. *arXiv:1903.08689 [cs, stat]*, June 2020. arXiv: 1903.08689.
- [42] B.D. Dudson, M.V. Umansky, X.Q. Xu, P.B. Snyder, and H.R. Wilson. BOUT++: A framework for parallel plasma fluid simulations. *Computer Physics Communications*, 180(9):1467–1480, September 2009.
- [43] Andreas Döpp, Christoph Eberle, Sunny Howard, Faran Irshad, Jinpu Lin, and Matthew Streeter. Data-driven science and machine learning methods in laser–plasma physics. *High Power Laser Science and Engineering*, 11:e55, 2023.
- [44] D. A. D’Ippolito and B. Hafizi. Low- m ballooning stability of an axisymmetric sharp-boundary tandem mirror. *The Physics of Fluids*, 24(12):2274–2279, December 1981.
- [45] J. Egedal, D. Endrizzi, C.B. Forest, and T.K. Fowler. Fusion by beam ions in a low collisionality, high mirror ratio magnetic mirror. *Nuclear Fusion*, 62(12):126053, 11 2022.
- [46] D. Endrizzi, J.K. Anderson, M. Brown, J. Egedal, B. Geiger, R.W. Harvey, M. Ialovega, J. Kirch, E. Peterson, Yu.V. Petrov, J. Pizzo, T. Qian, K. Sanwalka, O. Schmitz, J. Wallace, D. Yakovlev, M. Yu, and C.B. Forest. Physics basis for the Wisconsin HTS Axisymmetric Mirror (WHAM). *Journal of Plasma Physics*, 89(5):975890501, October 2023.

- [47] E. T. Everson, P. Pribyl, C. G. Constantin, A. Zylstra, D. Schaeffer, N. L. Kugland, and C. Niemann. Design, construction, and calibration of a three-axis, high-frequency magnetic probe (B-dot probe) as a diagnostic for exploding plasmas. *Review of Scientific Instruments*, 80(11):113505, November 2009.
- [48] A. Fasoli, B. Labit, M. McGrath, S. H. Müller, G. Plyushchev, M. Podestà, and F. M. Poli. Electrostatic turbulence and transport in a simple magnetized plasma. *Physics of Plasmas*, 13(5):055902, May 2006.
- [49] John R. Ferron and Alfred Y. Wong. The dependence of the drift cyclotron loss cone instability on the radial density gradient. *Physics of Fluids*, 27(5):1287, 1984.
- [50] John R. Ferron, Alfred Y. Wong, Guy Dimonte, and Bernard J. Leikind. Interchange stability of an axisymmetric, average minimum- B magnetic mirror. *The Physics of Fluids*, 26(8):2227–2233, August 1983.
- [51] C.B. Forest, J.K. Anderson, D. Endrizzi, J. Egedal, S. Frank, K. Furlong, M. Ialovega, J. Kirch, R.W. Harvey, B. Lindley, Yu.V. Petrov, J. Pizzo, T. Qian, K. Sanwalka, O. Schmitz, J. Wallace, D. Yakovlev, and M. Yu. Prospects for a high-field, compact break-even axisymmetric mirror (BEAM) and applications. *Journal of Plasma Physics*, 90(1):975900101, February 2024.
- [52] S. Fornaca, Y. Kiwamoto, and N. Rynn. Experimental stabilization of interchange mode by surface line tying. *Phys. Rev. Lett.*, 42:772–776, Mar 1979.
- [53] S. J. Frank, J. Viola, Yu V. Petrov, J. K. Anderson, D. Bindl, B. Biswas, J. Caneses, D. Endrizzi, K. Furlong, R. W. Harvey, C. M. Jacobson, B. Lindley, E. Marriott, O. Schmitz, K. Shih, and C. B. Forest. Integrated modelling of equilibrium and transport in axisymmetric magnetic mirror fusion devices, November 2024. arXiv:2411.06644 [physics].
- [54] E. Fransson, A. Gillgren, A. Ho, J. Borsander, O. Lindberg, W. Rieck, M. Åqvist, and P. Strand. A fast neural network surrogate model for the eigenvalues of QuaLiKiz. *Physics of Plasmas*, 30(12):123904, December 2023.
- [55] B. Friedman, T. A. Carter, M. V. Umansky, D. Schaffner, and I. Joseph. Nonlinear instability in simulations of Large Plasma Device turbulence). *Physics of Plasmas*, 20(5):055704, 05 2013.
- [56] Brett Cory Friedman. *Simulation Analysis of Zero Mean Flow Edge Turbulence in LAPD*. PhD thesis, 2013.
- [57] Yichen Fu, David Eldon, Keith Erickson, Kornee Kleijwegt, Leonard Lupin-Jimenez, Mark D. Boyer, Nick Eidietis, Nathaniel Barbour, Olivier Izacard, and Egemen Kolemen. Machine learning control for disruption and tearing mode avoidance. *Physics of Plasmas*, 27(2):022501, February 2020.

- [58] Jim A. Gaffney, Scott T. Brandon, Kelli D. Humbird, Michael K. G. Kruse, Ryan C. Nora, J. Luc Peterson, and Brian K. Spears. Making inertial confinement fusion models more predictive. *Physics of Plasmas*, 26(8):082704, August 2019.
- [59] Ruiqi Gao, Yang Lu, Junpei Zhou, Song-Chun Zhu, and Ying Nian Wu. Learning Generative ConvNets via Multi-grid Modeling and Sampling. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9155–9164, Salt Lake City, UT, USA, June 2018. IEEE.
- [60] W. Gekelman, P. Pribyl, Z. Lucky, M. Drandell, D. Leneman, J. Maggs, S. Vincena, B. Van Compernelle, S. K. P. Tripathi, G. Morales, T. A. Carter, Y. Wang, and T. De-Haas. The upgraded Large Plasma Device, a machine for studying frontier basic plasma physics. *Review of Scientific Instruments*, 87(2):025105, February 2016.
- [61] W. Gekelman, P. Pribyl, Z. Lucky, M. Drandell, D. Leneman, J. Maggs, S. Vincena, B. Van Compernelle, S. K. P. Tripathi, G. Morales, T. A. Carter, Y. Wang, and T. De-Haas. The upgraded Large Plasma Device, a machine for studying frontier basic plasma physics. *Review of Scientific Instruments*, 87(2):025105, February 2016.
- [62] S. Ghazaryan, M. Kaloyan, W. Gekelman, Z. Lucky, S. Vincena, S. K. P. Tripathi, P. Pribyl, and C. Niemann. Thomson scattering on the large plasma device. *Review of Scientific Instruments*, 93(8):083514, August 2022.
- [63] R.J. Goldston. *Introduction to Plasma Physics*. CRC Press, Boca Raton, 1995.
- [64] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks, June 2014. arXiv:1406.2661 [stat].
- [65] V. Gopalaswamy, R. Betti, J. P. Knauer, N. Luciani, D. Patel, K. M. Woo, A. Bose, I. V. Igumenshchev, E. M. Campbell, K. S. Anderson, K. A. Bauer, M. J. Bonino, D. Cao, A. R. Christopherson, G. W. Collins, T. J. B. Collins, J. R. Davies, J. A. Delettretz, D. H. Edgell, R. Epstein, C. J. Forrest, D. H. Froula, V. Y. Glebov, V. N. Goncharov, D. R. Harding, S. X. Hu, D. W. Jacobs-Perkins, R. T. Janezic, J. H. Kelly, O. M. Mannion, A. Maximov, F. J. Marshall, D. T. Michel, S. Miller, S. F. B. Morse, J. Palastro, J. Peebles, P. B. Radha, S. P. Regan, S. Sampat, T. C. Sangster, A. B. Sefkow, W. Seka, R. C. Shah, W. T. Shmyada, A. Shvydky, C. Stoeckl, A. A. Solodov, W. Theobald, J. D. Zuegel, M. Gatu Johnson, R. D. Petrasso, C. K. Li, and J. A. Frenje. Tripled yield in direct-drive laser fusion through statistical modelling. *Nature*, 565(7741):581–586, January 2019.
- [66] G. E. Guest and E. G. Harris. Flute stabilization via electrostatically confined cold electrons. *Phys. Rev. Lett.*, 27:1500–1503, Nov 1971.
- [67] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On Calibration of Modern Neural Networks. <http://arxiv.org/abs/1706.04599>, August 2017. arXiv:1706.04599 [cs].

- [68] H. W. Hendel, T. K. Chu, and P. A. Politzer. Collisional Drift Waves—Identification, Stabilization, and Enhanced Plasma Transport. *The Physics of Fluids*, 11(11):2426–2439, November 1968.
- [69] Geoffrey E. Hinton. Training Products of Experts by Minimizing Contrastive Divergence. *Neural Computation*, 14(8):1771–1800, August 2002.
- [70] A. Ho, J. Citrin, C. Bourdelle, Y. Camenen, F. J. Casson, K. L. Van De Plassche, H. Weisen, and JET Contributors. Neural network surrogate of QuaLiKiz using JET experimental data to populate training space. *Physics of Plasmas*, 28(3):032305, March 2021.
- [71] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models, December 2020. arXiv:2006.11239 [cs].
- [72] Jr. Hooper, E. B., D. E. Baldwin, T. K. Fowler, R. J. Kane, and W. C. Turner. Radial transport reduction in tandem mirrors using end-wall boundary conditions. *The Physics of Fluids*, 27(9):2264–2267, 09 1984.
- [73] J J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, April 1982.
- [74] W. Horton. Drift waves and transport. *Reviews of Modern Physics*, 71(3):735–778, April 1999.
- [75] W. Horton, Jean C. Perez, Troy Carter, and Roger Bengtson. Vorticity probes and the characterization of vortices in the Kelvin–Helmholtz instability in the large plasma device experiment. *Physics of Plasmas*, 12(2):022303, February 2005.
- [76] K. D. Humbird, J. L. Peterson, J. Salmonson, and B. K. Spears. Cognitive simulation models for inertial confinement fusion: Combining simulation and experimental data. *Physics of Plasmas*, 28(4):042709, April 2021.
- [77] A A Ivanov and V V Prikhodko. Gas-dynamic trap: an overview of the concept and experimental results. *Plasma Physics and Controlled Fusion*, 55(6):063001, may 2013.
- [78] Samuel Jackson, Saiful Khan, Nathan Cummings, James Hodson, Shaun De Witt, Stanislas Pamela, Rob Akers, and Jeyan Thiyagalingam. FAIR-MAST: A fusion device data management system. *SoftwareX*, 27:101869, September 2024.
- [79] Azarakhsh Jalalvand, Joseph Abbate, Rory Conlin, Geert Verdoolaege, and Egemen Kolemen. Real-time and adaptive reservoir computing with application to profile prediction in fusion plasma. *IEEE Transactions on Neural Networks and Learning Systems*, 33(6):2630–2641, June 2022.

- [80] Azarakhsh Jalalvand, SangKyeun Kim, Jaemin Seo, Qiming Hu, Max Curie, Peter Steiner, Andrew Oakleigh Nelson, Yong-Su Na, and Egemen Kolemen. Multimodal Super-Resolution: Discovering hidden physics and its application to fusion plasmas, November 2024. arXiv:2405.05908 [physics].
- [81] D. L. Jassby. Transverse Velocity Shear Instabilities within a Magnetically Confined Plasma. *The Physics of Fluids*, 15(9):1590–1604, September 1972.
- [82] John Wesson. *Tokamaks*. Clarendon Press, 3rd edition, 2004.
- [83] Alexis Juven, Marie-Hélène Aumeunier, and Julien Marot. Generative Models and Simulation to Assess Uncertainties for Tokamak Infrared Thermography. In *2024 IEEE 34th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6, London, United Kingdom, September 2024. IEEE.
- [84] B.I. Kanaev. Stabilization of drift loss-cone instability (dci) by addition of cold ions. *Nuclear Fusion*, 19(3):347, mar 1979.
- [85] Andrej Karpathy. A recipe for training neural networks. <https://web.archive.org/web/20240709000647/http://karpathy.github.io/2019/04/25/recipe/>, April 2019. Accessed: 2024-07-12.
- [86] Julian Kates-Harbeck, Alexey Svyatkovskiy, and William Tang. Predicting disruptive instabilities in controlled fusion plasmas through deep learning. *Nature*, 568(7753):526–531, April 2019.
- [87] J Kesner, MJ Gerver, BG Lane, BD Vey, RE Aamodt, PJ Catto, DA Ippolito, and JR Myra. Introduction to tandem mirror physics. 1983.
- [88] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization, January 2017. arXiv:1412.6980 [cs].
- [89] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes, December 2013. arXiv:1312.6114 [stat].
- [90] Akio Komori, Yoshiya Higuchi, Yusuke Suetsugu, Akira Yonesu, and Yoshinobu Kawai. Flute stabilization of a mirror-confined plasma by a positive ambipolar potential. *Journal of the Physical Society of Japan*, 56(8):2607–2610, 1987.
- [91] Igor Kotelnikov, Andrej Lizunov, and Qiusun Zeng. On the stability of small-scale ballooning modes in axisymmetric mirror traps. *Plasma Science and Technology*, 24(1):015102, 11 2021.
- [92] Igor A. Kotelnikov, Ivan S. Chernoshtanov, and Vadim V. Prikhodko. Electrostatic instabilities in a mirror trap revisited. *Physics of Plasmas*, 24(12):122512, December 2017.

- [93] Igor A. Kotelnikov, Ivan S. Chernoshtanov, and Vadim V. Prikhodko. Electrostatic instabilities in a mirror trap revisited. *Physics of Plasmas*, 24(12):122512, 12 2017.
- [94] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. ImageNet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, May 2017.
- [95] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. <http://arxiv.org/abs/1612.01474>, November 2017. arXiv:1612.01474 [cs, stat].
- [96] Yann LeCun, Yoshua Bengio, and T Bell Laboratories. Convolutional Networks for Images, Speech, and Time-Series. April 1995.
- [97] Yann LeCun, Sumit Chopra, Raia Hadsell, Marc’Aurelio Ranzato, and Fu Jie Huang. A Tutorial on Energy-Based Learning. page 59, 2006.
- [98] Lennart van Rijn. *Minimizing neoclassical transport in the Wendelstein 7-X stellarator using variational autoencoders*. PhD thesis, Eindhoven University of Technology, July 2022.
- [99] H. Li, L. Wang, Y.L. Fu, Z.X. Wang, T.B. Wang, and J.Q. Li. Surrogate model of turbulent transport in fusion plasmas using machine learning. *Nuclear Fusion*, 65(1):016015, January 2025.
- [100] Paulett C. Liewer. Measurements of microturbulence in tokamaks and comparisons with theories of turbulence and anomalous transport. *Nuclear Fusion*, 25(5):543–621, May 1985.
- [101] B G Logan, L J Perkins, and J D Gordon. Mirror advanced reactor study (mars): executive summary and overview.
- [102] J. E. Maggs, T. A. Carter, and R. J. Taylor. Transition from Bohm to classical diffusion due to edge rotation of a cylindrical plasma. *Physics of Plasmas*, 14(5):052507, 05 2007.
- [103] J. E. Maggs and G. J. Morales. Fluctuations associated with a filamentary density depletion. *Physics of Plasmas*, 4(2):290–299, February 1997.
- [104] J. E. Maggs, G. J. Morales, and T. A. Carter. An Alfvén wave maser in the laboratory. *Physics of Plasmas*, 12(1):013103, January 2005.
- [105] Andrew Maris, Cristina Rea, Alessandro Pau, Wenhui Hu, Bingjia Xiao, Robert Granetz, Earl Marmor, the EUROfusion Tokamak Exploitation team, the Alcator C.-Mod team, the ASDEX Upgrade team, the DIII-D. team, the EAST team, and the TCV team. Correlation of the L-mode density limit with edge collisionality, June 2024. arXiv:2406.18442 [physics].
- [106] A. Mase, A. Itakura, M. Inutake, K. Ishii, J.H. Jeong, K. Hattori, and S. Miyoshi. Control of the radial electric field and of turbulent fluctuations in a tandem mirror plasma. *Nuclear Fusion*, 31(9):1725, sep 1991.

- [107] O Meneghini, B C Lyons, J McClenaghan, S P Smith, E Belli, J Candy, P B Snyder, G M Staebler, G Snoep, C S Imai, B Grierson, J M Park, and S Mordijck. Neural-network accelerated coupled core-pedestal simulations with self-consistent transport of impurities.
- [108] R.W. Moir and R.F. Post. Yin-yang minimum- |B| magnetic-field coil. *Nuclear Fusion*, 9(3):243–251, October 1969.
- [109] Shira Morosohk and Eugenio Schuster. Real-time estimation of the electron temperature profile in DIII-D by leveraging neural-network surrogate models. *Contributions to Plasma Physics*, 63(5-6):e202200153, June 2023.
- [110] Andrea Murari, Emmanuele Peluso, Michele Lungaroni, Riccardo Rossi, Michela Gelfusa, and JET Contributors. Investigating the Physics of Tokamak Global Stability with Interpretable Machine Learning Tools. *Applied Sciences*, 10(19):6683, September 2020.
- [111] Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep Double Descent: Where Bigger Models and More Data Hurt. <http://arxiv.org/abs/1912.02292>, December 2019. arXiv:1912.02292 [cs, stat].
- [112] Erik Nijkamp, Mitch Hill, Tian Han, Song-Chun Zhu, and Ying Nian Wu. On the Anatomy of MCMC-Based Maximum Likelihood Learning of Energy-Based Models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):5272–5280, April 2020.
- [113] Erik Nijkamp, Mitch Hill, Song-Chun Zhu, and Ying Nian Wu. Learning Non-Convergent Non-Persistent Short-Run MCMC Toward Energy-Based Model, November 2019. arXiv:1904.09770 [cs, stat].
- [114] D.A. Nix and A.S. Weigend. Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, pages 55–60 vol.1, Orlando, FL, USA, 1994. IEEE.
- [115] Ryan Nora, Jayson Luc Peterson, Brian Keith Spears, John Everett Field, and Scott Brandon. Ensemble simulations of inertial confinement fusion implosions. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 10(4):230–237, August 2017.
- [116] V.P. Pastukhov. Collisional losses of electrons from an adiabatic trap in a plasma with a positive potential. *Nuclear Fusion*, 14(1):3–6, 1 1974.
- [117] A. Pau, A. Fanni, S. Carcangiu, B. Cannas, G. Sias, A. Murari, F. Rimini, and the JET Contributors. A machine learning approach based on generative topographic mapping for disruption prevention and avoidance at JET. *Nuclear Fusion*, 59(10):106017, October 2019.
- [118] A Pavone, A Merlo, S Kwak, and J Svensson. Machine learning and Bayesian inference in nuclear fusion research: an overview. *Plasma Physics and Controlled Fusion*, 65(5):053001, May 2023.

- [119] Conor Perks, Saskia Mordijck, Troy Carter, Bart Van Compernelle, Stephen Vincena, Giovanni Rossi, and David Schaffner. Impact of the electron density and temperature gradient on drift-wave turbulence in the Large Plasma Device. *Journal of Plasma Physics*, 88(4):905880405, August 2022.
- [120] E.J. Doyle (Chair Transport Physics), W.A. Houlberg (Chair Confinement Da Modelling), Y. Kamada (Chair Pedestal Edge), , V. Mukhovatov (co-Chair Transport Physics), T.H. Osborne (co-Chair Pedestal Edge), , A. Polevoi (co-Chair Confinement Da Modelling), G Bateman, J.W Connor, J.G. Cordey (Retired), T Fujita, X Garbet, T.S Hahm, L.D Horton, A.E Hubbard, F Imbeaux, F Jenko, J.E Kinsey, Y Kishimoto, J Li, T.C Luce, Y Martin, M Ossipenko, V Parail, A Peeters, T.L Rhodes, J.E Rice, C.M Roach, V Rozhansky, F Ryter, G Saibene, R Sartori, A.C.C Sips, J.A Snipes, M Sugihara, E.J Synakowski, H Takenaga, T Takizuka, K Thomsen, M.R Wade, H.R Wilson, Itpa Transport Physics Topical Group, Itpa Confinement Database And Model Group, and Itpa Pedestal And Edge Topical Group. Chapter 2: Plasma confinement and transport. *Nuclear Fusion*, 47(6):S18–S127, June 2007.
- [121] F. M. Poli, S. Brunner, A. Diallo, A. Fasoli, I. Furno, B. Labit, S. H. Müller, G. Plyushchev, and M. Podestà. Experimental characterization of drift-interchange instabilities in a simple toroidal plasma. *Physics of Plasmas*, 13(10):102104, October 2006.
- [122] P. Popovich, M. V. Umansky, T. A. Carter, and B. Friedman. Analysis of plasma instabilities and verification of the BOUT code for the Large Plasma Device. *Physics of Plasmas*, 17(10):102107, October 2010.
- [123] R. F. Post. Electrostatic instabilities in finite mirror-confined plasmas. *Physics of Fluids*, 9(4):730, 1966.
- [124] R.F. Post. The magnetic mirror approach to fusion. *Nuclear Fusion*, 27(10):1579, 10 1987.
- [125] E.J. Powers. Spectral techniques for experimental investigation of plasma diffusion due to polychromatic fluctuations. *Nuclear Fusion*, 14(5):749, nov 1974.
- [126] Yuchen Qian, Walter Gekelman, Patrick Pribyl, Tom Sketchley, Shreekrishna Tripathi, Zoltan Lucky, Marvin Drandell, Stephen Vincena, Thomas Look, Phil Travis, Troy Carter, Gary Wan, Mattia Cattelan, Graeme Sabiston, Angelica Ottaviano, and Richard Wirz. Design of the Lanthanum hexaboride based plasma source for the large plasma device at UCLA. *Review of Scientific Instruments*, 94(8):085104, August 2023.
- [127] Yuchen Qian, Walter Gekelman, Patrick Pribyl, Tom Sketchley, Shreekrishna Tripathi, Zoltan Lucky, Marvin Drandell, Stephen Vincena, Thomas Look, Phil Travis, Troy Carter, Gary Wan, Mattia Cattelan, Graeme Sabiston, Angelica Ottaviano, and Richard Wirz. Design of the Lanthanum hexaboride based plasma source for the large plasma device at UCLA. *Review of Scientific Instruments*, 94(8):085104, 08 2023.
- [128] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving Language Understanding by Generative Pre-Training.

- [129] C. Rea, K. J. Montes, A. Pau, R. S. Granetz, and O. Sauter. Progress Toward Interpretable Machine Learning–Based Disruption Predictors Across Tokamaks. *Fusion Science and Technology*, 76(8):912–924, November 2020.
- [130] C. Rea, K.J. Montes, K.G. Erickson, R.S. Granetz, and R.A. Tinguely. A real-time machine learning-based disruption predictor in DIII-D. *Nuclear Fusion*, 59(9):096016, September 2019.
- [131] C.M. Roach, M. Walters, R.V. Budny, F. Imbeaux, T.W. Fredian, M. Greenwald, J.A. Stillerman, D.A. Alexander, J. Carlsson, J.R. Cary, F. Ryter, J. Stober, P. Gohil, C. Greenfield, M. Murakami, G. Bracco, B. Esposito, M. Romanelli, V. Parail, P. Stubberfield, I. Voitsekhovitch, C. Brickley, A.R. Field, Y. Sakamoto, T. Fujita, T. Fukuda, N. Hayashi, G.M.D. Hogewij, A. Chudnovskiy, N.A. Kinerva, C.E. Kessel, T. Aniel, G.T. Hoang, J. Ongena, E.J. Doyle, W.A. Houlberg, A.R. Polevoi, ITPA Confinement Database and Modelling Topical Group, and ITPA Transport Physics Topical Group. The 2008 Public Release of the International Multi-tokamak Confinement Profile Database. *Nuclear Fusion*, 48(12):125001, December 2008.
- [132] R. Rossi, M. Gelfusa, T. Craciunescu, I. Wyss, J. Vega, and A. Murari on behalf of JET Contributors. A hybrid physics/data-driven logic to detect, classify, and predict anomalies and disruptions in tokamak plasmas. *Nuclear Fusion*, 64(4):046017, April 2024.
- [133] Riccardo Rossi, Michela Gelfusa, Andrea Murari, and on behalf of JET contributors. On the potential of physics-informed neural networks to solve inverse problems in tokamaks. *Nuclear Fusion*, 63(12):126059, December 2023.
- [134] D. D. Ryutov, H. L. Berk, B. I. Cohen, A. W. Molvik, and T. C. Simonen. Magneto-hydrodynamically stable axisymmetric mirrors. *Physics of Plasmas*, 18(9):092301, 09 2011.
- [135] Ruslan Salakhutdinov and Geoffrey Hinton. Deep boltzmann machines. In David van Dyk and Max Welling, editors, *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, pages 448–455, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA, 16–18 Apr 2009. PMLR.
- [136] Rylan Schaeffer, Mikail Khona, Zachary Robertson, Akhilan Boopathy, Kateryna Pistunova, Jason W. Rocks, Ila Rani Fiete, and Oluwasanmi Koyejo. Double Descent Demystified: Identifying, Interpreting & Ablating the Sources of a Deep Learning Puzzle. <http://arxiv.org/abs/2303.14151>, March 2023. arXiv:2303.14151 [cs, stat].
- [137] D. A. Schaffner, T. A Carter, G. D. Rossi, D. S. Guice, J. E. Maggs, S. Vincena, and B. Friedman. Modification of Turbulent Transport with Continuous Variation of Flow Shear in the Large Plasma Device. *Physical Review Letters*, 109(13):135002, September 2012.

- [138] D. A. Schaffner, T. A. Carter, G. D. Rossi, D. S. Guice, J. E. Maggs, S. Vincena, and B. Friedman. Turbulence and transport suppression scaling with flow shear on the Large Plasma Device. *Physics of Plasmas*, 20(5):055907, May 2013.
- [139] David Schaffner. *Study of Flow, Turbulence and Transport on the Large Plasma Device*. PhD thesis, 2013.
- [140] J. A. Schwartz, W. Ricks, E. Kolemen, and J. D. Jenkins. The value of fusion energy to a decarbonized United States electric grid, January 2023. arXiv:2209.09373 [physics].
- [141] Prem Seetharaman, Gordon Wichern, Bryan Pardo, and Jonathan Le Roux. AutoClip: Adaptive Gradient Clipping for Source Separation Networks. <http://arxiv.org/abs/2007.14469>, July 2020. arXiv:2007.14469 [cs, eess, stat].
- [142] Maximilian Seitzer, Arash Tavakoli, Dimitrije Antic, and Georg Martius. On the Pitfalls of Heteroscedastic Uncertainty Estimation with Probabilistic Neural Networks. <http://arxiv.org/abs/2203.09168>, April 2022. arXiv:2203.09168 [cs, stat].
- [143] J. Seo, I.H. Kim, and H. Nam. Leveraging physics-informed neural computing for transport simulations of nuclear fusion plasmas. *Nuclear Engineering and Technology*, 56(12):5396–5404, December 2024.
- [144] J. Seo, Y.-S. Na, B. Kim, C.Y. Lee, M.S. Park, S.J. Park, and Y.H. Lee. Development of an operation trajectory design algorithm for control of multiple OD parameters using deep reinforcement learning in KSTAR. *Nuclear Fusion*, 62(8):086049, August 2022.
- [145] Jaemin Seo, SangKyeun Kim, Azarakhsh Jalalvand, Rory Conlin, Andrew Rothstein, Joseph Abbate, Keith Erickson, Josiah Wai, Ricardo Shousha, and Egemen Kolemen. Avoiding fusion plasma tearing instability with deep reinforcement learning. *Nature*, 626(8000):746–751, February 2024.
- [146] Jaemin Seo, Y.-S. Na, B. Kim, C.Y. Lee, M.S. Park, S.J. Park, and Y.H. Lee. Feedforward beta control in the KSTAR tokamak by deep reinforcement learning. *Nuclear Fusion*, 61(10):106010, October 2021.
- [147] T. C. Simonen. Measurements of ion cyclotron instability characteristics in a mirror-confined plasma. *The Physics of Fluids*, 19(9):1365–1370, 09 1976.
- [148] E. J. Strait. Stability of high beta tokamak plasmas*. *Physics of Plasmas*, 1(5):1415–1431, May 1994.
- [149] Jack Taylor, Jony Castagna, Luke Mason, and Vassil Alexandrov. How surrogate models can enable integration of experiments, big data, modelling and simulation. 2021. Publisher: STFC.

- [150] Tijmen Tieleman. Training restricted Boltzmann machines using approximations to the likelihood gradient. In *Proceedings of the 25th international conference on Machine learning - ICML '08*, pages 1064–1071, Helsinki, Finland, 2008. ACM Press.
- [151] Brendan D. Tracey, Andrea Michi, Yuri Chervonyi, Ian Davies, Cosmin Paduraru, Nevena Lazic, Federico Felici, Timo Ewalds, Craig Donner, Cristian Galperti, Jonas Buchli, Michael Neunert, Andrea Huber, Jonathan Evens, Paula Kurylowicz, Daniel J. Mankowitz, and Martin Riedmiller. Towards practical reinforcement learning for tokamak magnetic control. *Fusion Engineering and Design*, 200:114161, March 2024.
- [152] Phil Travis. physicistphil/lapd-isat-predict: 2025-3-11, March 2025. doi:10.5281/zenodo.15007853.
- [153] C. K. Tsui, J. A. Boedo, P. C. Stangeby, and TCV Team. Accounting for Debye sheath expansion for proud Langmuir probes in magnetic confinement fusion plasmas. *Review of Scientific Instruments*, 89(1):013505, January 2018.
- [154] G R Tynan, A Fujisawa, and G McKee. A review of experimental drift turbulence studies. *Plasma Physics and Controlled Fusion*, 51(11):113001, November 2009.
- [155] Matias Valdenegro-Toro and Daniel Saromo. A Deeper Look into Aleatoric and Epistemic Uncertainty Disentanglement. 2022. arXiv:2204.09308 [cs.LG].
- [156] B. Van Compernelle, W. Gekelman, P. Pribyl, and C. M. Cooper. Wave and transport studies utilizing dense plasma filaments generated with a lanthanum hexaboride cathode. *Physics of Plasmas*, 18(12):123501, 12 2011.
- [157] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. 2017.
- [158] J. Vega, A. Murari, S. Dormido-Canto, G. A. Rattá, M. Gelfusa, and JET Contributors. Disruption prediction with artificial intelligence techniques in tokamak plasmas. *Nature Physics*, 18(7):741–750, July 2022.
- [159] Stephen Vincena and Walter Gekelman. Drift-Alfvén wave mediated particle transport in an elongated density depression. *Physics of Plasmas*, 13(6):064503, June 2006.
- [160] J M Vos. *Discovery of hidden Neoclassical Transport variables in Wendelstein 7-X through Variational AutoEncoder Latent Space Exploration*. PhD thesis, 2024.
- [161] Masahiro Wakatani and Akira Hasegawa. A collisional drift wave description of plasma edge turbulence. *The Physics of Fluids*, 27(3):611–618, 03 1984.
- [162] Chenguang Wan, Zhi Yu, Alessandro Pau, Xiaojuan Liu, and Jiangang Li. EAST discharge prediction without integrating simulation results. *Nuclear Fusion*, 62(12):126060, December 2022.

- [163] Y. Wei, J.P. Levesque, C.J. Hansen, M.E. Mauel, and G.A. Navratil. A dimensionality reduction algorithm for mapping tokamak operational regimes using a variational autoencoder (VAE) neural network. *Nuclear Fusion*, 61(12):126063, December 2021.
- [164] J. Wesson. *4.9 - Radiation losses*, page 100–101. Clarendon Press, 1987.
- [165] Michael Wickham and Guy Vandegrift. Curvature-induced interchange mode in an axisymmetric plasma. *The Physics of Fluids*, 25(1):52–58, January 1982.
- [166] Samuel E. Wurzel and Scott C. Hsu. Progress toward fusion energy breakeven and gain as measured against the Lawson criterion. *Physics of Plasmas*, 29(6):062103, June 2022.
- [167] M. Yoshikawa, Y. Miyata, M. Mizuguchi, N. Imai, H. Hojo, M. Ichimura, T. Kariya, I. Katanuma, Y. Nakashima, R. Minami, H. Shidara, Y. Yamaguchi, Y. Shima, Y. Ohno, F. Yaguchi, and T. Imai. Use of a Gold Neutral Beam Probe to Study Fluctuation Suppression During Potential Formation in the GAMMA 10 Tandem Mirror. *Fusion Science and Technology*, 57(4):312–319, May 2010.
- [168] Yehor Yudin. *Uncertainty Quantification and Machine Learning Surrogate Models for Multi-Scale High-Performance-Computing Plasma Physics Turbulent Transport Simulations*. PhD thesis, Technical University of Munich, April 2024.
- [169] Vít Škvára, Václav Šmídl, Tomáš Pevný, Jakub Seidl, Aleš Havránek, and David Tskhakaya. Detection of Alfvén Eigenmodes on COMPASS with Generative Neural Networks. *Fusion Science and Technology*, 76(8):962–971, November 2020.