

THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING

Data-driven electrostatic heat flux closures with sparse regression

EMIL RAAHOLT INGELSTEN

Department of Physics
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden, 2026

Data-driven electrostatic heat flux closures with sparse regression

EMIL RAAHOLT INGELSTEN

© Emil Raaholt Ingelsten, 2026

Plasma Theory research group
Division of Subatomic, High Energy and Plasma Physics
Department of Physics
Chalmers University of Technology
SE-412 96 Göteborg
Sweden
Phone: +46(0)31 772 1000

Cover:

A comparison of the heat flux predicted by the data-driven closure developed in the thesis to the true heat flux calculated from kinetic particle-in-cell data for a two-stream unstable setup. The leftmost column is the true heat flux, the middle column is the model prediction and the rightmost column is the difference between them. The three rows show data for the beam, core and combined species, respectively, from top to bottom. The data comes from one of the simulations performed in Paper B, and is also shown in Figure 8 of that article.

Printed by Chalmers Digitaltryck,
Gothenburg, Sweden 2026.

Data-driven electrostatic heat flux closures with sparse regression

EMIL RAAHOLT INGELSTEN

Department of Physics

Chalmers University of Technology

Abstract

Developing models which are both accurate and computationally efficient is a long-standing goal in plasma physics. For systems where small-scale kinetic phenomena significantly influence global dynamics, this can be challenging, as global kinetic simulation – though highly accurate – is often computationally intractable. Fluid models often provide a viable alternative, but their accuracy depends on how well the utilised *closure* captures the neglected kinetic physics.

For collisional plasmas, closures can be constructed rigorously through perturbative expansion around thermal equilibrium. However, no similarly general framework exists in the collisionless case, prompting the development of a variety of collisionless closures with limited ranges of validity. Examples include the Chew-Goldberger-Low closure, valid in the strongly magnetised limit, and Hammett-Perkins-like closures for linear Landau damping. Many problems of interest lie outside these regimes, however, often necessitating ad-hoc or theoretically poorly justified closures.

This motivates the exploration of alternative approaches to closure construction. In this thesis, we use a *data-driven* approach, employing machine learning methods to systematically construct closures from first-principles kinetic simulation data. Initially restricting ourselves to one-dimensional setups, we develop a two-step *sparse regression* procedure based on the SINDy framework and use it to discover a new six-term closure for electrostatic phenomena. The closure accurately captures both linear and nonlinear regimes of the electron two-stream instability, as well as Landau damping of Langmuir waves. Additionally, the closure generalises naturally to multi-species modelling, relevant for streaming instabilities.

Having identified the closure form in the first step of the procedure, we further show how its free parameters can be dynamically estimated from fluid quantities. For this, we illustrate how both neural networks and more interpretable methods can be leveraged – in the latter case through a newly developed framework for nonlinear sparse regression.

Keywords: Plasma physics, data-driven, heat flux, closures, sparse regression, neural networks, electrostatic, two-stream instability, Landau damping

List of Publications

Appended publications

- A** E. R. Ingelsten, M. C. McGrae-Menge, E. P. Alves and I. Pusztai,
Data-driven discovery of a heat flux closure for electrostatic plasma phenomena
Journal of Plasma Physics **91**(2) E64 (2025).
<https://doi.org/10.1017/S0022377825000285>
- B** E. R. Ingelsten, M. C. McGrae-Menge, E. P. Alves and I. Pusztai,
Data-driven multi-species heat flux closures for two-stream-unstable plasmas with nonlinear sparse regression
Accepted for publication in Physics of Plasmas.
<https://arxiv.org/abs/2511.10147>

Related publications, not included in the thesis

- C E. R. Ingelsten, M. C. McGrae-Menge, E. P. Alves and I. Pusztai,
Data-driven discovery of a heat-flux closure in a two-stream unstable plasma
Proceedings of the 50th EPS Conference on Plasma Physics, Salamanca, Spain, O.227 (2024).
<https://lac913.epfl.ch/epsppd3/2024/html/PDF/O5-227.pdf>

Statement of contribution

- Paper A** Using simulation setups designed in conjunction with I. Pusztai and E. P. Alves, I ran particle-in-cell simulations of Landau-damped Langmuir waves and two-stream instability using the OSIRIS code. I then modified SINDy-based sparse regression code previously written by M. C. McGrae-Menge and E. P. Alves to analyse the kinetic simulation data and used it to discover the six-term heat flux closure identified in the paper. I also wrote the majority of the article, with support from my co-authors, and created the figures.
- Paper B** Following up on Paper A, I ran further simulations of two-stream unstable setups, with a wider range of initial conditions. An initial analysis of the data was then performed by M. C. McGrae-Menge, showing that the three most important free closure parameters could be estimated accurately with multi-layer perceptrons. I then developed the nonlinear sparse regression methodology described in the paper, and applied it to identify a somewhat more interpretable model with largely equivalent performance. Most of the manuscript was written by me, with significant contributions to the section on neural networks from M. C. McGrae-Menge, and help from I. Pusztai and E. P. Alves. I additionally created the majority of the figures.

Acknowledgements

The thesis you see before you may only have a single listed author, but in reality there are a great many people without which this work would not have been possible. Thank you, from the bottom of my heart, to all of you.

First of all, I would like to thank my main supervisor, István Pusztai, for his unwavering support – and in particular for always being available to talk to about whatever problems I might have, be they scientific or personal. Without his help, both my initial move from quantum computing to plasma physics and the ups and downs of the research thereafter would have been immeasurably more difficult. I also want to extend my sincerest gratitude to Tünde Fülöp, my examiner, who always goes to incredible lengths to help everyone under her wing reach their full potential. Furthermore, I want to thank my co-supervisor, Daniel Graham – not just for his invaluable support and wave dynamics expertise, but also for his patience with the unexpected twists and turns of the project. Two other people who had a significant hand in the work contained in this thesis are Paulo Alves and Madox McGrae-Menge at UCLA; thank you both for all of your help and many interesting discussions, over zoom and in person.

On a slightly less research-related note, I would also like to give my heartfelt thanks to all of the people in our group, past and present. Listing everybody is impossible within these margins, but just to name a few of you: Björn, Dominika, Ida, Ida (you decide which is which), Konrad, Linus, Lise, Mufei and Oskar – thank you. Without you, I would be very surprised if I had made it this far. Talking with you at lunch and fika, whether about academia, the state of the world, the weather, or just whatever nonsense is making it into our quote book that day is genuinely a highlight of my stay at Chalmers.

Finally, I want to thank my family and non-work friends, who stick by me despite my somewhat poor track record when it comes to keeping in touch. I especially want to thank everyone at the Linguistics Olympiad for giving me both a home away from home and welcome shakeups to my everyday routine – and the same goes for Young Science Sport (UVS) more broadly, as well as for the Chalmers choir. There are many I would like to name for a more personal thank you, but fitting any representative fraction of you on the remainder of this page would require novel physics. Hopefully you know who you are.

Thank you all!

Emil Raaholt Ingelsten, Göteborg, March 2026

Contents

Abstract	i
List of Publications	iii
Acknowledgements	vii
I Overview	1
1 Introduction	3
1.1 Why Study Plasma Physics?	3
1.2 Plasma Modelling and the Closure Problem	5
1.3 Data-Driven Closure Discovery	7
1.4 An Outline of this Thesis	8
2 Plasma Physics	9
2.1 First Principles: Kinetic Modelling	9
2.2 Fluid Models and the Moment Equations	10
2.2.1 The Momentum Equation	11
2.2.2 The Pressure Equation	13
2.3 Plasma Phenomena: The Necessities	14
2.3.1 Waves in Plasmas	15
2.3.2 Landau Damping and Instabilities	19
2.4 Closures from Theory	27
2.4.1 Collisional Closures	27
2.4.2 Collisionless Closures	29
3 Machine Learning	37
3.1 Artificial Neural Networks	37
3.2 Interpretability and Sparse Regression	40
3.2.1 STLS and the SINDy Framework	42
3.2.2 Nonlinear Sparse Regression	45
4 Summary and Outlook	49
4.1 Summary of papers	49
4.2 Outlook	54

References 57

II Included papers 73

Paper A - Data-driven discovery of a heat flux closure for electrostatic plasma phenomena

Paper B - Data-driven multi-species heat flux closures for two-stream-unstable plasmas with nonlinear sparse regression

Part I

Overview

Chapter 1

Introduction

1.1 Why Study Plasma Physics?

As every plasma physicist worth their salt will inevitably tell you at some point, the overwhelming majority of visible matter in the universe is plasma. The exact number given varies from 99 % to as high as 99.999 % [1–3], but suffice it to say that only a very small percentage of the universe is actually made up of solids, liquids and gases – despite these usually being introduced simply as “the three states of matter” in pre-university physics. With this in mind, it is unsurprising that when a fourth state of matter is mentioned, it is usually in reference to plasmas.

Whether “percentage of the universe’s matter content” is a particularly important metric on its own is seldom elaborated upon. If it is, dark matter scientists have us beat – according to the reigning standard model of cosmology, Lambda-CDM, only about 15 % of the universe’s matter content is visible in the first place, with the remaining 85 % being dark matter [4, 5]. However, to paraphrase a quote [6] by mathematician G. H. Hardy (1877-1947) which I recently heard from my examiner, Prof. Tünde Fülöp, it is the duty of any scientist to slightly exaggerate both the importance of their field and their own role within it. As such, I nevertheless feel obligated to repeat the abovementioned statistic, with the qualifier “visible” added before the word “matter” to discreetly remove dark matter from the equation.

Luckily, there are – facetiousness aside – many good reasons to care about plasma physics, meaning there is little need for exaggeration in our case. While the prevalence of plasmas in the universe might not be a strong argument on its own, it nevertheless hints at a better argument: Plasmas occur throughout the universe under a vast array of different conditions, and exhibit a diverse range of intricate behaviours and interactions. This means that understanding plasmas well is key to understanding the dynamics of the universe – both at large and more locally – which is fundamentally the core aim of physics as a whole. For example, understanding the interaction between electromagnetic radiation and plasmas was vital for the development of the modern field of

astronomy, since all radiation reaching us from space is filtered through the plasma of the interstellar and intergalactic media [7, 8].

Even in the modern day, the frontiers of high-energy physics are being explored through the study of plasma phenomena. For example, the intersection between general relativity and quantum mechanics, and thus the long-standing unresolved issue of quantum gravity, is currently being studied by examining black hole accretion disks [9–11] and the electron-positron pair-plasma coronas of neutron stars [12–14]. The high-energy frontiers are also being probed in settings closer to home, e.g. through particle accelerator and laser plasma experiments here on Earth [14–17]. Apart from shining a light on the previously unexplored strong-field limits of quantum electrodynamics, the setups developed through such exploration also lend themselves to the construction of more efficient particle accelerators, with both fundamental physics and medical applications [18–21].

Outside of the laboratory, one might think plasma physics is of less immediate concern in our everyday life. In some ways, this is of course true – one rarely comes into contact with plasmas directly, after all. There is, however, one very prominent feature of our lives through which we do come into contact with plasma physics, albeit indirectly: our sun. This ball of hot plasma is the ultimate source of almost all of the heat and energy we use on a daily basis, produced through hydrogen fusion in its core. And replicating this mode of energy generation in fusion power plants like tokamaks [22–26] and stellarators [27–29] is another highly active area of study within plasma physics [30–34]. Should these efforts succeed, they may very well prove vital to successfully addressing the climate crisis, paving the way for a future of practically unlimited clean energy.

The sun does not only affect us through the light it gives us, however. It is also constantly emitting charged particles in the form of the so-called solar wind, giving rise to *space weather* here on Earth [35–39]. For the most part, we can live our lives blissfully ignorant of this, due to the strong magnetic field generated by the Earth’s molten iron core, which shields us from nearly all of the sun’s particle emissions. The only way we usually experience space weather at all is through the spectacular displays of aurora which occur when the solar wind is particularly strong or an outburst from the sun hits the Earth. In such cases, the geomagnetic field guides the incoming energetic charged particles towards the north and south pole, where they hit and excite oxygen and nitrogen molecules in the atmosphere, giving rise to the aurora as they deexcite and emit the additional energy as visible light.

In some cases, space weather can be severe enough to punch through the Earth’s magnetic shielding even at lower latitudes. Such *geomagnetic storm* conditions are often caused by especially large solar eruptions – *coronal mass ejections* – hitting the Earth. In severe cases, these magnetic storms can knock out power grids on subcontinental scales by inducing large currents in power lines, as happened in 1989, when the Hydro-Québec grid failed, causing 9-hour blackouts throughout eastern Canada and the northeastern US [40, 41]. Additionally, the solar flares occurring in conjunction with the eruption of material from the solar surface can cause communications blackouts

by disturbing the Earth’s ionosphere, which is usually used as a mirror for long-distance high-frequency radio transmission [38–40].

Accurate modelling of the plasma conditions, both in the sun’s interior and in the solar wind, is thus of vital importance for predicting and mitigating disturbances to the power and communications infrastructure on which so much of our modern lives depend. Not least because we know the sun can produce even more severe magnetic storms than the one which occurred in 1989 – a repeat of the 1859 so-called Carrington event [42, 43], potentially more than twice as strong as the 1989 storm [44], would likely incur damages on the order of trillions of USD were it to occur today [45].

1.2 Plasma Modelling and the Closure Problem

Having now hopefully conveyed why plasma physics – and accurate modelling of it in particular – is important, it is high time to define what a plasma actually *is*. Per the standard textbook definition, a plasma is “a quasi-neutral ionised gas exhibiting collective behaviour”. In other words, it is a state of matter much like a regular gas, except that a significant percentage of the atoms have been stripped of (at least some) electrons. Since neutral gas particles have thus been split into collections of charged particles, which nevertheless still have a net charge of zero over large enough scales, the fluid in question is *quasi-neutral*.

The *collective behaviour* then stems from the charged nature of the constituent particles, letting them interact over long distances via the electromagnetic field – a characteristic which regular neutral gases notably lack. As we shall see, this long-range interaction between particles (and between particles and waves) gives rise to a long list of complex phenomena, making plasmas a particularly challenging beast to model – even more so than neutral fluids, which are themselves highly nontrivial to model efficiently.

In both cases, much of the complexity simply arises from the fact that there are so many particles involved. In fact, modelling them all individually is completely infeasible, even for relatively simple setups. Thus, one must resort to statistical models. The most fundamental such description of plasmas, termed the “kinetic” approach, describes the plasma using a distribution function $f(t, \mathbf{x}, \mathbf{v})$, giving the local particle density in *phase space* – in other words, it tells you how many particles are passing through a given point \mathbf{x} with velocity \mathbf{v} at a given time t for each possible \mathbf{x} , \mathbf{v} and t . The time evolution of this distribution function can then be calculated exactly, provided one has a suitable statistical model encoding the correlations caused by individual particle collisions. The differential equation one has to solve to do this is known as the Boltzmann equation, and looks like

$$\partial_t f + \mathbf{v} \cdot \nabla f + \mathbf{a} \cdot \nabla_{\mathbf{v}} f = \left(\frac{df}{dt} \right)_{\text{coll}}. \quad (1.1)$$

Here, the right-hand side of the equation is the rate of change of f due to collisions and ∇ and $\nabla_{\mathbf{v}}$ are the gradients in physical space and velocity space,

respectively, with \mathbf{a} being the (non-collision-driven) acceleration experienced by a particle with a given position and velocity. Depending on plasma conditions, the importance of the collisional term can vary widely: while it is very important for fusion applications, in space and astrophysical contexts it is often completely negligible, due to the extremely low density of the plasma. For example, electrons in the solar wind can travel for more than a day without experiencing a single one. In fact, the *mean free path* of a solar wind electron is of the order of 1 au, meaning they on average collide only about *once* during their entire journey from the sun to the Earth’s orbit [46].

The complexity of the collisional term might make collisionless plasmas seem significantly easier to model than collisional ones. In the kinetic picture, this is true. However, for many plasma phenomena of interest – especially ones where multiple scales are important simultaneously, the Boltzmann equation ends up being too expensive to solve numerically in any reasonable time frame. Thus, one needs to take a further step down the model ladder, reaching the world of fluid modelling. And for such models, the situation is rather the reverse – collisionless plasmas are actually *harder* to model than collisional ones. To see why, we must take a closer look at the way a fluid model is derived from the kinetic picture. The way this simplification is made is by *integrating away* the three dimensions of velocity space. The distribution function then turns into the particle density $n(t, \mathbf{x})$, and the Boltzmann equation turns into the continuity equation, $\partial_t n + \nabla \cdot (n\mathbf{V}) = 0$.

Unfortunately, the second term involves the particle flux $n\mathbf{V}$, the time evolution of which is unknown. To calculate it, we can multiply the Boltzmann equation by \mathbf{v} and integrate over velocity space again – but this introduces a new unknown quantity, namely the pressure, \mathbf{p} . And this pattern continues: the time evolution of \mathbf{p} depends on the heat flux, \mathbf{q} , the time evolution of which in turn depends on the *kurtosis*, sometimes denoted \mathbf{r} , and so on. The exact fluid description is thus really an infinite tower of coupled partial differential equations, together containing all of the complicated velocity-space information of the Boltzmann equation.

Needless to say, this is not very practical to work with. To reach a workable model, we need to stop at some floor of the tower and somehow express the unknown quantity at that floor in terms of the quantities we already know from the lower floors. Since this entails conjuring an additional equation to make the total set of equations closed, the additional equation in question is known as a *closure*, and the problem of how it should look in order to capture as much as possible of the neglected kinetic physics is perhaps *the* central problem in fluid modelling of plasmas.

And it is at this step that the collisions come to the rescue for collisional plasmas: while the collisional physics do yield some complications (e.g. friction terms) in the fluid equations at floors above the continuity equation, one of the main effects of all of the collisions is *thermalisation*, meaning collisional plasmas are constantly driven towards local thermal equilibrium. The physics of this state is both well-studied and well-understood: the distribution function’s velocity dependence simply becomes a normal distribution – commonly referred to as a “Maxwell-Boltzmann” or “Maxwell(ian)” distribution in plasma physics.

Specifically, it becomes a normal distribution centred on the flow velocity \mathbf{V} and with a width corresponding to the thermal speed $v_{\text{th}} = \sqrt{k_B T/m}$ of the particles involved ($k_B T$ being the thermal energy and m being the mass of the particles). Deviations from equilibrium can be assumed to be small and transient, and widely applicable closures can be derived [47–51] using perturbation theory.

Due to the lack of thermalising collisions, closures for collisionless plasmas are far less straightforward to derive – in fact, there is no known generally applicable framework for this case. The work included in this thesis is part of the currently ongoing efforts to address this issue. While closures have traditionally been derived by starting from kinetic theory – often in a linearised limit, ours are instead discovered through the use of machine learning methods. This approach, termed *data-driven* closure discovery, has seen a rapid growth in interest over the past few years, tracking the explosive growth of the field of machine learning at large. Most famous among these methods are probably those involving artificial neural networks [52–54], which e.g. underlie the recent upwelling of generative AI chatbots and virtual assistants. But while we do make use of neural networks to a limited extent, our focus is on a different machine learning framework, known as sparse regression, or SR [55–59].

1.3 Data-Driven Closure Discovery

Aiming to find collisionless heat flux closures (i.e. expressions for \mathbf{q} in terms of known lower-order quantities), we run kinetic simulations of a phenomenon of interest and calculate both \mathbf{q} and all lower-order fluid moments as well as the electromagnetic field at every spacetime point in the simulation. Using these, we then create a library of terms we think might be of interest for modelling \mathbf{q} .

In the simplest case [58–60], we assume the heat flux can be well approximated as a linear combination of the terms in the term library, and perform linear sparse regression to identify the coefficients which should be placed in front of all the terms to optimally fit the data. The term deemed “least important” according to some metric is then deleted, and a new round of linear regression is performed to find the optimal model without the deleted term. This process is then repeated, yielding a sequence of successively simpler models which approximately trace the so-called *Pareto front* between simplicity and accuracy.

For simplicity, we start by considering plasma phenomena which occur already in one-dimensional simulations, such as Landau damping of Langmuir waves and electron hole generation through two-stream instability – these two phenomena in particular are the ones featured in the appended papers. In both cases, SR consistently identifies the same six-term model for the (mass-normalised) electron heat flux. The coefficients for terms associated with spatial derivatives, which are free parameters, tend to correlate heavily with the instantaneous growth rate of heat flux perturbations, while the remaining coefficients vary more slowly depending on large-scale conditions. Interestingly, the three terms which do not contain spatial derivatives are generally far more

important for accurate q modelling, regularly accounting for more than 90 % of the variation in the heat flux on their own.

To actually make use of this closure in fluid simulations, the six free parameters need to be dynamically estimated from known fluid quantities – and as we show in the second of the two appended papers, this is at the very least possible to achieve for the three most important terms. To this end, both neural networks and a new framework for nonlinear sparse regression we have developed based on rational function fitting can be used. The latter approach yields models with higher interpretability, including e.g. analytically computable derivatives, as well as somewhat decreased exposure to overfitting – but at the cost of slightly lowered performance.

We additionally demonstrate that the six-term heat flux model can be generalised to multi-species modelling – particularly useful for e.g. two-stream instability simulations – in a straightforward manner. In future work, we hope to apply similar methodology to discover closures for phenomena requiring two- and three-dimensional simulation, such as pressure anisotropy-driven generation of whistler waves and magnetic reconnection, and analyse to what extent such closures can be viewed as generalisations of the six-term closure discussed above.

1.4 An Outline of this Thesis

The rest of this thesis is structured as follows: Chapter 2 gives an overview of the fundamentals of plasma physics modelling, from kinetic approaches to fluid models, and additionally describes a selection of closures and their respective ranges of validity. Furthermore, several plasma phenomena relevant to the subject are introduced and discussed, including those mentioned without explanation in this chapter, such as Landau damping and the fundamentals of plasma wave theory. In Chapter 3, the machine learning methods used in the appended articles – artificial neural networks and sparse regression – are outlined. Both more traditional linear SR methods and our newly developed nonlinear SR approach are discussed. In Chapter 4, the two previous chapters are then connected, outlining how machine learning methods were applied to closure discovery in the appended papers. Finally, we also give an outlook towards future avenues of investigation.

Chapter 2

Plasma Physics

2.1 First Principles: Kinetic Modelling

Reading the title of this section, the concerned theoretician might object that kinetic modelling of plasmas – i.e. modelling based on the use of phase-space distribution functions – is hardly the first principles of the physics involved: an *actual* first-principles analysis would start with the standard model of particle physics and general relativity. This is of course true in the strictest sense of the term. There are situations where taking quantum and/or general relativistic effects into account is necessary, such as when modelling the black hole accretion disks and neutron star pair-plasma coronas, as mentioned in Chapter 1. Closer to home, there is also the regime of extremely high-density plasma, which is reached in e.g. inertial confinement fusion and necessitates taking quantum degeneracy pressure into account. Nevertheless, for the space plasma applications this thesis is mainly concerned with, quantum effects can be safely neglected, gravity is weak, and the particles in the plasma travel at speeds which are at most weakly relativistic.

In such environments, treating the plasma as a collection of classical particles is the natural starting point. However, as we already briefly discussed, even this is far too computationally expensive to be viable for most plasma phenomena. Thus, one resorts to the statistical kinetic approach encoded in the distribution function $f(t, \mathbf{x}, \mathbf{v})$ and the Boltzmann equation written out in Equation (1.1). As long as we don't care about tracking individual particles, this view is actually *exact*, apart from any simplifications made with regards to the collisional term. In other words, for a perfectly collisionless plasma, where the collisional term is identically zero, the kinetic approach can for all intents and purposes be treated as first-principles. Without the collisional term, the Boltzmann equation is known as the Vlasov equation, and if we neglect gravity, the only important force is the electromagnetic Lorentz force. The set of partial differential equations we

need to solve in order to model our plasma is thus the Vlasov-Maxwell system:

$$\begin{cases} 0 = \partial_t f_\sigma + \mathbf{v} \cdot \nabla f_\sigma + \frac{q_\sigma}{m_\sigma} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_\sigma \\ \nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \\ \nabla \cdot \mathbf{B} = 0 \\ \nabla \times \mathbf{E} = -\partial_t \mathbf{B} \\ \nabla \times \mathbf{B} = c^{-2} \partial_t \mathbf{E} + \mu_0 \mathbf{J}, \end{cases} \quad (2.1)$$

Here, we use a separate distribution function f_σ for each particle species σ – meaning e.g. that the protons and electrons making up most of the solar wind would be described by distribution functions f_p and f_e , respectively. Furthermore, q_σ and m_σ refer to the charge and mass for particles of the species in question, and \mathbf{E} and \mathbf{B} denote the electric and magnetic field. The electric charge and current densities ρ and \mathbf{J} can be calculated from the distribution functions of the various species via

$$\begin{cases} \rho = \sum_{\sigma} q_{\sigma} \int d^3 \mathbf{v} f_{\sigma} \\ \mathbf{J} = \sum_{\sigma} q_{\sigma} \int d^3 \mathbf{v} \mathbf{v} f_{\sigma}. \end{cases} \quad (2.2)$$

For all of the small-to-intermediate-scale collisionless phenomena we will be considering, this is a *very* accurate model. We will get into what kinds of phenomena one can discern from these equations later on in Section 2.3, but first, let us take a closer look at how they can be used to derive collisionless fluid theory.

2.2 Fluid Models and the Moment Equations

As we discussed in Chapter 1, the kinetic approach to plasma modelling, while highly accurate, is also very computationally expensive. In large part, this is due to the fact that the distribution function has a seven-dimensional domain, comprising both time, physical space and velocity space. To resolve this, we can integrate the Vlasov equation over velocity space to derive a fluid theory:

$$0 = \underbrace{\int d^3 \mathbf{v} \partial_t f_\sigma}_{(0a)} + \underbrace{\int d^3 \mathbf{v} \mathbf{v} \cdot \nabla f_\sigma}_{(0b)} + \underbrace{\frac{q_\sigma}{m_\sigma} \int d^3 \mathbf{v} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_\sigma}_{(0c)}. \quad (2.3)$$

Since t and \mathbf{x} are independent of \mathbf{v} when viewed as coordinates in $(t, \mathbf{x}, \mathbf{v})$ phase space, terms (0a) and (0b) can be simplified by moving the differentiations with respect to t and \mathbf{x} outside of the integrals:

$$(0a) = \partial_t \int d^3 \mathbf{v} f_\sigma = \partial_t n_\sigma, \quad (0b) = \nabla \cdot \left(\int d^3 \mathbf{v} \mathbf{v} f_\sigma \right) = \nabla \cdot (n_\sigma \mathbf{V}_\sigma). \quad (2.4)$$

Here, we have made the usual physicist's assumption of all functions being sufficiently continuous (\mathcal{C}^∞ unless otherwise noted) for integration and differentiation to be freely reorderable. Additionally, we have introduced the particle density n_σ and particle flux $n_\sigma \mathbf{V}_\sigma$. These two quantities are the zeroth- and first-order *raw moments* of the distribution function – a term which might be familiar from statistics, where one uses similar methodology to extract information from probability distributions. In this analogy, the density corresponds to the total probability (in statistics usually normalised to 1), and the particle flux – or alternatively the flow velocity \mathbf{V}_σ – corresponds to the mean of the distribution. With these definitions, the expressions in Equation (2.2) can be simplified into $\rho = \sum_\sigma q_\sigma n_\sigma$ and $\mathbf{J} = \sum_\sigma q_\sigma n_\sigma \mathbf{V}_\sigma$, underscoring their respective interpretations as charge and current densities.

For term (0c), we need to perform partial integration, making use of the fact that

$$\frac{m_\sigma}{q_\sigma}(0c) = \int d^3\mathbf{v} \nabla_{\mathbf{v}} \cdot (f_\sigma \mathbf{E} + \mathbf{v} f_\sigma \times \mathbf{B}) - \int d^3\mathbf{v} f_\sigma \nabla_{\mathbf{v}} \cdot (\mathbf{E} + \mathbf{v} \times \mathbf{B}). \quad (2.5)$$

By the divergence theorem, the first of these expressions can be exchanged for a surface integral over the boundary of the integration domain. But since we are integrating over all of velocity space, this boundary is actually at $|\mathbf{v}| \rightarrow \infty$, and in this limit f_σ must go to zero. In fact, for all cases of interest to us, f_σ will go to zero fast enough that $|\mathbf{v}|^k f_\sigma \rightarrow 0$ for all k . Thus, the first term on the right-hand side of Equation (2.5) will evaluate to zero. As it turns out, this is actually also true for the second term: Firstly, the \mathbf{E} and \mathbf{B} fields clearly do not depend on \mathbf{v} . Secondly, the divergence of \mathbf{v} cross-multiplied by any \mathbf{v} -independent vector is identically zero. Provided one is familiar with index notation, this can be seen clearly by writing out the indices of the summation:

$$\nabla_{\mathbf{v}} \cdot (\mathbf{v} \times \mathbf{B}) = \varepsilon_{ijk} \partial_{v_k} v_i B_j = \varepsilon_{ijk} \delta_{ik} B_j = 0, \quad (2.6)$$

since the Kronecker delta δ_{ij} is symmetric, while the Levi-Civita tensor ε_{ijk} is fully antisymmetric. For conciseness, we are using the Einstein summation convention, where repeated indices are implicitly summed over. In other words, term (0c) vanishes entirely in the fluid equation we have derived, and we are left with the continuity equation,

$$\partial_t n_\sigma + \nabla \cdot (n_\sigma \mathbf{V}_\sigma) = 0. \quad (2.7)$$

Note that Maxwell's equations are not affected by this procedure – after all, they depend on integrals of f_σ over velocity space (specifically on n_σ and $n_\sigma \mathbf{V}_\sigma$) even in the kinetic case, through ρ and \mathbf{J} .

2.2.1 The Momentum Equation

This is not the end of the story, however, as we know from Chapter 1. The continuity equation tells us how the zeroth moment n_σ evolves only in terms of the divergence of the first moment $n_\sigma \mathbf{V}_\sigma$, for which the time evolution is still unknown. To derive the governing equation for its evolution, we must

multiply the Vlasov equation by \mathbf{v} and then integrate over velocity space again, effectively taking the first moment of the entire equation:

$$0 = \underbrace{\int d^3\mathbf{v} \mathbf{v} \partial_t f_\sigma}_{(1a)} + \underbrace{\int d^3\mathbf{v} \mathbf{v} \mathbf{v} \cdot \nabla f_\sigma}_{(1b)} + \underbrace{\frac{q_\sigma}{m_\sigma} \int d^3\mathbf{v} \mathbf{v} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_\sigma}_{(1c)}. \quad (2.8)$$

For terms (1a) and (1b), we can use exactly the same line of reasoning as for the “zeroth moment” (i.e. the continuity equation), giving us

$$(1a) = \partial_t \int d^3\mathbf{v} \mathbf{v} f_\sigma = \partial_t (n_\sigma \mathbf{V}_\sigma), \quad (1b) = \nabla \cdot \left(\int d^3\mathbf{v} \mathbf{v} \mathbf{v} f_\sigma \right) = \nabla \cdot \mathbf{P}_\sigma. \quad (2.9)$$

Here, we end up defining a new kind of mathematical object: the symmetric second-order *stress-energy tensor* \mathbf{P}_σ . As we have done here, we will make use of so-called *dyadic notation* going forward, where two vectors written next to each other are implicitly tensorially multiplied. In other words, $\mathbf{ab} = \mathbf{a} \otimes \mathbf{b}$ is a second-order tensor with elements $[\mathbf{ab}]_{ij} = a_i b_j$. As for term (1c), we must again perform partial integration, and the boundary term vanishes just like before. The other term, however, does not, instead yielding

$$\begin{aligned} \frac{m_\sigma}{q_\sigma} (1c) &= - \int d^3\mathbf{v} f_\sigma (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} \mathbf{v} = \{\mathbf{a} \cdot \nabla_{\mathbf{v}} \mathbf{v} = \mathbf{a}\} = \\ &= - \int d^3\mathbf{v} f_\sigma (\mathbf{E} + \mathbf{v} \times \mathbf{B}) = -n_\sigma (\mathbf{E} + \mathbf{V}_\sigma \times \mathbf{B}), \end{aligned} \quad (2.10)$$

where we have used the fact that the Lorentz force $\mathbf{E} + \mathbf{v} \times \mathbf{B}$ has no \mathbf{v} -divergence, as we discussed previously. In other words, the first-order moment equation – the *momentum equation* – looks like

$$\partial_t (n_\sigma \mathbf{V}_\sigma) + \nabla \cdot \mathbf{P}_\sigma = \frac{q_\sigma}{m_\sigma} n_\sigma (\mathbf{E} + \mathbf{V}_\sigma \times \mathbf{B}). \quad (2.11)$$

It should be noted that the stress-energy tensor \mathbf{P}_σ introduced here is really the *mass-normalised* stress-energy tensor. The more standard definition, which has the expected units of energy density, is equal to our \mathbf{P}_σ multiplied by the particle mass. For convenience, we will use similar mass-normalised definitions for several other fluid quantities introduced in this chapter as well, including the pressure and temperature tensors.

These two tensors in particular are related to the stress-energy tensor in a straightforward manner: the pressure tensor \mathbf{p} is the second *central* moment, while \mathbf{P} is the second *raw* moment. In other words, the pressure is the part of the stress-energy tensor which is not from overall plasma convection, i.e. the part which is nonzero even in the plasma’s local centre-of-momentum frame. Mathematically, we can calculate it through

$$\begin{aligned} \mathbf{p} &= \int d^3\mathbf{v} (\mathbf{v} - \mathbf{V}_\sigma)^{(2)} f_\sigma = \int d^3\mathbf{v} \left(\mathbf{v}^{(2)} - 2\{\mathbf{v} \mathbf{V}_\sigma\} + \mathbf{V}_\sigma^{(2)} \right) f_\sigma = \\ &= \mathbf{P}_\sigma - n_\sigma \mathbf{V}_\sigma^{(2)}. \end{aligned} \quad (2.12)$$

Here, we have introduced the shorthand notation $\mathbf{a}^{(n)}$ for n times repeated tensor multiplication, so that e.g. $\mathbf{v}^{(3)} = \mathbf{v}\mathbf{v}\mathbf{v}$. We have also introduced the notation $\{\cdot\}$ for symmetrisation. This operation returns the “fully symmetric version” of a given tensor – or, more specifically, the average of all tensor permutations, so that e.g. $\{\mathbf{a}\mathbf{P}\}_{ijk} = \frac{1}{3}(a_i P_{jk} + a_j P_{ik} + a_k P_{ij})$ with \mathbf{P} symmetric. From the pressure tensor \mathbf{p}_σ we can then derive the temperature tensor \mathbf{T}_σ for each species by normalising with respect to density: $\mathbf{T}_\sigma = \frac{1}{n_\sigma}\mathbf{p}_\sigma$. This mass-normalised definition of \mathbf{T}_σ also absorbs the k_B prefactor which would normally be present, as is common practice in plasma physics. Thus, this tensor really corresponds to (mass-normalised) thermal energy.

It is sometimes convenient to rewrite Equation (2.11) in terms of the pressure instead of the stress-energy tensor. Inserting Equation (2.12), the divergence term in Equation (2.11) can be expanded into

$$\nabla \cdot \mathbf{P}_\sigma = \nabla \cdot \mathbf{p}_\sigma + [\nabla \cdot (n_\sigma \mathbf{V}_\sigma) + n_\sigma \mathbf{V}_\sigma \cdot \nabla] \mathbf{V}_\sigma. \quad (2.13)$$

Expanding the time derivative term as well, we can cancel out a copy of the continuity equation multiplied by \mathbf{V}_σ , yielding the “centralised” version of the equation:

$$n_\sigma D_t \mathbf{V}_\sigma + \nabla \cdot \mathbf{p}_\sigma = \frac{q_\sigma}{m_\sigma} n_\sigma (\mathbf{E} + \mathbf{V}_\sigma \times \mathbf{B}). \quad (2.14)$$

Here, we have further introduced the convective derivative $D_t = \partial_t + \mathbf{V}_\sigma \cdot \nabla$.

2.2.2 The Pressure Equation

Of course, we now have the same problem as before: we only know how to evolve $n_\sigma \mathbf{V}_\sigma$ or \mathbf{V}_σ given the divergence of \mathbf{P}_σ or \mathbf{p}_σ , and the time evolution of these quantities is unknown. To find the differential equation governing their behaviour, we need to take the second moment of the Vlasov equation:

$$0 = \underbrace{\int d^3\mathbf{v} \left[\mathbf{v}^{(2)} \partial_t + \mathbf{v}^{(3)} \cdot \nabla \right] f_\sigma}_{(2ab)} + \underbrace{\frac{q_\sigma}{m_\sigma} \int d^3\mathbf{v} \mathbf{v}^{(2)} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_\sigma}_{(2c)}. \quad (2.15)$$

By the same line of reasoning we used when deriving the continuity and momentum equations, term (2ab) can be rewritten as $\partial_t \mathbf{P}_\sigma + \nabla \cdot \mathbf{Q}_\sigma$, where $\mathbf{Q}_\sigma = \int d^3\mathbf{v} \mathbf{v}^{(3)} f_\sigma$ is the third-order moment of f_σ , known as the *energy flux* tensor (sometimes also called the *energy flow* or *stress flow* tensor). Furthermore, partial integration of the remaining term gives

$$\begin{aligned} (2c) &= -\frac{q_\sigma}{m_\sigma} \int d^3\mathbf{v} \left[(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} \mathbf{v}^{(2)} \right] f_\sigma = \\ &= \left\{ \mathbf{a} \cdot \nabla_{\mathbf{v}} \mathbf{v}^{(2)} = 2\{\mathbf{a}\mathbf{v}\} \right\} = \\ &= -2\frac{q_\sigma}{m_\sigma} (n_\sigma \{\mathbf{V}_\sigma \mathbf{E}\} + \{\mathbf{P}_\sigma \times \mathbf{B}\}), \end{aligned} \quad (2.16)$$

where $\mathbf{P}_\sigma \times \mathbf{B}$ should be interpreted as the second-order tensor with indices $[\mathbf{P}_\sigma \times \mathbf{B}]_{ij} = \varepsilon_{jkl} P_{\sigma ik} B_l$. In other words, the complete pressure equation –

perhaps more accurately referred to as the stress-energy equation when written in terms of \mathbf{P}_σ and \mathbf{Q}_σ like this – looks like

$$\partial_t \mathbf{P}_\sigma + \nabla \cdot \mathbf{Q}_\sigma = 2 \frac{q_\sigma}{m_\sigma} (n_\sigma \{ \mathbf{V}_\sigma \mathbf{E} \} + \{ \mathbf{P}_\sigma \times \mathbf{B} \}). \quad (2.17)$$

Just like for the momentum equation, it is sometimes more convenient to work in terms of the central moments \mathbf{p}_σ and $\mathbf{q}_\sigma = \int d^3 \mathbf{v} (\mathbf{v} - \mathbf{V}_\sigma)^{(3)} f_\sigma$. This latter quantity is the *heat flux* tensor, which will feature heavily in the discussion of collisionless plasma closures below. Defining $\mathbf{v}' = \mathbf{v} - \mathbf{V}_\sigma$, we can expand

$$\mathbf{Q}_\sigma = \int d^3 \mathbf{v} (\mathbf{v}' + \mathbf{V}_\sigma)^{(3)} f_\sigma = \mathbf{q}_\sigma + 3 \{ \mathbf{V}_\sigma \mathbf{p}_\sigma \} + n_\sigma \mathbf{V}_\sigma^{(3)}, \quad (2.18)$$

where we have used the fact that $\int d^3 \mathbf{v} \mathbf{v}' f_\sigma = 0$. Inserting this expression into Equation (2.17) together with the previously derived expression for \mathbf{P}_σ in terms of \mathbf{p}_σ , we eventually reach the centralised *pressure equation*,

$$\partial_t \mathbf{p}_\sigma + \nabla \cdot (\mathbf{V}_\sigma \mathbf{p}_\sigma) + 2 \{ \mathbf{p}_\sigma \cdot \nabla \mathbf{V}_\sigma \} + \nabla \cdot \mathbf{q}_\sigma = 2 \frac{q_\sigma}{m_\sigma} \{ \mathbf{p}_\sigma \times \mathbf{B} \}, \quad (2.19)$$

after subtracting away one copy of the continuity equation multiplied by $\mathbf{V}_\sigma^{(2)}$ and two copies of the momentum equation multiplied by \mathbf{V}_σ along the way.

One could continue on further, deriving the time evolution of \mathbf{Q}_σ and \mathbf{q}_σ and so on, but for the purposes of this thesis, we will stop at the third moment equation. The resulting model is (somewhat confusingly) often referred to as the *10-moment model*, since n_σ , \mathbf{V}_σ and \mathbf{p}_σ together contain ten independent degrees of freedom. Collecting the three moment equations we have derived in one place for convenience, we can see that we also have ten scalar constraints:

$$\begin{cases} \partial_t n_\sigma + \nabla \cdot (n_\sigma \mathbf{V}_\sigma) = 0 \\ n_\sigma D_t \mathbf{V}_\sigma + \nabla \cdot \mathbf{p}_\sigma = \frac{q_\sigma}{m_\sigma} n_\sigma (\mathbf{E} + \mathbf{V}_\sigma \times \mathbf{B}) \\ \partial_t \mathbf{p}_\sigma + \nabla \cdot (\mathbf{V}_\sigma \mathbf{p}_\sigma) + 2 \{ \mathbf{p}_\sigma \cdot \nabla \mathbf{V}_\sigma \} + \nabla \cdot \mathbf{q}_\sigma = 2 \frac{q_\sigma}{m_\sigma} \{ \mathbf{p}_\sigma \times \mathbf{B} \}. \end{cases} \quad (2.20)$$

Unfortunately, these constraints are only helpful insofar as the heat flux \mathbf{q}_σ is known, which is often not the case, as we briefly discussed in Chapter 1. We will go into more of the details of how this problem – the closure problem – can be resolved in Section 2.4 below. But first, let us quickly go through some of the interesting phenomena which arise through the physics encoded in these equations.

2.3 Plasma Phenomena: The Necessities

As we have already alluded to, plasmas exhibit a very wide range of complex phenomena, not all of which fit neatly within the bounds of this thesis. For example, particle transport and drift will be mostly ignored, and turbulence will not be discussed in any significant detail either. Instead, we will focus on the phenomena which are more directly related to the aims of and motivation behind the appended papers – namely: waves, instabilities and Landau damping.

2.3.1 Waves in Plasmas

Much of the interesting collective behaviour which distinguishes plasmas from gases can be understood by looking at what kinds of oscillations are supported – i.e. what kinds of waves we might expect to see if we perturb the plasma in some way. One kind of wave we expect to see some version of is of course light (or, more generally, electromagnetic radiation). Though there is clearly some variance in how well such waves can propagate through a plasma depending on wave and plasma parameters, seeing as radio waves at frequencies used for communication are reflected back to Earth by the ionosphere, whereas visible light can pass through it without issue.

Additionally, we might expect to see oscillations of the electrons and ions making up the plasma, since the Coulomb attraction between the two kinds of particles acts as a restoring force if the particles are perturbed from equilibrium. As this is perhaps *the* defining characteristic separating a plasma from a neutral gas, this type of oscillation is often referred to as simply the “plasma oscillation” or “plasma wave”. In this thesis, however, we will not use the latter of these two terms, since it risks confusion with all of the other wave modes occurring in plasmas, though we will sometimes refer to them by their alternative name of *Langmuir waves*.

To see that both of these kinds of waves indeed occur, let us start by considering a plasma which is cold (zero-temperature to begin with) and homogeneous, filling all of space. We can then examine what happens if we make a slight perturbation δX to the background value \bar{X} of every quantity X . Assuming the background electromagnetic field is zero, we can derive the dispersion relation for the supported plasma waves by linearising Maxwell’s equations around the background conditions (with $\bar{\rho}$, $\bar{\mathbf{J}}$, $\bar{\mathbf{E}}$ and $\bar{\mathbf{B}}$ all zero):

$$\begin{cases} \nabla \cdot \delta \mathbf{E} = \frac{\delta \rho}{\epsilon_0} \\ \nabla \cdot \delta \mathbf{B} = 0 \\ \nabla \times \delta \mathbf{E} = -\partial_t \delta \mathbf{B} \\ \nabla \times \delta \mathbf{B} = c^{-2} \partial_t \delta \mathbf{E} + \mu_0 \delta \mathbf{J}. \end{cases} \quad (2.21)$$

Making a wave ansatz, we assume $\delta X = \tilde{X} e^{-i(\omega t - \mathbf{k} \cdot \mathbf{x})}$ for each quantity X , with \tilde{X} being some amplitude which is allowed to be complex to account for phase differences between the oscillations in the various quantities. Here, the quantities ω and \mathbf{k} are the (angular) frequency and wave vector of the oscillation. The ∇ operator when applied to a quantity δX then effectively reduces to $i\mathbf{k}$, and ∂_t reduces to $-i\omega$. Making a wave ansatz is thus equivalent to performing a Fourier transform. Comparing the curl of the third equation to the (negative) time derivative of the fourth, we find

$$-\mathbf{k} \times (\mathbf{k} \times \tilde{\mathbf{E}}) = \omega^2 c^{-2} \tilde{\mathbf{E}} + i\omega \mu_0 \tilde{\mathbf{J}}, \quad (2.22)$$

the Fourier-space Helmholtz equation.

Assuming the perturbation is small, we can linearise in \tilde{X} for all quantities, and in such a situation, there must be a *linear* relation between any two

perturbation amplitudes \tilde{X} and \tilde{Y} . In other words, to first order in the perturbation amplitudes, $\tilde{\mathbf{J}}$ and $\tilde{\mathbf{E}}$ must be related by some tensor conductivity $\boldsymbol{\sigma}$ according to $\tilde{\mathbf{J}} = \boldsymbol{\sigma} \cdot \tilde{\mathbf{E}}$. If we additionally define the vector index of refraction $\mathbf{n} = \mathbf{k}c/\omega$, we can rewrite this as

$$-\mathbf{n} \times (\mathbf{n} \times \tilde{\mathbf{E}}) = \left(\mathbf{I} + \frac{i\boldsymbol{\sigma}}{\varepsilon_0\omega} \right) \cdot \tilde{\mathbf{E}} = \mathbf{K} \cdot \tilde{\mathbf{E}}, \quad (2.23)$$

where \mathbf{I} is the identity tensor and $\mathbf{K} = \mathbf{I} + i\boldsymbol{\sigma}/\varepsilon_0\omega$ is the dielectric tensor. To find $\boldsymbol{\sigma}$, we must refer to the momentum equation, since $\delta\mathbf{J} = \mathbf{J} = \sum_{\sigma} q_{\sigma} n_{\sigma} \mathbf{V}_{\sigma}$:

$$\partial_t \mathbf{J} = \sum_{\sigma} q_{\sigma} \partial_t (n_{\sigma} \mathbf{V}_{\sigma}) = \sum_{\sigma} q_{\sigma} \left[-\nabla \cdot \mathbf{P}_{\sigma} + \frac{q_{\sigma}}{m_{\sigma}} n_{\sigma} (\mathbf{E} + \mathbf{V}_{\sigma} \times \mathbf{B}) \right]. \quad (2.24)$$

With our wave ansatz, which does not include a background magnetic field (meaning $\mathbf{B} = 0$), nor background particle streaming (meaning $\tilde{\mathbf{V}}_{\sigma} = 0$ for all species), this reduces to

$$-i\omega \tilde{\mathbf{J}} = \sum_{\sigma} q_{\sigma} \left[-i\mathbf{k} \cdot \tilde{\mathbf{p}}_{\sigma} + \frac{q_{\sigma}}{m_{\sigma}} \bar{n}_{\sigma} \tilde{\mathbf{E}} \right] \quad (2.25)$$

to first order in perturbation amplitudes. In the zero-temperature limit, \mathbf{p}_{σ} (and thus $\tilde{\mathbf{p}}_{\sigma}$) goes to zero for all species, meaning we are left with only

$$\tilde{\mathbf{J}} = \frac{i}{\omega} \sum_{\sigma} \frac{q_{\sigma}^2 \bar{n}_{\sigma}}{m_{\sigma}} \tilde{\mathbf{E}}. \quad (2.26)$$

In other words, the conductivity in this case is simply a scalar multiple of the identity tensor (or equivalently, just a scalar). The dispersion relation in Equation (2.23) thus simplifies greatly, into

$$-\mathbf{n} \times (\mathbf{n} \times \tilde{\mathbf{E}}) = \left(1 - \frac{1}{\omega^2} \sum_{\sigma} \frac{q_{\sigma}^2 \bar{n}_{\sigma}}{m_{\sigma} \varepsilon_0} \right) \tilde{\mathbf{E}} = K \tilde{\mathbf{E}}. \quad (2.27)$$

Without loss of generality, we can take $\hat{\mathbf{z}}$ to be parallel to \mathbf{n} , giving

$$-|\mathbf{n}|^2 \hat{\mathbf{z}} \times (\hat{\mathbf{z}} \times \tilde{\mathbf{E}}) = |\mathbf{n}|^2 \begin{bmatrix} \tilde{E}_x \\ \tilde{E}_y \\ 0 \end{bmatrix} = K \begin{bmatrix} \tilde{E}_x \\ \tilde{E}_y \\ \tilde{E}_z \end{bmatrix}. \quad (2.28)$$

This dispersion relation has two types of solutions:

$$\begin{cases} K = |\mathbf{n}|^2 = \frac{k^2 c^2}{\omega^2}, & \tilde{E}_z = 0 \\ K = 0, & \tilde{E}_x = \tilde{E}_y = 0. \end{cases} \quad (2.29)$$

Solving the dispersion relation of the first solution type for ω^2 , we find

$$\omega^2 = \omega_p^2 + k^2 c^2, \quad \omega_p = \sqrt{\sum_{\sigma} \frac{q_{\sigma}^2 \bar{n}_{\sigma}}{m_{\sigma} \varepsilon_0}}. \quad (2.30)$$

The mode described by this equation, known as the *ordinary mode*, is the unmagnetised plasma analogue of the electromagnetic radiation we are used to. To see this, we need only consider what happens in the zero-density limit (i.e. what happens if we remove the plasma). This corresponds to taking $\omega_p \rightarrow 0$, causing both the phase velocity $v_{\text{ph}} = \omega/k$ and the group velocity $v_g = \partial_k \omega$ to go to c , the speed of light. We can also see that two independent polarisations are supported (e.g. along $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$), since $\tilde{\mathbf{E}}$ can be polarised in any direction perpendicular to the direction of propagation for this type of solution. Unlike regular light, however, unmagnetised plasmas only allow waves of this type to propagate if they have frequencies $\omega > \omega_p$, as can be seen in Figure 2.1. This phenomenon, known as a *mode cutoff*, is the reason underlying why the ionosphere reflects some radio waves back to Earth while allowing visible light to pass through practically unaffected: $\omega_{\text{visible}} \gg \omega_p$, while $\omega_{\text{radio}} \lesssim \omega_{\text{pe}}$.

As for the second type of solution, it is clearly a longitudinal mode (i.e. with $\tilde{\mathbf{E}} \parallel \hat{\mathbf{z}}$), which has no analogue in vacuum – and this we can actually see from the dispersion relation, since solving for ω^2 nets us

$$\omega^2 = \omega_p^2, \quad (2.31)$$

which has no nontrivial solution in the $\omega_p \rightarrow 0$ limit. Calculating the phase and group velocity, we find $v_{\text{ph}} = \omega_p/k$ and $v_g = 0$, meaning the mode does not propagate.




Figure 2.1: The unmagnetised zero-temperature plasma dispersion relation, with the transverse, light-like ordinary mode in solid blue and the non-propagating plasma oscillation in dashed red. For comparison, the dispersion relation of electromagnetic radiation propagating through a vacuum is shown in dotted black.

But this *is* in agreement with what we would expect from the Coulomb-attraction-driven oscillation we anticipated finding, and it can be shown that this is indeed the mechanism whereby this mode is generated. As the mode in

question is commonly known simply as the plasma oscillation, its characteristic frequency ω_p is known as the *plasma frequency*. It is sometimes convenient to define separate plasma frequencies $\omega_{p\sigma} = q_\sigma \sqrt{\bar{n}_\sigma / m_\sigma \epsilon_0}$ for each species, in which case the overall plasma frequency can be computed via $\omega_p^2 = \sum_\sigma \omega_{p\sigma}^2$. Note however that since $\omega_{p\sigma} \sim 1/\sqrt{m_\sigma}$, the contribution from the lightest particle species dominates. Thus, ω_p is well approximated by the electron plasma frequency ω_{pe} for most applications.

One important plasma phenomenon we have missed in this analysis is *magnetisation*: the behaviour of a plasma changes significantly when an external magnetic field is applied (i.e. when $\bar{\mathbf{B}} \neq 0$). The underlying reason for this is that charged particles moving in a magnetic field will start to curl around the field lines due to the $\mathbf{v} \times \mathbf{B}$ term in the Lorentz force. This effectively restricts their motion to be largely along the field lines (neglecting the various sources of drift, such as field line curvature), while their motion in the plane perpendicular to the field lines is approximately circular, as illustrated in Figure 2.2.




Figure 2.2: The cyclotron motion of electrons (e^- , blue) and protons (p^+ , red) in a magnetised plasma. As the particles have opposite charge, their trajectories spiral in opposite directions.

The “orbital frequency” around the field lines is given by $\omega_{c\sigma} = q_\sigma \bar{B} / m_\sigma$, and is known as the *cyclotron frequency*. This is another important plasma parameter, and shows up not only in the single-particle view, but also in the dispersion relation for waves in a magnetised plasma. Note that with this definition, the electron and ion cyclotron frequencies will have different sign (the electron frequency being negative due to its negative charge), with $|\omega_{ce}| \gg |\omega_{ci}|$ due to the mass difference between electrons and any relevant ion species i .

In the dispersion relation for a magnetised plasma, these frequencies show up as resonances, giving rise to electron and ion cyclotron waves. Additionally, wave propagation both parallel and perpendicular to the background magnetic field is significantly affected. This gives rise to further wave modes, such as lower and upper hybrid waves, as well as so-called whistler waves, with behaviour varying strongly depending on the angle between \mathbf{k} and $\bar{\mathbf{B}}$. As

the details of waves in magnetised plasmas lie somewhat beyond the scope of this thesis, we will not go through the details of these modes here. The interested reader is referred to textbooks such as Stix's *Waves in Plasmas* [61] or Swanson's *Plasma Waves* [62] for a more in-depth treatment.

All of this discussion has been based on the zero-temperature limit, which is of course not completely faithful to an actual plasma, magnetised or otherwise. For example, a type of wave which we not seen so far is some version of sound waves, familiar from neutral gases. These waves do indeed exist in plasmas – and separately for ions and electrons, at that – but require nonzero temperature to be able to propagate. After all, they are fundamentally oscillations in pressure, travelling at a “speed of sound” proportional to the thermal speed of the species involved, which can be written as simply $v_{\text{th},\sigma} = \sqrt{T_\sigma}$ with our mass-normalised definition of T_σ . Due to the charged nature of the constituent particles, the plasma analogue of sound waves additionally couple to the electromagnetic field, giving rise to magnetised ion and electron sound waves at frequencies below the cyclotron frequencies of the respective species, with resonances at $|\omega_{c\sigma}|$. At higher frequencies, the sound waves are unmagnetised, more similar to the familiar sound waves in a neutral gas.

Apart from giving rise to these new wave modes, a nonzero temperature also modifies the dispersion relations for the previously discussed ones. Most relevant to our purposes, the Langmuir wave dispersion relation changes from $\omega^2 = \omega_p^2$ to approximately

$$\omega^2 = \omega_p^2 + 3 \frac{k^2 v_{\text{th},e}^2}{\omega^2} + \mathcal{O}\left(\frac{k^4 v_{\text{th},e}^4}{\omega^4}\right) \quad (2.32)$$

if electrons are allowed to have a nonzero temperature. This allows the mode to propagate, since the group velocity is no longer identically zero. Being the first terms of a Maclaurin expansion in $kv_{\text{th},e}/\omega$, the expression on the right-hand side of this dispersion relation is only useful at low k , however. To more fully understand waves in plasmas (and Langmuir waves in particular), we must also take into account wave-particle interactions. Specifically, a wave passing through a section of plasma may either deposit energy into the particles of the plasma, being damped in the process, or absorb energy and intensify, leading to an instability. To understand these processes, let us go back to the kinetic picture and see what happens to a wave passing through a plasma with non-negligible temperature.

2.3.2 Landau Damping and Instabilities

For simplicity, and because it is the case most relevant to this thesis, let us restrict ourselves to electrostatic waves in an unmagnetised (but nonzero-temperature) plasma. Largely following Section 4.2.2 of Swanson [62], $\bar{\rho}$, $\bar{\mathbf{J}}$, $\bar{\mathbf{E}}$ and $\bar{\mathbf{B}}$ are zero like before, but we additionally have $\delta\mathbf{B} = 0$ and $\delta\mathbf{E} = -\nabla\delta\phi$, with $\delta\phi$ being the perturbation in the electrostatic potential. The Vlasov-

Maxwell system thus effectively reduces to

$$\begin{cases} 0 = \partial_t \delta f_\sigma + \mathbf{v} \cdot \nabla \delta f_\sigma - \frac{q_\sigma}{m_\sigma} \nabla \delta \phi \cdot \nabla_{\mathbf{v}} (\bar{f}_\sigma + \delta f_\sigma) \\ -\nabla^2 \delta \phi = \frac{\delta \rho}{\varepsilon_0} = \varepsilon_0^{-1} \sum_\sigma q_\sigma \int d^3 \mathbf{v} \delta f_\sigma \end{cases} \quad (2.33)$$

Making a wave ansatz only in space to start with, i.e. taking $\delta X = \tilde{X}(t)e^{i\mathbf{k} \cdot \mathbf{x}}$, and neglecting higher-order terms, this further reduces to

$$\begin{cases} 0 = \partial_t \tilde{f}_\sigma + ikv_z \tilde{f}_\sigma - \frac{q_\sigma}{m_\sigma} ik \tilde{\phi} \partial_{v_z} \bar{f}_\sigma \\ \tilde{\phi} = \frac{1}{\varepsilon_0 k^2} \sum_\sigma q_\sigma \int d^3 \mathbf{v} \tilde{f}_\sigma, \end{cases} \quad (2.34)$$

where we like before have taken $\mathbf{k} \parallel \hat{\mathbf{z}}$ without loss of generality. Instead of making a wave ansatz (or Fourier-transforming) in time as well, let us perform a Laplace transform in order to take into account initial conditions. This transform, which may be considered a generalisation of the Fourier transform adapted for decaying and growing phenomena, is defined by

$$\begin{cases} \mathcal{L}[X(t)] := \hat{X}(s) = \int_0^\infty dt e^{-st} X(t) \\ \mathcal{L}^{-1}[\hat{X}(s)] = X(t) = \int_{r-i\infty}^{r+i\infty} \frac{ds}{2\pi i} e^{st} \hat{X}(s), \end{cases} \quad (2.35)$$

where the positive real number r is chosen so that $r > \text{Re } p_n$ for all poles p_n of $\hat{X}(s)$. With this definition, one can further derive that

$$\mathcal{L}[\partial_t X] = s\hat{X} + [X e^{-st}]_{t=0}^{t \rightarrow \infty} = \{\text{Re } s > \text{Re } p_n \forall n\} = s\hat{X} - X_0, \quad (2.36)$$

where we have introduced to notation $X_0 = X|_{t=0}$. Our two equations then become

$$\begin{cases} \tilde{f}_{\sigma,0} = s\hat{f}_\sigma + ikv_z \hat{f}_\sigma - \frac{q_\sigma}{m_\sigma} ik \hat{\phi} \partial_{v_z} \bar{f}_\sigma \\ \hat{\phi} = \frac{1}{\varepsilon_0 k^2} \sum_\sigma q_\sigma \int d^3 \mathbf{v} \hat{f}_\sigma, \end{cases} \quad (2.37)$$

where we are omitting the tildes on Laplace-transformed quantities to avoid clutter. Solving the first equation for \hat{f}_σ , we get

$$\hat{f}_\sigma = \frac{\tilde{f}_{\sigma,0} + \frac{ikq_\sigma}{m_\sigma} \hat{\phi} \partial_{v_z} \bar{f}_\sigma}{s + ikv_z}, \quad (2.38)$$

and inserting this into the second equation, we find

$$\begin{aligned} \hat{\phi} &= \frac{1}{\varepsilon_0 k^2} \sum_\sigma q_\sigma \int d^3 \mathbf{v} \frac{\tilde{f}_{\sigma,0} + \frac{ikq_\sigma}{m_\sigma} \hat{\phi} \partial_{v_z} \bar{f}_\sigma}{s + ikv_z} = \\ &= \frac{1}{\varepsilon_0 k^2} \sum_\sigma q_\sigma \left[\int d^3 \mathbf{v} \frac{\tilde{f}_{\sigma,0}}{s + ikv_z} + \frac{ikq_\sigma}{m_\sigma} \hat{\phi} \int d^3 \mathbf{v} \frac{\partial_{v_z} \bar{f}_\sigma}{s + ikv_z} \right]. \end{aligned} \quad (2.39)$$

Performing the integrations over v_x and v_y and using the notation $F(v_z) = \int d^2\mathbf{v}_\perp f(\mathbf{v})$ for the distribution function over v_z , we can solve for $\hat{\phi}$ to get

$$\hat{\phi} = \frac{\sum_\sigma \frac{q_\sigma}{\varepsilon_0 k^2} \int dv_z \frac{\tilde{F}_{\sigma,0}}{s + ikv_z}}{1 - \sum_\sigma \frac{iq_\sigma^2}{km_\sigma \varepsilon_0} \int dv_z \frac{\tilde{F}_\sigma}{s + ikv_z}}. \quad (2.40)$$

The question, then, is how this should be interpreted. By definition, we can return to t -space via

$$\tilde{\phi} = \int_{r-i\infty}^{r+i\infty} \frac{ds e^{st}}{2\pi i} \hat{\phi}. \quad (2.41)$$

Computing this contour integral directly is nontrivial. However, the integrand is analytic along the path of integration (by the definition of r) – and, indeed, analytic for all s such that $\text{Re } s \geq r$. Furthermore, the function is analytic almost everywhere on its domain for any reasonable \tilde{f}_σ and $\tilde{f}_{\sigma,0}$, and can be analytically extended even to s such that $\text{Re } s < 0$, where the definition of $\tilde{X}(s)$ given in Equation (2.35) may not converge.

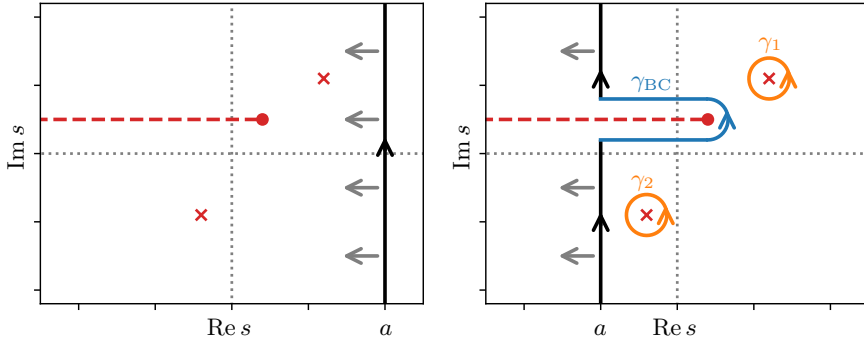


Figure 2.3: The contour deformation used in the derivation of Equation (2.42). As $a \rightarrow -\infty$, the black contour moves leftwards and is deformed to avoid the singularities shown in red. A detour γ_{BC} , shown in blue, is taken around the branch cut, and loops γ_n , shown in orange, are left behind encircling each pole as the main integration contour moves leftwards.

We can thus deform the integration contour leftwards towards $\text{Re } s \rightarrow -\infty$ without changing the value of the integral, as long as we make detours to avoid integrating over any singularities. The integral then becomes

$$\tilde{\phi} = \lim_{a \rightarrow -\infty} \left[\int_{a-i\infty}^{a+i\infty} \frac{ds e^{st}}{2\pi i} \hat{\phi} \right] + \sum_m \int_{\gamma_{BC,m}} \frac{ds e^{st}}{2\pi i} \hat{\phi} + \sum_n \oint_{\gamma_n} \frac{ds e^{st}}{2\pi i} \hat{\phi}, \quad (2.42)$$

where $\gamma_{BC,m}$ denotes the detour taken to avoid branch cut m and γ_n denotes a small contour encircling a pole at $z = p_n$ anticlockwise – see Figure 2.3. Since the absolute value of the integrand is proportional to $|e^{st}| = e^{t \text{Re } s}$, the first integral goes to zero when the limit is evaluated. Assuming for simplicity that

there are no branch cuts, we are then left with only the third term. By the residue theorem, this term can be evaluated as

$$\tilde{\phi} = \sum_n \operatorname{Res}_{s=p_n} \left[e^{st} \hat{\phi} \right] = \{\text{if simple poles}\} = \sum_n e^{p_n t} \lim_{s \rightarrow p_n} \left[(s - p_n) \hat{\phi} \right], \quad (2.43)$$

where $\operatorname{Res}_{s=p} \hat{X}(s)$ is the residue of $\hat{X}(s)$ at $s = p$ (the general definition of which is somewhat involved and will be left out of this discussion, but can be found in e.g. Ref. [63] – or for that matter in Appendix A of [64]). In any case: what matters, largely, are clearly the poles of $\hat{\phi}$. These correspond to the independent wave modes supported by the plasma. In contrast to our previous fluid analysis, however, we now allow them to decay or grow in time.

A priori, these may come either from poles in the numerator or from zeros in the denominator. Choosing our initial condition so that $\tilde{F}_{\sigma,0}(v_z)$ is analytic for all species, the numerator as written is well-defined and analytic for all s with positive real part. For $\operatorname{Re} s = 0$, however, the integrand in the numerator integral

$$I_{N,\sigma}(s) = \int dv_z \frac{\tilde{F}_{\sigma,0}}{s + ikv_z} \quad (2.44)$$

has a pole at $v_z = is/k \in \mathbb{R}$, meaning the integral is undefined. This can be remedied, however, since we can analytically continue $I_{N,\sigma}(s)$ to non-positive $\operatorname{Re} s$ by taking a detour around the pole of the integrand when needed, effectively redefining the function as

$$I_{N,\sigma}(s) = \begin{cases} \int dv_z \frac{\tilde{F}_{\sigma,0}}{s + ikv_z}, & \operatorname{Re} s > 0 \\ \lim_{\varepsilon \rightarrow 0^+} \int dv_z \frac{\tilde{F}_{\sigma,0}}{s + \varepsilon + ikv_z}, & \operatorname{Re} s = 0 \\ \int dv_z \frac{\tilde{F}_{\sigma,0}}{s + ikv_z} + \frac{2\pi}{k} \tilde{F}_{\sigma,0} \left(\frac{is}{k} \right), & \operatorname{Re} s < 0, \end{cases} \quad (2.45)$$

where we have again used the residue theorem in the final line. Since this analytic continuation is nonsingular for all $s \in \mathbb{C}$, the numerator of $\hat{\phi}$ has no poles, and all its poles must come from zeros in the denominator.

These occur when

$$1 = \sum_{\sigma} \frac{iq_{\sigma}^2}{km_{\sigma}\varepsilon_0} I_{D,\sigma}(s), \quad (2.46)$$

where $I_{D,\sigma}(s)$ is the analytical continuation of the denominator integral, which by much the same logic as we used for $I_{N,\sigma}(s)$ can be explicitly written out as

$$I_{D,\sigma}(s) = \begin{cases} \int dv_z \frac{\partial_{v_z} \bar{F}_{\sigma}}{s + ikv_z}, & \operatorname{Re} s > 0 \\ \lim_{\varepsilon \rightarrow 0^+} \int dv_z \frac{\partial_{v_z} \bar{F}_{\sigma}}{s + \varepsilon + ikv_z}, & \operatorname{Re} s = 0 \\ \int dv_z \frac{\partial_{v_z} \bar{F}_{\sigma}}{s + ikv_z} + \frac{2\pi}{k} \bar{F}'_{\sigma} \left(\frac{is}{k} \right), & \operatorname{Re} s < 0. \end{cases} \quad (2.47)$$

Alternatively, we may write this more concisely in terms of the average between the integrals above and below the pole, equivalent to the Cauchy principal value (denoted with \mathcal{P}) of the integral along a contour moved upward or downward to pass straight through the pole – i.e. along $\{z \in \mathbb{C} : z = v_z + i \operatorname{Re} s/k\}$, meaning

$$I_{D,\sigma}(s) = \mathcal{P} \int dv_z \frac{\bar{F}'_\sigma(v_z + i \operatorname{Re} s/k)}{i \operatorname{Im} s + ikv_z} + \frac{\pi}{k} \bar{F}'_\sigma\left(\frac{is}{k}\right), \quad (2.48)$$

which is well-defined for all s . Setting $is = \omega = \omega_r + i\gamma$ for $\omega_r, \gamma \in \mathbb{R}$ and assuming the only two species are electrons and protons (which must have equal unperturbed density $\bar{n}_e = \bar{n}_p = \bar{n}$), our poles lie at ω satisfying

$$1 = \frac{\omega_{\text{pe}}^2}{k^2 \bar{n}} \mathcal{P} \int dv_z \frac{\bar{F}'_e(v_z + i\gamma/k) + \frac{m_e}{m_p} \bar{F}'_p(v_z + i\gamma/k)}{v_z - \omega_r/k} + \frac{i\pi\omega_{\text{pe}}^2}{k^2 \bar{n}} \left[\bar{F}'_e\left(\frac{\omega}{k}\right) + \frac{m_e}{m_p} \bar{F}'_p\left(\frac{\omega}{k}\right) \right], \quad (2.49)$$

where we have used the fact that $\omega_{\text{pe}}^2 = q_\sigma^2 \bar{n}_\sigma / m_\sigma \varepsilon_0$. We saw previously that for a zero-temperature plasma, the supported wave modes are not damped – meaning $\omega = \omega_r$ is real, and the growth rate γ is zero. And reasonably, this also falls out of Equation (2.49) – the kinetic dispersion relation for electrostatic waves – if we go into this limit. The reason for this is that \bar{f}_σ (and thus \bar{F}_σ) becomes a delta function, which allows the right-hand side of Equation (2.49) to be nonzero only when $\gamma = 0$ exactly. It is thus not entirely unreasonable to assume that $|\gamma|$ should be small compared to ω_r for many situations of interest – especially when $\bar{v}_{\text{th},\sigma}$ is low compared to the phase velocity ω_r/k . If we assume not only $|\gamma| \ll \omega_r$ but also that $|\gamma|/k \ll \delta v_{\text{char}}$, where $\delta v_{\text{char}} \sim \bar{v}_{\text{th},\sigma}$ is the characteristic velocity scale over which significant variations in $\bar{F}''_\sigma(v_z)$ occur, we can expand the kinetic dispersion relation around $\gamma = 0$ via

$$\bar{F}'_\sigma\left(v_z + \frac{i\gamma}{k}\right) \approx \bar{F}'_\sigma(v_z) + \frac{i\gamma}{k} \bar{F}''_\sigma(v_z). \quad (2.50)$$

Splitting the dispersion relation into its real and imaginary parts, we get two real-valued constraints – one for ω_r and one for γ . Keeping only the lowest-order terms, we get

$$\begin{cases} \frac{k^2 \bar{n}}{\omega_{\text{pe}}^2} \approx \mathcal{P} \int dv_z \frac{\bar{F}'_e(v_z) + \frac{m_e}{m_p} \bar{F}'_p(v_z)}{v_z - \omega_r/k}, \\ 0 \approx \frac{\gamma}{k} \mathcal{P} \int dv_z \frac{\bar{F}''_e(v_z) + \frac{m_e}{m_p} \bar{F}''_p(v_z)}{v_z - \omega_r/k} + \pi \left[\bar{F}'_e\left(\frac{\omega_r}{k}\right) + \frac{m_e}{m_p} \bar{F}'_p\left(\frac{\omega_r}{k}\right) \right]. \end{cases} \quad (2.51)$$

Here, the first equation defines an approximate dispersion relation for the real frequency ω_r , and the second can be solved for γ in terms of ω_r to yield

$$\gamma \approx - \frac{\pi k \bar{F}'(\omega_r/k)}{\mathcal{P} \int dv_z \frac{\bar{F}''(v_z)}{v_z - \omega_r/k}}, \quad \bar{F}(v) = \bar{F}_e(v) + \frac{m_e}{m_p} \bar{F}_p(v). \quad (2.52)$$

Note that unless the protons are very hot compared to the electrons, their contribution can be neglected, since $m_e/m_p \ll 1$. In other words, it is often the case that $\bar{F} \approx \bar{F}_e$ (for example, this holds for the setups examined in the appended papers). In the nonresonant – or low-but-nonzero temperature – regime where $\omega_r/k \gg \bar{v}_{\text{th},\sigma}$ for both species, the integral in the denominator will be approximately proportional to $\bar{F}'''(\omega_r/k)$. This quantity will be negative in the asymptotic limit for any monotonically decreasing and strictly positive \bar{F} , since the high- v_z tail of $\bar{F}(v_z)$ slowly flattens out, tending towards zero curvature as $v_z \rightarrow \infty$ and $\bar{F}(v_z) \rightarrow 0$. Since $\bar{F}'(\omega_r/k)$ is also negative, γ is as well, meaning any such electrostatic wave will be damped, and gradually lose energy to the particles in the plasma.

This process is known as *Landau damping*, after Lev Landau [65], and the growth rate γ tending to have the same sign as the slope of \bar{F} turns out to be quite broadly applicable as a rule of thumb. In other words, a negative slope at v_z around the phase velocity ω_r/k of the wave will generally lead to damping, and a positive slope will generally lead to *inverse Landau damping*, or growth, corresponding to an instability. From an intuitive point of view, this can be understood through the lens that the wave mostly exchanges energy with particles whose velocities are close to the phase velocity, since these are the particles which spend a significant amount of time being pushed in the same direction by the local perturbation moving along with the wave as they travel. In our case, this push can be taken as coming from the co-propagating perturbation in the electrostatic potential, but as pointed out in Swanson [62], Landau damping is actually quite a general feature of waves in gases and plasmas. For pressure waves in neutral gases, we may instead interpret the “pushing mechanism” as coming from the local pressure gradients caused by the propagating wave.

Whatever the mechanism, the end result is that particles moving slightly slower than the phase velocity tend to be accelerated so that they catch up with the wave, while particles moving slightly faster than the phase velocity are instead decelerated. From a distribution function point of view, this “diffusion” in velocity space acts to flatten the distribution function around $v_z = \omega_r/k$, tending towards a situation where $\bar{F}'(\omega_r/k) = 0$, and propagation is stable (i.e. the wave is neither damped nor growing). If $\bar{F}'(\omega_r/k)$ was originally negative, more particles were accelerated than decelerated as a result, and the net effect is that the wave amplitude is decreased, while the plasma is heated. In the same way, an originally positive slope would lead to a net deceleration of particles, cooling the plasma and increasing the amplitude of the wave. Of course, wave growth can only continue as long as a positive slope is maintained – as the slope approaches zero, growth slows and eventually stops. In the damping case, there are two possible end states: If the initial wave amplitude is small, the wave is fully damped out, meaning the distribution function still has nonzero negative slope around $v_z = \omega_r/k$ when the wave has dissipated. If the initial wave amplitude is large enough, the damping will instead eventually saturate, leaving a non-decaying propagating excitation and a flattened distribution function.

Notably, these nonlinear saturation processes are not captured by the linear theory of Landau damping outlined here. And there are indeed many aspects of nonlinear Landau damping which we have left out of this discussion in the interest of brevity, such as particle trapping, bounce times and the origins of the oscillations in wave amplitude which occur during saturation. For more information, see e.g. Chapter 8 of Swanson [62].

If we specifically choose \bar{f}_σ to be Maxwellian, i.e.

$$\bar{f}_\sigma(\mathbf{v}) = \frac{\bar{n}_\sigma}{(2\pi\bar{v}_{\text{th},\sigma}^2)^{3/2}} \exp\left(-\frac{|\mathbf{v}|^2}{2\bar{v}_{\text{th},\sigma}^2}\right), \quad (2.53)$$

the dispersion relation of Equation (2.49) can be expressed as

$$1 = \sum_\sigma \frac{\omega_{\text{p}\sigma}^2}{2k^2\bar{v}_{\text{th},\sigma}^2} Z'\left(\frac{\omega}{\sqrt{2}k\bar{v}_{\text{th},\sigma}}\right), \quad (2.54)$$

where

$$Z(\zeta) = \frac{1}{\sqrt{\pi}} \mathcal{P} \int_{\gamma_\zeta} \frac{e^{-\xi^2} d\xi}{\xi - \zeta} + i\sqrt{\pi}e^{-\zeta^2}, \quad (2.55)$$

with γ_ζ running parallel to the real axis at $\text{Im } \xi = \text{Im } \zeta$, is the so-called *plasma dispersion function*. Expanding the integral term in the nonresonant large- $|\zeta|$ limit, we find

$$\begin{cases} Z(\zeta) = -\frac{1}{\zeta} \left(1 + \frac{1}{2\zeta^2} + \dots\right) + i\sqrt{\pi}e^{-\zeta^2} \\ Z'(\zeta) = \frac{1}{\zeta^2} \left(1 + \frac{3}{2\zeta^2} + \dots\right) - 2i\sqrt{\pi}\zeta e^{-\zeta^2}, \end{cases} \quad (2.56)$$

and keeping only terms up to $\mathcal{O}(\zeta^{-4})$ in the first term, the dispersion relation becomes

$$1 = \sum_\sigma \left[\frac{\omega_{\text{p}\sigma}^2}{\omega^2} \left(1 + \frac{3k^2\bar{v}_{\text{th},\sigma}^2}{\omega^2}\right) - i\sqrt{\pi} \frac{\omega_{\text{p}\sigma}^2\omega}{\sqrt{2}k^3\bar{v}_{\text{th},\sigma}^3} e^{-\omega^2/2k^2\bar{v}_{\text{th},\sigma}^2} \right]. \quad (2.57)$$

Assuming weak damping (or potentially growth), i.e. $|\gamma| \ll \omega_r, k\bar{v}_{\text{th},\sigma}$, and splitting into real and imaginary parts, we get

$$\begin{cases} 1 \approx \sum_\sigma \frac{\omega_{\text{p}\sigma}^2}{\omega_r^2} \left(1 + \frac{3k^2\bar{v}_{\text{th},\sigma}^2}{\omega_r^2}\right) \\ \gamma \approx -\sqrt{\frac{\pi}{8}} \frac{\omega_r^2}{k^3} \sum_\sigma \frac{\omega_{\text{p}\sigma}^2}{\bar{v}_{\text{th},\sigma}^3} e^{-\omega_r^2/2k^2\bar{v}_{\text{th},\sigma}^2}, \end{cases} \quad (2.58)$$

and neglecting the thermal correction to the proton terms (effectively treating the protons as stationary, which is a reasonable approximation when modelling electron-scale phenomena like Langmuir waves), we recover Equation (2.32), as well as a first-order estimate of the damping rate γ in the cold-proton limit.

It should be noted that when initialising the Langmuir waves in Paper A, we used the exact numerical solution to Equation (2.54) rather than any of these approximations. Nevertheless, the approximations are useful in that they are easier to interpret than the derivative of the plasma dispersion function.

On top of Langmuir waves, the appended papers also examine *electron two-stream instability*. As outlined above, electrostatic instabilities in general require effective distribution functions $\bar{F}(v_z)$ with positive slope at some $v_z > 0$, in order to drive wave growth. For the electron two-stream instability in particular, this occurs due to the presence of two counter-streaming electron populations. To model this, we can use two separate electron species, labelled with b and c for the lower-density *beam* and the higher-density *core*, respectively. Each species can then separately have a Maxwellian background distribution, but centred on different positions in velocity space, separated by their relative flow velocity $\bar{V}_{\text{rel}} = \bar{V}_c - \bar{V}_b$. If we still enforce $\bar{\mathbf{J}} = 0$ by choosing $\bar{V}_{b,c}$ such that $\bar{n}_b \bar{V}_b + \bar{n}_c \bar{V}_c = 0$, we effectively just need to doppler-shift $\omega \rightarrow \omega - k\bar{V}_\sigma$ for each species to account for the fact that the Maxwellians are centred on different flow velocities. The dispersion relation thus becomes

$$1 = \sum_{\sigma} \frac{\omega_{\text{p}\sigma}^2}{2k^2 \bar{v}_{\text{th},\sigma}^2} Z' \left(\frac{\omega - k\bar{V}_\sigma}{\sqrt{2}k\bar{v}_{\text{th},\sigma}} \right), \quad (2.59)$$

or in the cold limit ($\bar{v}_{\text{th},\sigma} \ll |\bar{V}_\sigma|$)

$$1 = \frac{\omega_{\text{p}i}^2}{\omega_r^2} + \frac{\frac{\bar{n}_b}{\bar{n}} \omega_{\text{p}e}^2}{(\omega_r - k\bar{V}_b)^2} + \frac{\frac{\bar{n}_c}{\bar{n}} \omega_{\text{p}e}^2}{(\omega_r - k\bar{V}_c)^2}, \quad (2.60)$$

where $\omega_{\text{p}i}$ is the ion (proton) plasma frequency.

As outlined in Ref. [66], this configuration has two high-frequency ($\omega_r \gg \omega_{\text{p}i}$) modes in the limit of relatively weak and slow beams: a Langmuir-like mode with $\omega_r \approx \omega_{\text{p}}$, and a beam mode with $\omega_r \approx k\bar{V}_b$. When \bar{V}_{rel} is increased from zero, one of the modes will eventually become unstable due to the peak of the beam Maxwellian becoming separated enough from the core distribution function to cause positive slope at some v_z . For very weak or warm beams, it is the Langmuir mode which first becomes unstable, whereas dense cold beams instead make the beam mode unstable first. It is this latter case which corresponds to two-stream instability, sometimes also referred to as *electron beam instability*, *electron/electron streaming instability* or *bump-on-tail instability*. The specific criterion, originally derived by O’Neil and Malmberg in 1968 [67], states that two-stream instability occurs for

$$\frac{\bar{n}_b \bar{V}_b^3}{\bar{n} \bar{v}_{\text{th},b}^3} \gtrsim 1, \quad (2.61)$$

while a smaller-than-unity value of this expression gives rise to Langmuir beam instability. Having now described most of the necessary theory behind the phenomena we develop fluid closures for in the appended papers, let us move on to discussing how fluid closures can be derived, and what types of closures already exist for modelling various phenomena.

2.4 Closures from Theory

As the closure problem is a long-standing challenge in plasma physics, it has seen many attempts at being addressed throughout the history of the field, from different angles and with different aims. In this section, we will attempt to give an overview of these efforts, some of the closures which have been developed, as well as their applications and ranges of validity. To keep some semblance of brevity, the overview will by no means be exhaustive, but rather try to focus on closures which have seen significant use – and in particular on those which are more closely connected to the motivation and aims of the appended papers.

To set the stage, we will begin with a brief discussion of collisional closures through the lens of the Chapman–Enskog procedure, also giving an overview of the isothermal and adiabatic closures. We will then move on to collisionless plasmas, where we will cover the CGL and Le closures as well as the Hammett–Perkins closure and its generalisations.

2.4.1 Collisional Closures

As alluded to in Chapter 1, closures for collisional plasmas can be derived by doing perturbation theory around local thermal equilibrium. Specifically, this is done by assuming purely binary collisions and expanding the Vlasov equation in terms of the Knudsen number $\text{Kn} = \lambda_{\text{mfp}}/L$, where λ_{mfp} is the mean free path of constituent particles with respect to Coulomb collisions and L is the characteristic length scale of the phenomena one is interested in modelling. This approach was originally developed in the context of neutral gases in 1916–1917, independently by Sydney Chapman and David Enskog [47, 48], and then applied to plasmas in 1939 by Chapman and Cowling [51]. In 1958, the framework was used by S. I. Braginskii to derive the so-called *Braginskii equations* [49, 50], which constitute a closed fluid model for collisional plasmas.

At zeroth order in Kn , the distribution function is Maxwellian (i.e. equal to \bar{f}_σ as defined in Equation (2.53)), and it can be shown [51] that the first-order correction is given by

$$\delta f_\sigma = - \left[\frac{v_{\text{th},\sigma}}{n_\sigma} \mathbf{A} \cdot \nabla \ln T_\sigma + \frac{2}{n_\sigma} \mathbf{B} : \nabla \mathbf{V}_\sigma \right] \bar{f}_\sigma, \quad (2.62)$$

with $:$ denoting double contraction, so that e.g. $\mathbf{T} : \mathbf{ab} = T_{ij} a_i b_j$ for an arbitrary second-order tensor \mathbf{T} and vectors \mathbf{a} and \mathbf{b} . In this expression for δf_σ , $\mathbf{A} = \mathbf{A}(\mathbf{v})$ is a vector and $\mathbf{B} = \mathbf{B}(\mathbf{v})$ is a second-order tensor, both of which can be calculated to desired accuracy by expanding them into convergent series in terms of so-called *Sonine* (or *generalised Laguerre*) polynomials, and then solving a set of integral equations. The details of such calculations – as well as the integral equations in question themselves – will be left out of the present discussion for brevity (they can be found in e.g. [51]). One important consequence is nevertheless that to first order in Kn , the pressure tensor in the unmagnetised case is given by

$$\mathbf{p}_\sigma = p_\sigma \mathbf{I} - 2\eta_\sigma \{ \nabla \mathbf{V}_\sigma \} + \frac{2}{3} \eta_\sigma \nabla \cdot \mathbf{V}_\sigma \mathbf{I}, \quad (2.63)$$

where the (mass-normalised) viscosity $\eta_\sigma \sim \bar{n}_\sigma \lambda_{\text{mfp},\sigma} \bar{v}_{\text{th},\sigma}$ can be calculated from **B**. Since the only unknown quantity in the momentum equation thus becomes the scalar pressure $p_\sigma = \frac{1}{3} \text{Tr } \mathbf{p}_\sigma$, we only need a closure for (one-third) the trace of the tensorial pressure equation. In other words, we need to find an expression for the heat flux vector \mathbf{q}_σ defined so that $[\mathbf{q}_\sigma]_i = \frac{1}{2} \sum_j [\mathbf{q}_\sigma]_{ijj}$, i.e. half the contraction of the full heat flux tensor. As luck would have it, such an expression can be calculated, yielding the closure

$$\mathbf{q}_e = -\kappa_e \nabla T_e + \alpha p_e (\mathbf{V}_e - \mathbf{V}_i), \quad \mathbf{q}_i = -\kappa_i \nabla T_i, \quad (2.64)$$

where the thermal force coefficient α and the thermal conductivities κ_σ can be calculated from **A**. As we can see, we have two types of contributions to the heat flux – terms proportional to temperature gradients, and for the electrons also a friction-related term $\propto (-)T_e \mathbf{J}$. In principle, such a friction term also exists for the ions, but can be neglected since there is very little deflection and deceleration of ions due to collisions with electrons (as opposed to the reverse, which is quite prevalent). As we shall see, terms similar to these will continue to crop up in the other heat flux closures we will discuss – but with some significant differences. For a full treatment of the Braginskii equations, including also the magnetised case which we have left out of the discussion here, see e.g. Ref. [68].

Before moving on to collisionless closures, we would be remiss not to mention the fact that in certain collisional limits, one can derive even simpler closures, relating the pressure directly to the particle density without even needing to involve the heat flux [46]. In other words, the moment equation hierarchy can be truncated already at the momentum equation. The first of these limits is the *isothermal* one, where one is interested only in phenomena with characteristic velocities $v_{\text{char}} \ll \bar{v}_{\text{th},\sigma}$, meaning the heat flux is strong enough to practically instantly equilibrate the temperature over large distances. In an unmagnetised plasma, this would correspond to having $p_\sigma = n_\sigma T_\sigma$ with T_σ constant in 3D space. Such closures are particularly relevant for electron species when considering phenomena occurring on ion time scales.

The opposite limit, $v_{\text{char}} \gg \bar{v}_{\text{th},\sigma}$, also simplifies in a similar manner. In this regime, known as the *adiabatic* limit, the comoving rate of change in the pressure will be dominated by the terms $\propto \nabla \mathbf{V}_\sigma$, meaning we effectively have

$$\begin{aligned} \partial_t \mathbf{p}_\sigma + \nabla \cdot (\mathbf{V}_\sigma \mathbf{p}_\sigma) + 2\{\mathbf{p}_\sigma \cdot \nabla \mathbf{V}_\sigma\} &\approx 0 \\ \xrightarrow{\text{Tr}_N[\cdot]} N D_t p_\sigma + (N+2)p_\sigma \nabla \cdot \mathbf{V}_\sigma &\approx 0, \end{aligned} \quad (2.65)$$

where N is the dimensionality of the system (i.e. the number of dimensions over which \mathbf{p}_σ is isotropic), and Tr_N denotes the trace over these dimensions. Using the continuity equation, we can then rearrange this as

$$\begin{aligned} p_\sigma^{-1} D_t p_\sigma &= \frac{N+2}{N} n_\sigma^{-1} D_t n_\sigma \\ \Leftrightarrow D_t \left(\frac{p_\sigma}{n_\sigma^\Gamma} \right) &= 0, \quad \Gamma = \frac{N+2}{N} \end{aligned} \quad (2.66)$$

In other words, the ratio p_σ/n_σ^Γ is constant in the frame moving with the plasma. The ratio $\Gamma = \frac{N+2}{N}$ (more commonly denoted by γ , which we reserve for the growth rate) is known as the *adiabatic index* of the plasma, and relates the heat capacity at constant pressure C_P to the heat capacity at constant volume C_V via $\Gamma = C_P/C_V$. Because of this, the quantity in question is also known as the *heat capacity ratio*. In the highly collisional regime, \mathbf{p}_σ is fully isotropic, meaning $N = 3$ and $\Gamma = 5/3$.

2.4.2 Collisionless Closures

Going into the adiabatic limit and prescribing $\mathbf{q}_\sigma = 0$ was also one of the first closures considered for collisionless plasmas. In the unmagnetised collisionless case, the line of thinking outlined above for collisional plasmas cannot easily be applied, since isotropisation does not occur. In the *magnetised* collisionless case, however, the plasma becomes *gyrotropic* to leading order, i.e. locally cylindrically symmetric around \mathbf{B} (assuming relatively slow variations in time and space). Because of this, one *can* perform a similar derivation to the one outlined above, as shown by Chew, Goldberger and Low in an influential 1956 paper [69], treating the directions $\parallel \mathbf{B}$ and $\perp \mathbf{B}$ separately.

In general, the gyrotropic approximation is valid at length scales L significantly larger than the characteristic Larmor radius $r_L = \bar{v}_{\text{th},\sigma\perp}/|\omega_{c\sigma}|$ associated with the cyclotron oscillations of species σ . In the $r_L/L \rightarrow 0$ limit, the pressure and heat flux tensors reduce to

$$\begin{cases} \mathbf{p}_\sigma = p_{\sigma\parallel} \hat{\mathbf{b}}\hat{\mathbf{b}} + p_{\sigma\perp} (\mathbf{I} - \hat{\mathbf{b}}\hat{\mathbf{b}}) \\ \mathbf{q}_\sigma = q_{\sigma\parallel} \hat{\mathbf{b}}^{(3)} + 3q_{\sigma\perp} \{ (\mathbf{I} - \hat{\mathbf{b}}\hat{\mathbf{b}}) \hat{\mathbf{b}} \}, \end{cases} \quad (2.67)$$

where $\hat{\mathbf{b}}$ is the unit vector in the direction of \mathbf{B} .

Additionally taking the adiabatic limit by postulating $q_{\sigma\parallel} = q_{\sigma\perp} = 0$, the pressure equation can be simplified into the two scalar equations

$$\begin{cases} D_t p_{\sigma\parallel} + p_{\sigma\parallel} (\nabla \cdot \mathbf{V}_\sigma + 2\hat{\mathbf{b}}\hat{\mathbf{b}} : \nabla \mathbf{V}_\sigma) = 0 \\ D_t p_{\sigma\perp} + p_{\sigma\perp} (2\nabla \cdot \mathbf{V}_\sigma - \hat{\mathbf{b}}\hat{\mathbf{b}} : \nabla \mathbf{V}_\sigma) = 0, \end{cases} \quad (2.68)$$

and we have reached a closed set of equations. We have not yet reached an equation of state for the pressure, however. To find such an expression, representing the magnetised collisionless analogue of Equation (2.66), one needs to make several additional simplifying assumptions, due to the presence of the terms $\propto \hat{\mathbf{b}}\hat{\mathbf{b}} : \nabla \mathbf{V}_\sigma$ in the pressure equations.

For gyrotropic pressure \mathbf{p}_σ , it can in general be shown that the adiabatic limit $\mathbf{q}_\sigma \rightarrow 0$ yields

$$D_t \ln \left(\frac{p_{\sigma\parallel} |\mathbf{B}|^2}{n_\sigma^3} \right) = -2D_t \ln \left(\frac{p_{\sigma\perp}}{n_\sigma |\mathbf{B}|} \right) \propto \hat{\mathbf{b}} \cdot [\nabla \times (\mathbf{E} + \mathbf{V}_\sigma \times \mathbf{B})], \quad (2.69)$$

meaning that if the expression on the right is zero, then $p_{\sigma\parallel}|\mathbf{B}|^2/n^3$ and $p_{\sigma\perp}/n|\mathbf{B}|$ are both constant, giving us the equations of state we are after. One way this can hold approximately is if $\mathbf{E} \approx -\mathbf{V}_\sigma \times \mathbf{B}$.

As outlined in e.g. Ref. [70], one can reach precisely this relation for the proton species by neglecting electron inertia – i.e. focusing on time scales $\tau \gg |\omega_{ce}|^{-1}$ – and assuming relatively cold electrons, so that

$$\frac{m_e |D_t \mathbf{V}_e|}{e |\mathbf{V}_e \times \mathbf{B}|} \sim \frac{1}{|\omega_{ce}| \tau} \ll 1, \quad \frac{m_e |\nabla \cdot \mathbf{p}_e|}{en_e |\mathbf{V}_e \times \mathbf{B}|} \sim \frac{r_{L,e} \bar{v}_{th,e}}{L \bar{V}_e} \ll 1. \quad (2.70)$$

One then further neglects the displacement current in Ampère’s law so that $\mathbf{J} \propto \nabla \times \mathbf{B}$, as well as the resulting Hall term $\propto (\nabla \times \mathbf{B}) \times \mathbf{B}$, to find

$$\mathbf{E} \approx -\mathbf{V}_i \times \mathbf{B}, \quad (2.71)$$

satisfying our requirement. In this limit, we thus have

$$\frac{p_{i\parallel} |\mathbf{B}|^2}{n^3}, \frac{p_{i\perp}}{n|\mathbf{B}|} = \text{const.}, \quad (2.72)$$

and these equations of state, derived by Chew, Goldberger and Low, are known as the CGL closure. Additionally, other adiabatic closures for magnetised plasmas are often considered CGL variants – for example, closures which do not neglect the Hall term are usually denoted *Hall-CGL* closures.

It should be noted that even without assuming anything apart from gyro-tropy and adiabaticity, Equation (2.69) implies that $p_{\sigma\parallel}$ and $p_{\sigma\perp}$ are related through

$$\frac{\sqrt[3]{p_{\sigma\parallel} p_{\sigma\perp}^2}}{n^{5/3}} = \text{const.} \quad (2.73)$$

This is a more widely applicable generalisation of the unmagnetised collisional expression, reducing to Equation (2.66) if the pressure is isotropic. But notably, it is not a complete closure in its own right, since one of the pressure components still needs to be evolved explicitly by solving the corresponding component of the pressure equation with $\mathbf{q}_\sigma = 0$.

There are many situations where these adiabatic and isothermal closures are not sufficient to adequately model important phenomena, however. One example is that of magnetic reconnection, where magnetic field geometry is rapidly rearranged, leading to fast acceleration and heating of particles. This process plays a central role in both solar flares and coronal mass ejections, and also underlies the production of auroras when the solar emissions reach and reconfigure the Earth’s magnetic field.

Schematically, a reconnecting region looks like Figure 2.4. Particles are initially largely tied to their respective field lines due to magnetisation, being advected towards the X-point (or, more realistically, X-line) where the separatrices meet. Eventually, the particles come close enough that their cyclotron motion is sufficient to reach the field lines on the other side of the X-line, breaking magnetisation. The region where this occurs is known as the *diffusion region* for the species in question. Since the Larmor radius is proportional

to the particle momentum (which typically scales with the square root of the particle mass), the ion diffusion region (IDR) is generally much larger than the electron diffusion region (EDR).




Figure 2.4: A schematic view of a reconnecting magnetic field. Two magnetic field lines oriented in opposite directions (black) move towards the separatrices (dotted grey), reconnect at the X-point and move apart following the orange arrows.

The breaking of magnetisation in this way allows the magnetic field lines to reconnect at a point along the X-line – something which would otherwise be forbidden by the frozen-in flux theorem [46, 71, 72] – changing the global geometry according to Figure 2.4. The reconnected field line is then expelled outwards, accelerating the newly re-magnetised particles towards the outflow region.

A little less than two decades ago [73, 74], a more accurate closure for guide-field collisionless reconnection (i.e. reconnection with significant \bar{B}_y) was developed by Le, Egedal and collaborators. The chief aim of this closure is to account for the high electron temperature anisotropy ($T_{e\parallel} \gg T_{e\perp}$) which often characterises reconnection in Earth’s magnetotail and similar environments. To derive it, one still needs to make several simplifying assumptions, but not quite as many as for the CGL closures.

Specifically, the closure is derived from kinetic theory by assuming the electrons involved are very hot, so that $\bar{v}_{\text{th},e}$ is much larger than all other relevant velocities. All distribution functions can then be expanded around a gyrotropic background in terms of $|\mathbf{V}_\sigma|/\bar{v}_{\text{th},e}$. Additionally, the electrons are assumed to be well-magnetised, so that their magnetic moment μ_e is conserved – an assumption which is valid everywhere except in the EDR very close to the X-line. This quantity $\mu_\sigma = m_\sigma \bar{v}_{\text{th},e\perp}^2 / 2|\mathbf{B}|$ is also known as the *first adiabatic invariant*.

Using Liouville’s theorem, which states that $f_\sigma(\mathbf{x}, \mathbf{v})$ is conserved along particle trajectories, the value of the distribution functions in the reconnecting plasma can be related to their values f_σ^∞ far away from the reconnection site, following the field lines. Sufficiently far away, the ambient plasma is assumed to be uniform with $\mathbf{E} \cdot \mathbf{B} = 0$ and $B = |\mathbf{B}|$ constant. Additionally,

the ambient distribution function f_e^∞ feeding electrons into the reconnection region is assumed to be fully gyrotropic.

As electrons propagate towards the reconnection site, they will be affected by both electrostatic and magnetic mirror forces. Because of this, particles with very low initial velocity may become trapped on their way through the reconnection site, while particles with higher initial velocity are able to pass through without being significantly affected. Assuming f_e^∞ is Maxwellian and treating trapped and untrapped electrons separately, one can derive the following relation between n_e , B and the parallel potential $\phi_{\parallel}(\mathbf{x}) = \int_{\mathbf{x}}^{\infty} dl \hat{\mathbf{b}} \cdot \mathbf{E}$:

$$\frac{n_e}{n_e^\infty} = \frac{2}{\sqrt{\pi}}(1-b)\sqrt{u} + e^u \operatorname{erfc}(\sqrt{u}) - b^{3/2} e^{u/b} \operatorname{erfc}\left(\sqrt{\frac{u}{b}}\right), \quad (2.74)$$

where $u = e\phi_{\parallel}/T_e^\infty > 0$ and $b = 1 - B/B^\infty > 0$. This can then be numerically solved for ϕ_{\parallel} , which can in turn be used to derive an equation of state relating p_{\parallel} and p_{\perp} to lower-order fluid quantities and \mathbf{B} , analogous to the CGL closure.

As discussed in [74], the exact equation of state is well approximated by the analytical expressions

$$\begin{cases} p_{\parallel*} = \frac{2n_*}{2+\alpha} + \frac{\pi^3}{6} \frac{2\alpha^2}{2\alpha+1}, \\ p_{\perp*} = \frac{n_*}{1+\alpha} + \frac{\alpha}{1+\alpha} n_* B_*, \end{cases} \quad (2.75)$$

where $\alpha = n_*^3/B_*^2$ and the * subscript denotes normalisation with respect to ambient values, so that e.g. $B_* = B/B^\infty$ and $p_{\parallel*} = p_{\parallel}/p^\infty$. All fluid quantities in this expression are electron quantities (with species index omitted for brevity). Recently, significant progress has also been made towards generalising this closure to the EDR, where μ_e is no longer conserved, by instead leveraging the z -component of the *second* adiabatic invariant, $\mathcal{J}_z \propto \oint dz v_z$ [75–77].

Another important case where the CGL-type closures discussed above are insufficient is when Landau damping or growth is important – which is very much true for the phenomena considered in the appended papers. Fluid models which attempt to model these effects are known as *Landau fluid* models, and were pioneered by Hammett and Perkins in a widely cited 1990 article [78], originally in the context of the so-called *ion temperature gradient* (or *ITG*) instability. The simplest closure developed in that paper is now known as the Hammett-Perkins closure, and relates the heat flux to the temperature perturbation in Fourier space. In the original derivation for the 1D electrostatic case, a closure of the form

$$\tilde{q}_\sigma = -\bar{n}_\sigma \chi \frac{\sqrt{2}\bar{v}_{\text{th},\sigma}}{|k|} ik \tilde{T}_\sigma \quad (2.76)$$

is postulated as an ansatz through analogy with Fick’s law, $\mathbf{q}_\sigma \propto \nabla T_\sigma$, from the collisional limit, with χ being a dimensionless proportionality constant (later shown to equal $2/\sqrt{\pi}$). This ansatz is then justified by comparing resulting expressions to the three-pole approximation of the plasma response function

$R(\zeta) = 1 + \zeta Z(\zeta)$ one can derive by linearising the three lowest-order moment equations.

The essential difference between this closure and the temperature gradient closures which appear in the collisional limit lies in the factor $|k|^{-1}$ present in the proportionality constant on the right-hand side of Equation (2.76). Because of this term, the real-space closure becomes nonlocal, since multiplication by $-ik/|k| = -i \operatorname{sgn}(k)$ in Fourier space corresponds to a so-called *Hilbert transform* in real space, usually denoted with H . Using the notation $\mathcal{H} = -H$ for the negative Hilbert transform, the real-space closure looks like

$$\begin{aligned} q_\sigma(z) &= -\sqrt{\frac{8}{\pi}} \bar{n}_\sigma \bar{v}_{\text{th},\sigma} \mathcal{H} \delta T_\sigma := \\ &:= -\sqrt{\frac{8}{\pi^3}} \bar{n}_\sigma \bar{v}_{\text{th},\sigma} \int_0^\infty dz' \frac{\delta T_\sigma(z+z') - \delta T_\sigma(z-z')}{z'}. \end{aligned} \quad (2.77)$$

In other words, to calculate the real-space heat flux q_σ at a given point in time and space, one needs to know the temperature *over all of space* at that time. This is very different from the collisional case, where we only needed to know ∇T_σ , which can be calculated locally. This type of closure, where one needs global information to evaluate the otherwise unknown quantity, is usually referred to as a *nonlocal* closure, and is generally disfavoured if local alternatives exist with similar accuracy, since computing global integrals is computationally expensive and makes any implementation of the closure into a fluid code more difficult to parallelise. In this particular case, however, the closure *is* local in Fourier space, meaning parallelisability is still possible if one's setup geometry allows for periodic boundary conditions, so that time evolution can be performed entirely within Fourier space.

While there is certainly a marked difference in behaviour between the Hammett-Perkins closure and Fick's law, there are also clear structural similarities, which are particularly visible in the Fourier domain. This is no coincidence. In fact, it is possible to formulate a more general version of the Hammett-Perkins closure which can handle arbitrary collisionality, as measured by λ_{mfp} . To accomplish this, one starts by prescribing a collision operator and a background distribution function, from which a more exact relation between \tilde{q}_σ and e.g. \tilde{T}_σ can then be computed, valid for arbitrary λ_{mfp} . In Ref. [79], the relatively simple *Krook-model* collision operator is used along with a Maxwellian background to derive

$$\tilde{q}_\sigma \propto Q(\zeta) \bar{n}_\sigma ik \tilde{T}_\sigma. \quad (2.78)$$

Here, $Q(\zeta)$ is a somewhat involved expression involving the plasma dispersion function which can be shown to reduce to a constant in the highly collisional limit (yielding Fick's law), and to an expression $\propto \bar{v}_{\text{th},\sigma}/|k|$ in the collisionless one (corresponding to the Hammett-Perkins closure).

Notably, Hammett-Perkins-like closures can also be applied to 3D magnetised dynamics $\parallel \mathbf{B}$, i.e. not only in the 1D electrostatic case. In the original Hammett-Perkins paper, this is illustrated for slab geometry used to model the ITG instability. If one does desire to use the closure for such cases, however,

the integrals of the Hilbert transform must be performed along magnetic field lines [80]. This raises additional issues with fluid code implementation, since the line integration may become poorly defined numerically for chaotic field line structures.

The dynamics $\perp \mathbf{B}$ are significantly more difficult to model well if one wants to include Landau damping – in fact, as discussed in Ref. [81], constructing an accurate equation of state for q_{\perp} is likely only possible if all relevant phase velocities ω_r/k are significantly smaller than $\bar{v}_{\text{th},\sigma}$. To more fully resolve the perpendicular heat flux, a higher-order closure involving the fourth-order moment \mathbf{r}_{σ} – or, alternatively, a dynamical \mathbf{q}_{σ} closure, supplying an expression for $\partial_t \mathbf{q}_{\sigma}$ – is required. Significant work has been devoted to developing such kurtosis-level Landau closures for various applications in the years since the original Hammett-Perkins paper (see e.g. Refs. [80, 82–86]).

Finally, returning to the electrostatic case, it should be underlined that the three-pole approximation of $R(\zeta)$ used to derive the Hammett-Perkins closure is not the only potentially viable such approximation. Indeed, any n -pole Padé approximant of $R(\zeta)$ can in principle be used to derive an equation of state for the n^{th} -order moment of the distribution function, but the resulting closure is not always well-behaved. Even seemingly well-behaved Padé approximants may give rise to unphysical higher-order wave modes, however – though closures where all such modes are heavily damped may still be useful. In Ref. [81], Padé approximants with up to 8 poles are considered, and all closures which accurately reproduce the kinetic dispersion relation are showcased.

In Refs. [81, 87], the notation $R_{m,n}(\zeta)$ is used for the m -pole Padé approximant which agrees with the asymptotic expansion of $R(\zeta)$ to n^{th} order in $1/\zeta$ and uses the remaining degrees of freedom to agree with the small- $|\zeta|$ expansion to highest possible order. In this view, the Hammett-Perkins closure can be derived from kinetic theory by selecting specifically the $R_{3,2}$ approximant and linearising the moment equations. However, both the $R_{3,0}$ and $R_{3,1}$ approximants are more accurate for small $|\zeta|$ – and have smaller maximum relative errors than $R_{3,2}$ on the real line. Choosing instead the $R_{3,1}$ approximant, which still has correct asymptotic behaviour to leading order, a different closure is found, namely

$$\tilde{q}_{\sigma} = \frac{3\pi - 8}{4 - \pi} \bar{p}_{\sigma} \tilde{V}_{\sigma} - i \frac{\sqrt{2\pi}}{4 - \pi} \bar{n}_{\sigma} \bar{v}_{\text{th},\sigma} \text{sgn}(k) \tilde{T}_{\sigma}, \quad (2.79)$$

or in real space

$$\delta q_{\sigma} = \frac{3\pi - 8}{4 - \pi} \bar{p}_{\sigma} \delta V_{\sigma} - \frac{\sqrt{2\pi}}{4 - \pi} \bar{n}_{\sigma} \bar{v}_{\text{th},\sigma} \mathcal{H} \delta T_{\sigma}. \quad (2.80)$$

As we can see, this closure still has a term proportional to the Hilbert transform of the temperature perturbation, analogous to the Hammett-Perkins closure. Now, however, there is also an additional advective term proportional to the ambient pressure times the flow velocity perturbation, reminiscent of the friction-related additional term present in the electron heat flux closure of the Braginskii equations. Interestingly, this type of closure is more similar to those we find for 1D electrostatic phenomena using data-driven techniques, which we will discuss more in Chapter 4.

It should be noted that regardless of how many poles we include in our Padé approximants, the underlying methodology used to derive all of these Landau closures is still fundamentally limited by the fact that it involves linearisation of the fluid equations. In other words, there is no guarantee that nonlinear effects will be modelled correctly. This is one of the chief motivations for exploring data-driven closure construction for Landau fluid models – the underlying methodology of which we will delve into in Chapter 3.

Chapter 3

Machine Learning

You would be hard pressed to find anyone on the street today who hasn't heard of machine learning in some form. The most common point of reference is likely generative AI and large language models. After all, these have for better or worse come to play an increasingly prevalent role in society over the past few years – whether in the context of customer service, hiring procedures, faster coding or disinformation on social media.

However, machine learning as a field is quite a bit broader than just generative AI. In general, it can refer to any methodology which infers the rules governing the behaviour of some system from data. And contrary to how it may seem, the field did not suddenly appear out of thin air with the launch of ChatGPT in November 2022. Needless to say, we will not cover every facet of machine learning in this chapter – one could write entire books on the subject – but we will cover some of the history of the field, and outline the basics of the machine learning methodology used in the appended papers.

3.1 Artificial Neural Networks

While some basic inference methods, such as linear regression [88–91], were known already in the early 20th century, much of the architecture underlying the modern advancements in machine learning is built upon another idea, first outlined in a 1943 paper by McCulloch and Pitts [92]: that of using artificial neurons, inspired by the biological ones which make up our own brains. In their original biological form, neurons can largely be viewed as binary functions, which given a set of incoming signals from other neurons, will either fire, corresponding to an output of 1, or not, corresponding to an output of 0. Assuming that the incoming signals are combined linearly, we can formalise it mathematically as a binary function $f(\mathbf{x})$ of the form

$$f(\mathbf{x}) = \Theta(\mathbf{w} \cdot \mathbf{x} + b), \quad (3.1)$$

where $\mathbf{x} \in \mathbb{R}^d$ is the input vector containing the strengths of the incoming signals, \mathbf{w} is a vector containing a set of tunable weights, b is a bias term

setting the activation threshold and Θ is the Heaviside step function. This formalisation of the neuron is known as a *perceptron* [52, 53].

For neurons to do anything useful, however, a large number of them must be connected together into a *neural network*, or *NN*. In the simplest case, we may imagine collecting N of them together to form a layer, corresponding to a function $f : \mathbb{R}^d \rightarrow \{0, 1\}^N$. Each neuron is allowed to have its own set of weights and biases, meaning the function will look like

$$f(\mathbf{x}) = \Theta(\mathbf{W}\mathbf{x} + \mathbf{b}), \quad (3.2)$$

where \mathbf{W} is now a matrix containing the weights of all the neurons, \mathbf{b} is a vector containing the biases, and the Heaviside function is understood to act element-wise.

It is often advantageous to replace the binary Heaviside function with some other function $\sigma(\mathbf{x})$, generally known as the *activation function* of the neuron. Popular alternatives include various types of sigmoids like $\tanh(x)$ and $\arctan(x)$, which can be viewed as smooth analogues of $\Theta(x)$, as well as functions which heavily damp negative inputs while leaving positive inputs largely unchanged, like the so-called ReLU and GELU functions. Choosing $\sigma(\mathbf{x})$ to be differentiable at least almost everywhere is necessary for gradient-based optimisation of the free parameters. This is particularly important for the training of deep, multi-layer networks, where *backpropagation* enables efficient computation of gradients [93–96]. To increase performance of such training algorithms, it is also often useful to choose activation functions with finite range [97].

To construct larger networks, we can simply connect several layers of neurons in series. This architecture – perhaps the simplest neural network architecture in regular use today – is known as the *multi-layer perceptron*, or *MLP*. In such a network with $L + 1$ layers, the output of layer l can be expressed recursively as

$$f^l(\mathbf{x}) = \sigma(\mathbf{W}^l f^{l-1}(\mathbf{x}) + \mathbf{b}^l), \quad (3.3)$$

as illustrated in Figure 3.1.

With M_l inputs and N_l outputs feeding into and from layer l , the weight and bias matrices \mathbf{W}^l and \mathbf{b}^l then have dimensions $N_l \times M_l$ and $N_l \times 1$, respectively. Commonly, the input and output layers, with indices $l = 0$ and $l = L$, may have a slightly different structure than the intermediate *hidden layers* with indices $1 \leq l \leq L - 1$. For example, the MLPs used in Paper B omit the activation function $\tanh(x)$ in the final output layer, to allow the final output to take values with arbitrary magnitude. Despite their relatively simple structure, MLPs still see application today (albeit often as part of more complex architectures), owing to their capacity for arbitrary function approximation [98].

For many applications, however, more sophisticated network structures are called for. For example, in audio and image processing, it has proven useful to restrict the neurons in a layer to depend only on “nearby” neurons in the previous layer. In this way, the number of free parameters is significantly reduced, which speeds up optimisation and alleviates some issues related to divergent




Figure 3.1: A schematic picture of a single layer in a multi-layer perceptron, using $M^l = 4$ inputs and $N^l = 5$ outputs. First, an affine transformation defined by the weight matrix \mathbf{W}^l and bias vector \mathbf{b}^l is applied to the four-dimensional input $f^{l-1}(\mathbf{x})$. Then, the activation function σ is applied element-wise, yielding the five-dimensional output $f^l(\mathbf{x})$.

or near-zero gradients which may otherwise occur during backpropagation [99, 100]. Such network structures are known as *convolutional neural networks*, or *CNNs*, since every layer effectively computes a convolution of the output from the previous layer [101].

In large language models and other generative AI (and increasingly also for applications where alternative architectures like CNNs previously dominated), a more complex structure known as a *transformer* is commonly used. This structure, first introduced in the influential 2017 paper “Attention is all you need” by Vaswani et al [102], uses a combination of MLP layers and a so-called *attention* block [103], which helps the network take into account nonlocal information – essentially, how the meaning of any one element of the input is affected by the *context* of all the other input values. As transformers are ultimately at most tangentially related to the aims of this thesis, we will not discuss their inner workings in detail, but the interested reader is referred to e.g. Refs. [104] and [105].

All of these kinds of neural network structures have started to see use in plasma physics over the past few years [106–114], sometimes with specific adaptations to more efficiently handle data defined on graphs, such as kinetic or fluid quantities evaluated on discretised grids [115–117]. Architectures adapted in this way are referred to as *graph neural networks*, or *GNNs* – see e.g. Ref. [118]. Precisely how neural networks are leveraged to achieve more accurate simulation varies greatly across different studies, however. Approaches range from attempting to fully replace traditional simulation with neural-network time evolution prediction [111, 112, 115–117] to using them to provide a fluid closure wholesale [106–110, 113, 114] to estimating free parameters in a fixed-form data-driven closure, as in Paper B.

With a naive implementation, using neural networks for direct prediction of time evolution runs the risk of yielding unphysical results due to violation of known theoretical constraints, such as conservation laws and physical symmetries. A promising framework for avoiding these issues is that of *physics-informed*

neural networks, or *PINNs* [119], where such constraints are explicitly taken into account in each simulation time step. For example, a PINN fluid solver might explicitly heavily penalise prediction of physical states which are not in local agreement with the moment equations to desired order. Another approach which warrants mentioning is that of Fourier Neural Operators (or *FNOs*) [120, 121], which may be considered an alternative to the usual PINN approach more suited to nonlocal phenomena (and has seen some use recently for both data-driven closures [113, 122] and tokamak disruption prediction [123]).

Of particular interest are adaptations of such approaches which are more robust with regards to extrapolation outside of the convex hull of the training data, such as the recently developed BEACONS framework (short for *Bounded-Error, Algebraically-COMposable Neural Solvers*) of Ref. [124]. Alternatively, one may of course use traditional simulation, and enforce the constraints through the form of the closure, as is done in Papers A and B, and discussed in Ref. [125].

3.2 Interpretability and Sparse Regression

While neural networks are very useful for many applications, they do have one major drawback: due to their complex structure, large number of free parameters and opaque optimisation process, their inner workings are very difficult to interpret. Thus, they are mainly useful for applications where understanding the exact details of how they work are less important. Some alternative non-neural-network machine learning approaches which have also seen significant use in plasma physics, such as Bayesian optimisation [126–128] (commonly used for optimisation problems where the cost function is particularly expensive to evaluate [129–134]) exhibit the same black-box opaqueness problem. Luckily, there are alternative machine learning methods which are more interpretable.

In large part, these more interpretable approaches can trace their lineages back to more traditional regression-type machine learning. This is the case both for the *sparse regression (SR)* we use in the appended papers, and for the *symbolic regression* methodology of Refs. [135–138]. In both of these frameworks, focus is placed not only on finding accurate models, but specifically on finding accurate models which are *as simple as possible*. This can significantly help with understanding the behaviour of the system being studied, and also facilitates further more careful analytical examination.

Specifically, the goal is to model a target quantity y in terms of a set of d other quantities $\{x_n\}$. The difference between the two frameworks lies in how y is allowed to depend on $\{x_n\}$. In symbolic regression, one selects a set of allowed operations and functions for combining the x_n quantities, e.g. $\left\{+, -, \times, \div, \textcircled{\cdot}^2, \sin \textcircled{\cdot}, \cos \textcircled{\cdot}\right\}$, but otherwise leaves the functional form free. By defining some metric for the complexity of a given combination of these operations and combining it with a more traditional error-based cost, one can optimise for both simplicity and accuracy simultaneously, while leaving the functional form semi-arbitrary. Commonly, this has been done through

encoding the functional form in a tree-like data structure and using “genetic” evolution-inspired optimisation algorithms [135, 139, 140].

However, symbolic regression generally scales poorly both with d and with the complexity of the true functional dependence of y upon the x_n quantities, since the number of possible functions increases combinatorially with the number of included operations and quantities. Sparse regression remedies this to a significant extent, at the cost of requiring a pre-defined library of M allowed terms $\theta_j(x_1, \dots, x_d)$, and specifying in advance how these terms should be combined to yield a candidate model.

Historically, linear combination of terms has generally been used, since the optimal model utilising some given subset of the allowed terms θ_j can then be found exactly and efficiently through e.g. linear least-squares regression. Each model $\hat{y} = \sum_j \xi_j \theta_j$ is then completely described by its coefficient vector $\boldsymbol{\xi}$. For convenience, we can collect all the measured pairs of y and θ_j data into a target vector \mathbf{y} and term library matrix Θ such that $y_i = y|_{\text{point } i}$ and $\Theta_{ij} = \theta_j|_{\text{point } i}$. The least-squares optimal model then corresponds to the coefficient vector that minimises

$$C(\boldsymbol{\xi}) = \|\mathbf{y} - \Theta\boldsymbol{\xi}\|^2 \quad (3.4)$$

under the constraint that only the coefficients ξ_j corresponding to θ_j in the subset of interest are allowed to be nonzero.

Needless to say, the selection of terms included in the term library becomes very important for these methods – and as is discussed in Paper B, the restriction to linear combination of terms can make the creation of a viable term library very difficult for certain problems. Additionally, even with linearly combined terms, finding *the* optimal model at a given complexity s (defined as the number of library terms θ_j which are used) is expensive, even for only moderately high s and M , since the number of candidate models $\binom{M}{s}$ still grows combinatorially. Indeed, the only way to be sure of finding it is to go through each value of s from 1 to M and manually check every possibility – performing what is known as *best-subset selection* [56, 57]. As such a procedure quickly becomes untenable at higher s and M , it is in practice necessary to somehow focus the search.

For this, many different approaches have been tried, as outlined in e.g. Ref. [57]. Largely, these can be split into three categories: *selection* methods, where one restricts the number of library subsets considered, *shrinkage* methods, where the cost function is modified to penalise high-amplitude coefficient values, and *dimension reduction* methods, where the function space spanned by $\{\theta_j\}$ is transformed into one of lower dimension. Examples of the latter two approaches include e.g. *lasso* or *ridge* regression, and *principal component analysis/regression (PCA/PCR)* or *partial least squares (PLS)* regression, respectively. As the methodology used in the appended papers is best thought of as a selection-type approach, we will restrict our more in-depth discussion to these methods. For more information about the other types of SR, the interested reader is referred to the sources in the bibliography.

Perhaps the most intuitive selection method is *forward stepwise selection*, where one starts from the optimal 1-term model and adds terms sequentially, always choosing the additional term which yields the biggest accuracy improve-

ment. Analogously, one might instead start with the optimal M -term model and successively remove terms based on which one yields the smallest decrease in accuracy, which is known as *backward stepwise selection*. Both of these methods decrease the total number of least-squares regressions which need to be performed from $2^M - 1$ to $\frac{1}{2}M(M + 1)$. It should be underlined that this reduction in computational complexity from exponential to quadratic results in very significant speedup even for relatively small term libraries – with 20 terms, $2^M - 1 \approx 10^6$, while $\frac{1}{2}M(M + 1) \approx 200$. Additionally, the fact that the very best model at each complexity is no longer always found can actually be helpful, in that it reduces overfitting.

3.2.1 STLS and the SINDy Framework

If M , or the number of samples N in the dataset, is very large, even this $\mathcal{O}(M^2)$ complexity may be impractically high. Using the forward or backward stepwise selection algorithm as a starting point, we can address this by removing the exhaustive search through all $\mathcal{O}(s)$ terms at each complexity to find the optimal term to add or remove. If we do so, the complexity is reduced to $\mathcal{O}(M)$. But we then need to replace the exhaustive search with some heuristic for selecting which term should be added or removed.

One such heuristic-based selection method is the *sequentially thresholded least-squares* algorithm, or *STLS*, introduced in Ref. [58] as a more computationally tractable alternative to other options (such as the lasso) for especially large datasets. This is particularly relevant for applications to cases like ours, where the unknown quantity y is a field, with a separate value at each point in spacetime – and even more so in two- and three-dimensional settings.

Being a modified version of backward stepwise selection, the STLS algorithm starts from the least-squares-optimal full M -term model. But instead of performing an exhaustive search to find the optimal next simpler model, in STLS the term marked for deletion is simply the one whose coefficient ξ_j has the smallest absolute value. As long as the term library is normalised [58, 141, 142], so that the variance of each term θ_j is the same over the dataset considered, negligible optimal coefficients generally signify that the corresponding terms are unimportant or spurious. After all, if a term θ_j can be leveraged to account for a significant portion of the variation in y , its optimal coefficient will tend to be order unity, while uncorrelated terms will generally have small optimal coefficients.

There are, however, a few caveats. In particular, having a subset of terms in the term library which are heavily correlated with each other can cause issues: A pair of such terms could in principle have relatively large coefficients with opposite sign without significantly decreasing the accuracy of the model, even if they are ultimately spurious. Additionally, if they *do* have explanatory power, they will at least to some degree be interchangeable. This may distort their perceived importance according to the heuristic, since e.g. $0.6\theta_a + 0.4\theta_b$ or $1.2\theta_b - 0.2\theta_a$ would be approximately equivalent to using either θ_a or θ_b on their own if $\theta_a \sim \theta_b$, but each case would lead to very different “importance” scores for the two terms in question. To address this, one can calculate the correlations

between terms already when constructing the term library, and if necessary remove highly correlated terms – either based on physics considerations or on some other rule.

Ultimately, this rarely causes severe problems in practice, but can lead to inconsistent (i.e. dataset-dependent) model identification, as there are multiple viable models at each given complexity, or to the inclusion of pairs of spurious terms. Spurious terms can nevertheless be identified as such by manually calculating the change in model accuracy which would be induced by their removal – corresponding to the δ_{FVU} measure of Paper A for the variance-normalised least-squared error used in STLS (commonly abbreviated *FVU*, for *fraction of variance unexplained*). For a model \hat{y} giving a set of predictions $\hat{\mathbf{y}}$ across a dataset, the FVU of the model over said dataset is defined as

$$\text{FVU}[\hat{\mathbf{y}}] = \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}, \quad (3.5)$$

where \bar{y} is the mean of y over the dataset, y_i is the true value of the target quantity at point i , and the summation runs over all points in the dataset.

This type of SR, applied to the modelling of dynamical systems, is known as *Sparse Identification of Nonlinear Dynamics (SINDy)* [58] – or *PDE-FIND* when applied to systems described by partial differential equations [59, 60]. Over the past decade, this framework has been adapted for a wide variety of applications, in many different fields [143–147] – including plasma physics [125, 148–151]. We use methodology based on this framework throughout the appended papers, with various modifications.

The main modification to the standard SINDy methodology we make is that instead of using the raw Θ_{ij} and $y_i = q|_{\text{point } i}$ data from our kinetic simulations, we integrate over small spacetime environments surrounding each point. This is useful for reducing the particle noise which is otherwise present due to the particle-in-cell methodology we use for our kinetic simulations. As outlined in Ref. [152], this both increases the probability convergence for the STLS algorithm and improves the accuracy of the estimated optimal coefficients, at very limited additional computational cost. For terms involving derivatives this is particularly valuable, since numerical differentiation tends to amplify high-frequency noise. Integrating over environments of each spacetime point in this way may be considered a special case of the weak SINDy (WSINDy) approach introduced in Refs. [153, 154], using a box function integration kernel.

Another deviation from the standard SINDy approach in our work is that we do not perform STLS over the entire dataset at once for simulations where the dynamics differ greatly between different points in time. Ideally, one would of course want to discover a closure which is applicable to arbitrary phenomena, but the linear combination of terms in SINDy limits expressive power too much for this to be viable for our purposes. To get around this, we split our data-driven closure discovery into two steps: a first step, where we try to identify what the overall form of the closure should be, and a second step, where models for the free parameters in the identified closure form are constructed. To our knowledge, this two-step procedure has not been used previously, and seems

like a promising avenue for future data-driven closure discovery endeavours. In particular, this split into two discovery steps further encourages model sparsity and interpretability, since the full complexity of the data does not need to be captured in a single SR step.

In the first step, we split our simulation data into time slices such that the dynamics across a single time slice is broadly similar, and perform volume-integrated SINDy over samples from each time slice separately. This is what nets us the six-term heat flux closure of Paper A, where the coefficients ξ_j in front of the six identified terms (relabelled A_1 through A_6) are allowed to vary with plasma conditions. In this step we additionally perform 10-fold cross-validation to verify that these six terms are identified consistently, and that the model is not overfitting the data.

The second step is significantly less straightforward, which is why it was left out of Paper A. Since the optimal values for the free parameters may depend on local (or global) conditions in complex ways, defining a suitable term library with sufficient expressive power under the linear combination of regular SINDy is very difficult. To get around this, one alternative is to make use of less interpretable but freer functional forms like MLPs, as we illustrate in Paper B. After all, since the approximate functional dependence of the free parameters upon any set of candidate dependent quantities x_n (e.g. time-slice-averaged fluid quantities) is expected to be complex, finding an easily interpretable model may be difficult.

Alternatively, we may insist on trying to find a more interpretable model. As we show in Paper B, this is also possible, at least to some extent. As regular linear combination of terms is insufficient, a natural generalisation is to consider rational functional forms. Versions of SINDy modified along these lines have been considered previously [155, 156]. Specifically for our problem, however, where we seek an explicit rational model for $y = q$, these approaches effectively reduce to applying e.g. STLS to sparsely minimise the cost function

$$C_L(\boldsymbol{\xi}, \boldsymbol{\zeta}) = \sum_i \left[\sum_j \Theta_{ij} (y_i \zeta_j - \xi_j) \right]^2, \quad (3.6)$$

using this linear problem as a proxy for the more difficult nonlinear problem of minimising

$$C_{NL}(\boldsymbol{\xi}, \boldsymbol{\zeta}) = \sum_i \left(y_i - \frac{\sum_j \Theta_{ij} \xi_j}{\sum_k \Theta_{ik} \zeta_k} \right)^2, \quad (3.7)$$

which is what we are actually aiming to do. As can be seen in these equations, we now have two coefficient vectors $\boldsymbol{\xi}$ and $\boldsymbol{\zeta}$ to optimise – containing the free parameters in the numerator and denominator of our rational model, respectively.

Using a linear proxy cost function in this way is in theory a very clever approach, since we retain the linear least-squares optimisation which makes SINDy so computationally efficient. Additionally, any nontrivial ideal model which fully satisfies $C_{NL}(\boldsymbol{\xi}, \boldsymbol{\zeta}) = 0$ will also necessarily satisfy $C_L(\boldsymbol{\xi}, \boldsymbol{\zeta}) = 0$, and vice versa. In practice, however, there will always be some level of noise,

meaning no such ideal solution exists. And perhaps even more crucially, the x_n quantities used to construct our term library may simply be insufficient to fully explain the variations in our target quantity y . Indeed, we expect this to be the case when we attempt to estimate the free parameters in the closure.

To see why this is problematic, let us examine the error contribution from point i in both cost functions. Denoting the error contributions to C_L and C_{NL} by δ_i and ε_i respectively, so that $C_L(\boldsymbol{\xi}, \boldsymbol{\zeta}) = \sum_i \delta_i^2$ and $C_{NL}(\boldsymbol{\xi}, \boldsymbol{\zeta}) = \sum_i \varepsilon_i^2$, we have

$$\delta_i = \varepsilon_i \sum_k \Theta_{ik} \zeta_k. \quad (3.8)$$

The presence of the sum on the right hand side means that for models where the optimal ε_i is still typically relatively large, it will be advantageous for the linear optimiser to “cheat”, and decrease δ_i by minimising the sum instead of ε_i . And since the linear least squares algorithm always finds the global optimum, the linear proxy optimisation will always fail in this way once the typical difference between y and the C_{NL} -optimal \hat{y} is above some threshold.

Notably, this will be the case even if we enforce some restriction on the size of ζ_k , such as requiring that $\zeta_0 = 1$ or that $\sum_k \zeta_k = 1$ (which we would want to do regardless to ensure that the optimal pair of $\boldsymbol{\xi}$ and $\boldsymbol{\zeta}$ is unique). If the optimal \hat{y} differs from y too much, δ_i can almost always be decreased more by leveraging the correlations between the terms in θ_j to find a model with $\sum_k \Theta_{ik} \zeta_k \approx 0$ for all i than by actually minimising ε_i .

3.2.2 Nonlinear Sparse Regression

It seems, then, that there is no free lunch with regards to avoiding the nonlinear cost function C_{NL} . In Paper B, we thus chose to simply use this cost function as-is, and accept the increased cost of nonlinear least-squares regression. The details of the implementation are described in Sections II A 2 and II A 3 of the paper in question, but for convenience we will summarise the overall structure of the framework and discuss its advantages and disadvantages here as well.

Broadly, the nonlinear sparse regression (NLSR) framework we use is unchanged from the original SINDy method, except for specifically the fact that the linear regression performed at each complexity in the STLS procedure is now replaced with a global nonlinear optimisation step. The specific version we use in Paper B may also indirectly be considered a version of WSINDy, since the dependent quantities x_n we use are averaged over the time slices used in the first of our two SR steps. This is done to ensure that our dataset contains one datapoint per time slice for both our model terms θ_j (chosen to be multivariate monomials in x_n to yield rational models) and our target quantities $y = A_k$.

It should perhaps be emphasised that linear least-squares regression differs from its nonlinear counterpart not only in that the linear case is less computationally expensive, but on a more fundamental level: In the linear case, the global optimum can be found exactly through an analytical formula. This is very much not true in the nonlinear case – instead, finding the global minimum of C_{NL} can only be done by applying one’s nonlinear optimisation algorithm of choice and terminating the optimisation once further improvement is deemed

unlikely. Finding the global optimum can never be guaranteed, and one always has the option of spending more computational resources to potentially find a better model, with diminishing returns as more of the parameter landscape is explored. As was the case when switching from best subset selection to stepwise selection or STLS, however, this is not entirely a negative, in that overfitting is decreased.

The open-endedness of global nonlinear optimisation also entails a lot of freedom with regards to the choice of optimisation algorithm. In Paper B, we used the `basinhopping()` algorithm of SciPy [157] with an L-BFGS-B minimiser, to make use of the fact that the gradients of C_{NL} with respect to ξ and ζ are analytically computable. However, there is much room for further exploration of other nonlinear optimisation algorithms and tuning of relevant hyperparameters.

An additional difference between the second and first SR steps is the importance of culling lower-quality parts of the dataset. This is in principle also important in traditional SINDy if some parts of the dataset are significantly more noisy than others. But often (as is the case for us in the first step), the noisiest parts of the dataset are the parts where the dynamics of interest are very low-amplitude, and these regimes are automatically suppressed through our use of the (variance-normalised) mean-squared error as our error measure. In the second step, however, we must manually account for the relative reliability of the data from the various time slices.

There are a multitude of ways one could go about this, including e.g. weighting the data from a given time slice based on the accuracy of the optimal model found in step 1. For simplicity, however, we used a simple cutoff of 30% first-step FVU in Paper B: Coefficient values A_k from time slices with worse-performing optimal models than this were excluded from the dataset, since optimising our \hat{A}_k models to better fit these time slices is unlikely to significantly help the overall performance of the resulting q model. In fact, doing so could even decrease the overall performance, if the fit to data from more well-modelled time slices is worsened. The precise cutoff FVU could likely be optimised further, but doing so is not expected to yield any significant increase in performance compared to what we achieved in the paper.

A major hurdle for our NLSR framework which needs to be overcome when working specifically with rational functions is the fact that a naive implementation risks yielding models with poles within the parameter regime of interest, which could cause rare but potentially catastrophic failure of the resulting model. In principle, one could address this by simply throwing more data at the problem – if all parts of this domain are well-sampled, pole-ridden models would be eliminated by the optimisation algorithm itself.

However, such an approach is not particularly robust, as one can never be sure whether enough data has been supplied to eliminate all poles. Additionally, thoroughly sampling the entirety of the potential parameter space where the model might be applied somewhat defeats the purpose of constructing the closure in the first place – we want a model that can at least semi-reliably interpolate and extrapolate from the limited simulation data it is trained on.

As we show in section **II A 3** of Paper B, we can reliably address this issue in a straightforward manner without unnecessarily sacrificing functional flexibility by expressing the denominator polynomial $\sum_k \theta_k \zeta_k$ in the so-called *Bernstein basis* instead of the usual monomial one. In the single-variable case, the k^{th} degree- n Bernstein basis function looks like

$$B_{n,k}(x) = \binom{n}{k} x^k (1-x)^{n-k}. \quad (3.9)$$

Unlike the monomial basis, where every basis function has a different degree, all Bernstein basis functions have the same degree n as the polynomial being decomposed. They may be thought of as a set of $n + 1$ different degree- n polynomials “spread out as evenly as possible” over the interval $[0, 1]$, normalised so that

$$\sum_{k=0}^n B_{n,k}(x) = 1. \quad (3.10)$$

For $n = 2$, we thus have three second-degree basis functions – one which is largest close to $x = 0$, one which is largest close to $x = \frac{1}{2}$ and one which is largest close to $x = 1$.

There are a lot of details which need to be taken into account to fully implement this methodology – for example, defining sparsity in the Bernstein basis is not as straightforward as in the monomial basis. One natural definition (which is also the definition we use in the paper) is the number of terms with coefficient different from unity, since the simplest denominator polynomial may reasonably be taken as $f(x) = 1$, which corresponds to setting every coefficient equal to one, per Equation (3.10).

To avoid getting lost in minutiae, we will not go into further detail here. For a more thorough overview of our Bernstein basis methodology, see Section II A 3 of Paper B. At the end of the day, what is important is that applying our NLSR framework to the second sparse regression step yields rational models for the free closure parameters which are generally highly accurate and fully protected from divergences over the parameter regime of interest. Performance is comparable to that reached with MLP models, generally being a little less accurate but also less overfitted.

Since the dependence of the optimal A_k values on plasma conditions is complex (generally needing on the order of 50 terms in the numerator and denominator combined), the NLSR models are not quite as easily interpretable as one might be used to from linear SINDy. Nevertheless, the NLSR models are significantly more tractable to analyse analytically than the virtually completely opaque MLP models, and constitute a viable tool for future closure discovery.

Chapter 4

Summary and Outlook

Multi-scale collisionless plasma processes have long proven challenging to model both accurately and efficiently. In large part, this can be attributed to the lack of a general framework for collisionless closure construction akin to the collisional Chapman-Enskog procedure. As a result, while many theoretically derived closures have been proposed over the years, they are often quite limited in applicability. Over the past few years, however, a new line of attack has opened up in the form of *data-driven* closure construction, making use of recent advancements in the field of machine learning.

In this thesis, we restrict ourselves to one-dimensional electrostatic setups, focusing on Landau damping and electron two-stream instability, in order to lay the groundwork for more complex two- and three-dimensional applications. Through the first half of the two-step sparse regression (SR) approach outlined in Chapter 3, we identify a local six-term closure in Paper A which regularly accounts for upwards of 90% of the variation in the heat flux. In Paper B, we then generalise this closure to multi-species modelling and proceed to the second step, illustrating how the most important free parameters can be estimated from box-averaged fluid quantities. Below we give a more in-depth summary of the findings of the appended papers, and outline potential avenues of investigation for future work.

4.1 Summary of papers

Applying our two-step SR procedure requires a dataset. For both appended papers, this dataset is constructed by performing fully kinetic particle-in-cell simulations using the OSIRIS code, and saving all relevant fluid quantities at regular time intervals. To ensure any discovered closures apply over a range of plasma parameters, we perform simulations with a variety of initial conditions.

In Paper A, six Landau-damped Langmuir wave setups and seven two-stream-unstable setups are considered. For the former, Langmuir waves with fix wavenumber are excited in an initially Maxwellian plasma with thermal speeds v_{th} between $8 \times 10^{-3} c$ and $1.5 \times 10^{-2} c$, so as to yield decay rates compatible with the duration and temporal resolution of the simulations, after

which the waves are allowed to decay self-consistently. For the latter, the plasma is initialised as two counter-streaming Maxwellian electron populations with fix and equal initial thermal speed on top of a zero-temperature stationary ion (proton) background. The beam density ratio n_b/n is then varied from 1% to 50%, adjusting the flow velocities V_b and V_c to stay in the zero-current centre-of-momentum (CoM) frame with a fix relative flow velocity $V_{\text{rel}} = V_c - V_b$.

In Paper B, we focus wholly on two-stream-unstable setups, varying both the beam density ratio n_b/n , the relative flow velocity V_{rel} and the thermal speed of the two electron species. To more thoroughly examine the generality of the closure form discovered in Paper A and evaluate how accurately its free parameters may be estimated, we perform 76 simulations in total. Of these, 47 constitute parameter sweeps around the initial conditions considered in Paper A, and the remaining 29 use initial conditions sampled randomly from the four-dimensional parameter space exhibiting the desired instability.

While all particle-in-cell simulations are fully special-relativistic, at most weakly relativistic setups are considered, since the moment equations underlying our closure framework, described in Section 2.2, are fundamentally non-relativistic. For an exhaustive account of the precise simulation parameters used, see Section 2.2 of Paper A and Section II C of Paper B, respectively.

Having collected all data required, we apply our two-step SR procedure. In the first step, where we identify the closure form – chiefly the focus of Paper A – the target quantity y is the heat flux q , and the dependent quantities x_n are chosen to be the lower-order fluid quantities n , V and v_{th} . Since we initialise our simulations in the overall CoM frame, V should be interpreted as $V - v_{\text{CoM}}$. With this interpretation, all of our dependent quantities are Galilean-invariant, ensuring the same holds for any resulting q model.

Next, a term library is selected. In Paper A, we firstly include all terms of the form

$$\theta_j = nv_{\text{th}}^\alpha V^{3-\alpha} \quad (4.1)$$

with positive integer $\alpha \leq 3$, having the dimensions of heat flux. To additionally allow for the discovery of Hammett-Perkins-like closures, we also include terms where one or more of the fluid quantities is replaced with its spatial derivative (e.g. $n \rightarrow \partial_x n$). In Paper B, we generalise this to multiple species, replacing $\{n, V, v_{\text{th}}\} \rightarrow \{n_\sigma, V_\sigma - \bar{V}_\sigma, v_{\text{th},\sigma}\}$ for each species σ , with \bar{V}_σ being the spatially averaged flow velocity for the species in question.

Applying the SINDy-based approach outlined in Section 3.2.1 to temporal slices of simulation data, we consistently identify a six-term model, split into $q = q_{\text{even}} + q_{\text{odd}}$ with

$$\begin{cases} q_{\text{even}} = A_1 n v_{\text{th}}^2 V + A_2 v_{\text{th}}^3 \partial_x n + A_3 n v_{\text{th}}^2 \partial_x v_{\text{th}} \\ q_{\text{odd}} = A_4 + A_5 n v_{\text{th}}^3 + A_6 n v_{\text{th}}^2 \partial_x V. \end{cases} \quad (4.2)$$

Here, the coefficients in q_{even} (q_{odd}) have even (odd) parity under spatial reflections, meaning $A_{1,2,3}$ are invariant while $A_{4,5,6}$ pick up a minus sign under such a transformation. This also means that $A_{4,5,6} \rightarrow 0$ for symmetric setups, such as standing wave Landau damping or equal-density two-stream instability. As shown in Paper B, the same closure form is also identified in the

multi-species case, but with $\{n, V, v_{\text{th}}\} \rightarrow \{n_\sigma, V_\sigma - \bar{V}_\sigma, v_{\text{th},\sigma}\}$ in accordance with the changes made to the term library.

This six-term closure regularly accounts for 95% of the variation in q . As we discuss in Section 3.3 of Paper A, the vast majority of this explanatory power lies in the $A_{1,4,5}$ terms, typically accounting for upwards of 90% of the variation on their own. Since the constant A_4 term does not affect the heat flux divergence appearing in the pressure equation (only being present to correct the 0th order A_5 contribution), a two-term closure with only the A_1 and A_5 terms is likely sufficient for many applications. The A_1 term can be understood as a convective heat flux of the type used in the Landau closure of Equations (2.79) and (2.80). The A_5 term, on the other hand, is a bit more mysterious, but as discussed in Appendix C of Paper B it may help capture the nonlocal parts of the Hammett-Perkins term, and/or nonlinear effects unaccounted for in previously discovered linear Landau closures.

The coefficients for the remaining $A_{2,3,6}$ terms are strongly correlated with the instantaneous growth rate $\gamma_q(t)$ for the heat flux perturbation, and are thus only non-negligible when there is significant Landau damping or growth. This is consistent with the fact that these three terms correspond to gradient-driven heat fluxes, which either intensify or weaken existing perturbations. In particular, the A_3 term is proportional to a temperature gradient, equivalent to the Fick's law heat fluxes discussed in Section 2.4 – a local approximation of the full Hammett-Perkins heat flux. The A_2 term analogously corresponds to a density gradient-driven heat flux, whereas the A_6 term is driven by gradients in (or divergences of) flow velocity. As noted in Paper A, it is often the case that $A_3 \sim 2A_6$, in which case the $A_{2,3}$ terms together constitute a heat flux contribution driven by pressure gradients.

Even without describing their individual interpretations, the consistent identification of these terms by the sequentially thresholded least-squares procedure may be understood from another point of view. These six terms are precisely the terms in the term library which contribute to the heat flux at zeroth or first order in perturbation theory. There is, however, one further term which is identified. As outlined in Section 3.5 of Paper A, this term, $\propto nv_{\text{th}}V^2$, seems to help model the nonlinear processes involved in the slowing of growth or decay (such as particle trapping). As such, it is mainly of use when $\partial_t|\gamma_q| < 0$.

To reach a complete closure, we need to perform the second step of our SR approach and estimate the free parameters A_k . Already in Paper A, we outline a few heuristics for this which can be derived from linear theory. For example, in the low- γ_q regime one can set $A_2 = A_3 = A_6 = 0$ and derive e.g. $A_1 \sim -3 + \frac{1}{2} \text{sgn}(k)A_5$ for $\omega_r \sim \omega_{\text{pe}} \sim |k|\bar{v}_{\text{th}}$. More generally, linear theory gives two constraints on the six closure coefficients, as outlined in Appendix A of Paper A, though in their general form they are quite involved, and have limited accuracy for nonlinear dynamics.

For the three-term $A_{1,4,5}$ model, the linear constraints yield an explicit, complete closure as the A_4 term does not influence closure performance. In

the single-species case, this linear completion of the closure looks like

$$\begin{cases} A_1 = -3 - \zeta_{\text{pe}}^2 - |\zeta|^2 + 4\zeta_r^2 \frac{1 + 3\zeta_{\text{pe}}^2}{1 + 3\zeta_{\text{pe}}^2 + 3|\zeta|^2} \\ A_5 = \frac{4\zeta_r |\zeta|^2}{1 + 3\zeta_{\text{pe}}^2 + 3|\zeta|^2}, \end{cases} \quad (4.3)$$

where $\zeta_r = \omega_r/k\bar{v}_{\text{th}}$ is the real part of the plasma dispersion function argument $\zeta = \omega/k\bar{v}_{\text{th}}$, and $\zeta_{\text{pe}} = \omega_{\text{pe}}/k\bar{v}_{\text{th}}$. Since these expressions depend on wave parameters, they are impractical for use in local closures, however – and even more so in the multi-species case, where the complexity of the analogous expressions is even higher (see Appendix A of Paper B). Additionally, their accuracy remains limited outside the linear regime.

In Paper B, we show how models for the free parameters can be identified in a more holistic manner through data-driven methods. To this end, we use both multi-layer perceptrons (MLPs) as described in Section 3.1 and the newly developed nonlinear sparse regression (NLSR) framework outlined in Section 3.2.2 to complete the closure by finding more generally applicable models for the free closure parameters. For several reasons, we restrict these efforts to the three-term $A_{1,4,5}$ model in the paper.

Firstly, the $A_{2,3,6}$ terms, as discussed previously, only account for a limited portion of the total heat flux variation (usually less than 10% even during maximum growth or decay). Furthermore, while two-stream instability in single-electron species fluid models can only be driven through manual insertion of terms like these in the closure, the same is not true in the multi-species case. If the beam and core electron populations are treated individually, with separate closures for q_b and q_c , the instability is captured even with e.g. a naive adiabatic closure $q_b = q_c = 0$.

A further reason is the fact that the $A_{2,3,6}$ terms correlate strongly with $\gamma_q(t)$, and including this quantity as part of the closure without inducing non-physical numerical instabilities is nontrivial. Additionally, the growth of the q_b and q_c perturbations in two-stream unstable setups is fundamentally nonlinear (see Section 3 A of Paper B), and does not occur at the same rate as the perturbations in the lower-order fluid quantities. This means that even if one wants to include γ_q in the closure, estimating this quantity is not straightforward without access to the heat flux – which is the quantity we are trying to model.

Thus, in the second step of our closure construction, our target quantities are $y = A_{1,4,5}$. For both the MLP and NLSR models we explore, we proceed by choosing our dependent quantities $\{x_1, \dots, x_d\}$ to be the four lower-order fluid quantities which define our space of initial conditions – averaged over each time slice to match the size of our y quantity dataset, i.e. $\mathbf{x} = [\bar{n}_b/\bar{n}, \bar{V}_{\text{rel}}/c, \bar{v}_{\text{th},b}/c, \bar{v}_{\text{th},c}/c]$. Training one MLP with three 25-node hidden layers for each A_k and species, we find that this set of x_n quantities is sufficient for predicting the target quantities with FVU error scores of 3–21% on training data and 4–30% on test data.

With the more interpretable NLSR approach, the A_k model accuracy is lowered somewhat due to the more restrictive rational functional form. However, this reduced flexibility also helps limit overfitting, yielding training FVUs of 5–31 % and test FVUs of 5–38 % for the fully divergence-protected Bernstein-basis models. It should be noted that these errors do not necessarily translate into poor performance for the resulting heat flux closure

$$\hat{q}_\sigma = \hat{A}_{1,\sigma} n_\sigma v_{\text{th},\sigma}^2 (V_\sigma - \bar{V}_\sigma) + \hat{A}_{4,\sigma} + \hat{A}_{5,\sigma} n_\sigma v_{\text{th},\sigma}^3, \quad (4.4)$$

since values of $A_{1,4,5}$ which deviate significantly from the least-squares optimum may nevertheless yield an accurate closure.

Choosing what metric to use for evaluating the closure accuracy requires some thought, however, as discussed in Section IIIB of Paper B. The most straightforward choice is perhaps the q FVU evaluated across all spacetime points i in a simulation, i.e.

$$\text{FVU}_{\text{sim}}[\hat{q}_\sigma] = \frac{\sum_i (q_{\sigma,i} - \hat{q}_{\sigma,i})^2}{\sum_i (q_{\sigma,i} - \bar{q}_\sigma)^2}. \quad (4.5)$$

This measure is ultimately flawed, however. In particular, bias-type inaccuracies in \hat{A}_4 are strongly amplified, since such errors yield error contributions across the entire dataset, despite not actually affecting $\partial_x q_\sigma$. Especially for low-variance data, such as q_b in the weak-beam limit, this can result in very high FVUs ($\gg 1$), even when $\partial_x q_\sigma$ is captured quite accurately.

A simple way to address this is to instead consider the $\partial_x q$ FVU, though this measure is also imperfect. For example, complications arise with regards to how exactly the spatial derivative should be evaluated numerically, as discussed in Paper B. Additionally, the performance of the closure in an actual fluid code implementation *really* only depends on the accuracy with which the pressure is evolved. Thus, there is strong reason to measure closure accuracy by how well we predict the entire right-hand side of the pressure equation

$$\partial_t p_\sigma = -V_\sigma \partial_x p_\sigma - 3p_\sigma \partial_x V_\sigma - \partial_x q_\sigma. \quad (4.6)$$

Notably, all errors will still come from the errors in $\partial_x q_\sigma$. In other words, the corresponding $\partial_t p$ FVU may be considered a re-normalised version of the $\partial_x q$ FVU which takes into account the fact that predicting $\partial_x q_\sigma$ poorly is acceptable as long as the other contributions to $\partial_t p$ dominate.

Generally, $\partial_t p$ FVU error rates lie between 5 and 30 % for the range of initial conditions we consider, with the core species being somewhat harder to model than the beam and combined species. For example, the median $\partial_t p$ FVU for the simulations with randomly sampled initial conditions is 15 % for both the beam and combined species but 20 % for the core species. The beam and combined $\partial_t p$ FVU furthermore correlate strongly with the beam density, being below 10 % for $n_b/n \sim 0.01$ and increasing for denser beams. On the other hand, the core-species FVU is largely uncorrelated with the beam density, instead being inversely correlated with the core thermal speed $\bar{v}_{\text{th},c}$.

This is connected to the main source of inaccuracies for the version of the closure implemented in Paper B – namely, strongly inhomogeneous conditions

across the simulation box (chiefly caused by fast, dense beams, especially in conjunction with cold backgrounds). In such cases, the box-averaged quantities in \mathbf{x} are no longer good predictors of local conditions, which significantly decreases the accuracy of the closure. This can be remedied by replacing the box averaging with more local averaging, preferably using a kernel with compact support to improve parallelisability. Notably, switching to local averaging in this way would be necessary regardless for simulation of large-scale phenomena – especially for two- and three-dimensional applications.

In summary, the two-step sparse regression procedure we have developed constitutes a promising framework for constructing data-driven closures. By initially letting the closure parameters vary freely and focusing on the closure form, we retain the interpretability advantage of sparse regression, while minimising the adverse effects of the restrictive function space. Furthermore, the nonlinear sparse regression method of Section 3.2.2 is a viable alternative to more opaque neural network models, as we show in Paper B. The electrostatic closure we identify is highly accurate over a large range of parameter values, significantly outperforming both the naive $q_\sigma = 0$ closure and the local approximation of the Hammett-Perkins closure. Though challenges remain, the methodology used and developed in the present work will likely prove useful for a number of future endeavours.

4.2 Outlook

Before switching focus to potential future applications of our data-driven closure discovery methodology, we should note that there is some additional work which could be done to more fully evaluate our electrostatic closure. In the appended papers, we show that the closure generally predicts $\partial_t p_\sigma$ data from our kinetic simulations accurately. However, a closure’s utility is ultimately determined by its performance when used in fluid simulations, through implementation either in a new purpose-built code or in an established code such as Gkeyll [158, 159]. In such efforts, it may be of interest to consider not only traditional finite difference or finite element methods, but also machine learning-inspired approaches such as physics-informed neural networks, or PINNs [119], and the more robust BEACONS approach of Ref. [124].

Generalising the closure, such as by extending it to handle relativistic effects or to higher-dimensionality simulations, is also a potential target of future endeavours, likely requiring significant analytical legwork. For example, fully Lorentz-covariant closure construction demands significant modification of the framework outlined in Section 2.2, since the Lorentz-tensor moment equations differ from their non-relativistic analogues in several ways which make closure construction more difficult (see e.g. Refs. [160–164]). Even outside of such explorations, however, a more thorough theoretical analysis of our electrostatic closure – beyond the comparisons to linear theory performed in the appended papers – is of interest.

For higher-dimensional scenarios, directly applying our two-step sparse regression methodology to kinetic simulation data from relevant setups – as

was done in the present work – is a promising way forward. After all, the development of our data-driven closure discovery methods is mainly motivated by finding new ways to construct collisionless closures for two- and three-dimensional collisionless phenomena, such as wave-particle interactions [165–169] and magnetic reconnection [169–178]. For such work, careful consideration of the tensorial nature of the pressure and heat flux is of vital importance.

In magnetic reconnection setups, it would be interesting to compare any data-driven closures discovered with those which have previously been derived analytically [73–77]. It would also be interesting to examine the effects of specific small-scale phenomena on global dynamics, eventually aiming to incorporate the net effects of important kinetic processes into a sub-grid scale closure, similar to the approaches of Refs. [179–183]. For example, such an approach could be applied to model the effects of pressure anisotropy-driven whistler waves in the outflow regions of magnetic reconnection sites [184, 185].

Finally, we would be remiss not to reiterate that there are many different potential data-driven avenues towards addressing the lack of generally applicable collisionless closures, and the sparse regression methodology which is the focus of this thesis is but one of them. Many of the alternatives we have already discussed in Chapter 3, such as symbolic regression, Bayesian optimisation and various neural network-based approaches. In the latter category, methods which are more constrained towards physical correctness, such as FNOs, PINNs and BEACONS [119, 120, 124] (the latter two also discussed above in the context of fluid-code closure evaluation), are especially of interest. Nevertheless, the methodology developed and used in this thesis remains a strong contender – chiefly due to its relative simplicity and interpretability, paving the way for further analytical study.

References

- [1] M.-B. Kallenrode, *Space Physics: An Introduction to Plasmas and Particles in the Heliosphere and Magnetospheres*, 3rd ed. Springer-Verlag Berlin, 2004, ISBN: 978-3-540-20617-0. DOI: 10.1007/978-3-662-09959-9.
- [2] W. Baumjohann and R. Treumann, *Basic Space Plasma Physics*, 3rd ed. World Scientific, Mar. 2022, ISBN: 978-981-12-5405-5. DOI: 10.1142/12771.
- [3] A. L. Peratt, ‘Advances in numerical modeling of astrophysical and space plasmas,’ *Astrophysics and Space Science*, vol. 242, pp. 93–163, 1996. DOI: 10.1007/BF00645112.
- [4] N. Jarosik et al., ‘Seven-year Wilkinson microwave anisotropy probe (WMAP*) observations: Sky maps, systematic errors, and basic results,’ *The Astrophysical Journal Supplement Series*, vol. 192, no. 2, p. 14, Jan. 2011. DOI: 10.1088/0067-0049/192/2/14.
- [5] Planck Collaboration et al., ‘Planck 2013 results. I. Overview of products and scientific results,’ *Astronomy & Astrophysics*, vol. 571, A1, 2014. DOI: 10.1051/0004-6361/201321529.
- [6] G. H. Hardy, *A mathematician’s apology*, Nov. 1940. DOI: 10.1017/CB09781107295599.
- [7] L. S. Sparke and J. S. Gallagher, *Galaxies in the Universe: An Introduction*, 2nd ed. Cambridge University Press, 2007.
- [8] B. Ryden and R. W. Pogge, *Interstellar and Intergalactic Medium*. Cambridge University Press, 2021. DOI: 10.1017/9781108781596.
- [9] M. A. Abramowicz and P. C. Fragile, ‘Foundations of black hole accretion disk theory,’ *Living Reviews in Relativity*, vol. 16, p. 1, 2013. DOI: 10.12942/lrr-2013-1.
- [10] B. Liu and E. Qiao, ‘Accretion around black holes: The geometry and spectra,’ *iScience*, vol. 25, no. 1, p. 103 544, 2022, ISSN: 2589-0042. DOI: 10.1016/j.isci.2021.103544.
- [11] J. Nättilä, ‘Radiative plasma simulations of black hole accretion flow coronae in the hard and soft states,’ *Nature Communications*, vol. 15, p. 7026, 2024. DOI: 10.1038/s41467-024-51257-1.

- [12] B. M. S. Hansen and M. Lyutikov, ‘Radio and X-ray signatures of merging neutron stars,’ *Monthly Notices of the Royal Astronomical Society*, vol. 322, no. 4, pp. 695–701, Apr. 2001, ISSN: 0035-8711. DOI: 10.1046/j.1365-8711.2001.04103.x.
- [13] A. M. Beloborodov and C. Thompson, ‘Corona of magnetars,’ *The Astrophysical Journal*, vol. 657, no. 2, p. 967, Mar. 2007. DOI: 10.1086/508917.
- [14] C. D. Arrowsmith et al., ‘Laboratory realization of relativistic pair-plasma beams,’ *Nature Communications*, vol. 15, p. 5029, 2024. DOI: 10.1038/s41467-024-49346-2.
- [15] X.-L. Zhu, T.-P. Yu, Z.-M. Sheng, Y. Yin, I. C. E. Turcu and A. Pukhov, ‘Dense GeV electron–positron pairs generated by lasers in near-critical-density plasmas,’ *Nature Communications*, vol. 7, p. 13686, 2016. DOI: 10.1038/ncomms13686.
- [16] K. Qu, S. Meuren and N. J. Fisch, ‘Creating pair plasmas with observable collective effects,’ *Plasma Physics and Controlled Fusion*, vol. 65, no. 3, p. 034007, Feb. 2023. DOI: 10.1088/1361-6587/acb080.
- [17] H. Chen and F. Fiuza, ‘Perspectives on relativistic electron–positron pair plasma experiments of astrophysical relevance using high-power lasers,’ *Physics of Plasmas*, vol. 30, no. 2, p. 020601, 2023, ISSN: 1070-664X. DOI: 10.1063/5.0134819.
- [18] E. Esarey, C. B. Schroeder and W. P. Leemans, ‘Physics of laser-driven plasma-based electron accelerators,’ *Reviews of Modern Physics*, vol. 81, pp. 1229–1285, 3 Aug. 2009. DOI: 10.1103/RevModPhys.81.1229.
- [19] J. Faure, D. Gustas, D. Guénot, A. Vernier, F. Böhle, M. Ouillé, S. Haessler, R. Lopez-Martens and A. Lifschitz, ‘A review of recent progress on laser-plasma acceleration at kHz repetition rate,’ *Plasma Physics and Controlled Fusion*, vol. 61, no. 1, p. 014012, Nov. 2018. DOI: 10.1088/1361-6587/aae047.
- [20] M. Labat et al., ‘Seeded free-electron laser driven by a compact laser plasma accelerator,’ *Nature Photonics*, vol. 17, pp. 150–156, 2023. DOI: 10.1038/s41566-022-01104-w.
- [21] D. Mašlářová, A. Hansson, M. Luo, V. Horný, J. Ferri, I. Pusztai and T. Fülöp, ‘Batch Bayesian optimization of attosecond betatron pulses from laser wakefield acceleration,’ *Communications Physics*, 2026. DOI: 10.1038/s42005-026-02542-6.
- [22] P.-H. Rebut and the ITER Joint Central Team, ‘ITER: The first experimental fusion reactor,’ *Fusion Engineering and Design*, vol. 30, no. 1, pp. 85–118, 1995, ISSN: 0920-3796. DOI: [https://doi.org/10.1016/0920-3796\(94\)00403-T](https://doi.org/10.1016/0920-3796(94)00403-T).
- [23] J. Paméla, E. R. Solano and JET EFDA Contributors, ‘Overview of JET results,’ *Nuclear Fusion*, vol. 43, no. 12, p. 1540, Dec. 2003. DOI: 10.1088/0029-5515/43/12/002.

- [24] S. Wu, ‘An overview of the EAST project,’ *Fusion Engineering and Design*, vol. 82, no. 5, pp. 463–471, 2007, Proceedings of the 24th Symposium on Fusion Technology, ISSN: 0920-3796. DOI: 10.1016/j.fusengdes.2007.03.012.
- [25] Y. Kamada, P. Barabaschi, S. Ishida, the JT-60SA Team and JT-60SA Research Plan Contributors, ‘Progress of the JT-60SA project,’ *Nuclear Fusion*, vol. 53, no. 10, p. 104010, Sep. 2013. DOI: 10.1088/0029-5515/53/10/104010.
- [26] A. J. Creely et al., ‘Overview of the SPARC tokamak,’ *Journal of Plasma Physics*, vol. 86, no. 5, p. 865860502, 2020. DOI: 10.1017/S0022377820001257.
- [27] F. S. B. Anderson, A. F. Almagri, D. T. Anderson, P. G. Matthews, J. N. Talmadge and J. L. Shohet, ‘The helically symmetric experiment, (HSX) goals, design and status,’ *Fusion Technology*, vol. 27, no. 3T, pp. 273–277, 1995. DOI: 10.13182/FST95-A11947086.
- [28] T. Klinger et al., ‘Overview of first Wendelstein 7-X high-performance operation,’ *Nuclear Fusion*, vol. 59, no. 11, p. 112004, Jun. 2019. DOI: 10.1088/1741-4326/ab03a7.
- [29] A. H. Boozer, ‘Stellarators as a fast path to fusion,’ *Nuclear Fusion*, vol. 61, no. 9, p. 096024, Aug. 2021. DOI: 10.1088/1741-4326/ac170f.
- [30] S. Meschini, F. Laviano, F. Ledda, D. Pettinari, R. Testoni, D. Torsello and B. Panella, ‘Review of commercial nuclear fusion projects,’ *Frontiers in Energy Research*, vol. 11, 2023, ISSN: 2296-598X. DOI: 10.3389/fenrg.2023.1157394.
- [31] E. R. Sadik-Zada, A. Gatto and Y. Weißnicht, ‘Back to the future: Revisiting the perspectives on nuclear fusion and juxtaposition to existing energy sources,’ *Energy*, vol. 290, p. 129150, 2024, ISSN: 0360-5442. DOI: 10.1016/j.energy.2023.129150.
- [32] C. D. Beidler et al., ‘Demonstration of reduced neoclassical energy transport in Wendelstein 7-X,’ *Nature*, vol. 596, pp. 221–226, 2021. DOI: 10.1038/s41586-021-03687-w.
- [33] J. Degraeve et al., ‘Magnetic control of tokamak plasmas through deep reinforcement learning,’ *Nature*, vol. 602, pp. 414–419, 2022. DOI: 10.1038/s41586-021-04301-9.
- [34] M. Hoppe, I. Ekmark, E. Berger and T. Fülöp, ‘Runaway electron generation during tokamak start-up,’ *Journal of Plasma Physics*, vol. 88, no. 3, p. 905880317, 2022. DOI: 10.1017/S002237782200054X.
- [35] C. Shi and Z. Gan, ‘Editorial: Solar wind turbulence: Its origins, evolution, and impacts,’ *Frontiers in Astronomy and Space Sciences*, vol. 11, 2024, ISSN: 2296-987X. DOI: 10.3389/fspas.2024.1513779.
- [36] D. Telloni et al., ‘Evolution of solar wind turbulence from 0.1 to 1 au during the first Parker Solar Probe–Solar Orbiter radial alignment,’ *The Astrophysical Journal Letters*, vol. 912, no. 2, p. L21, May 2021. DOI: 10.3847/2041-8213/abf7d1.

- [37] S. D. Bale et al., ‘Interchange reconnection as the source of the fast solar wind within coronal holes,’ *Nature*, vol. 618, pp. 252–256, 2023. DOI: 10.1038/s41586-023-05955-3.
- [38] R. Schwenn, ‘Space weather: The solar perspective,’ *Living Reviews in Solar Physics*, vol. 3, p. 2, 2006. DOI: 10.12942/lrsp-2006-2.
- [39] N. Buzulukova and B. Tsurutani, ‘Space weather: From solar origins to risks and hazards evolving in time,’ *Frontiers in Astronomy and Space Sciences*, vol. 9, 2022, ISSN: 2296-987X. DOI: 10.3389/fspas.2022.1017103.
- [40] D. Xue, L. Wu, T. Xu, C.-L. Wu, Z. Wang and Z. He, ‘Space weather effects on transportation systems: A review of current understanding and future outlook,’ *Space Weather*, vol. 22, no. 12, e2024SW004055, 2024. DOI: 10.1029/2024SW004055.
- [41] D. H. Boteler, ‘A 21st century view of the March 1989 magnetic storm,’ *Space Weather*, vol. 17, no. 10, pp. 1427–1441, 2019. DOI: 10.1029/2019SW002278.
- [42] R. C. Carrington, ‘Description of a singular appearance seen in the Sun on September 1, 1859,’ *Monthly Notices of the Royal Astronomical Society*, vol. 20, pp. 13–15, 1859.
- [43] R. Hodgson, ‘On a curious appearance seen in the Sun,’ *Monthly Notices of the Royal Astronomical Society*, vol. 20, no. 1, pp. 15–16, Nov. 1859, ISSN: 0035-8711. DOI: 10.1093/mnras/20.1.15a.
- [44] B. T. Tsurutani, W. D. Gonzalez, G. S. Lakhina and S. Alex, ‘The extreme magnetic storm of 1–2 September 1859,’ *Journal of Geophysical Research: Space Physics*, vol. 108, A7, 2003. DOI: 10.1029/2002JA009504.
- [45] Lloyd’s of London and Atmospheric and Environmental Research, Inc., *Solar storm risk to the north american electrical grid*, 2013. [Online]. Available: <https://assets.lloyds.com/assets/pdf-solar-storm-risk-to-the-north-american-electric-grid/1/pdf-Solar-Storm-Risk-to-the-North-American-Electric-Grid.pdf>.
- [46] P. M. Bellan, *Fundamentals of Plasma Physics*. Cambridge University Press, 2006.
- [47] S. Chapman, ‘The kinetic theory of simple and composite monatomic gases; viscosity, thermal conduction, and diffusion,’ *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 93, pp. 1–20, 1916.
- [48] D. Enskog, ‘Kineticische theorie der vorgänge in mässig verdünnten gasen.’ Ph.D. dissertation, Uppsala universitet, 1917.
- [49] S. I. Braginskii, ‘Transport phenomena in a completely ionized two-temperature plasma,’ *Sov. Phys. JETP*, vol. 6, no. 33, pp. 358–369, 1958.
- [50] S. I. Braginskii, ‘Transport processes in a plasma,’ *Reviews of Plasma Physics*, vol. 1, pp. 205–309, 1965.

- [51] S. Chapman and T. G. Cowling, *The mathematical theory of non-uniform gases: An account of the kinetic theory of viscosity, thermal conduction and diffusion in gases*. Cambridge, England: Cambridge University Press, Jan. 1991.
- [52] F. Rosenblatt, 'The perceptron: A probabilistic model for information storage and organization in the brain,' *Psychological Review*, vol. 65, no. 6, pp. 386–408, 1958. DOI: 10.1037/h0042519.
- [53] F. Rosenblatt, *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books, 1962.
- [54] K. Fukushima, 'Neocognitron: A Self-organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position,' *Biological Cybernetics*, vol. 36, pp. 193–202, 1980. DOI: 10.1007/BF00344251.
- [55] R. Tibshirani, 'Regression shrinkage and selection via the lasso,' *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996. DOI: 10.1111/j.2517-6161.1996.tb02080.x.
- [56] T. Hastie, R. Tibshirani and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York, NY, USA: Springer, 2009. DOI: 10.1007/978-0-387-84858-7.
- [57] G. James, D. Witten, T. Hastie and R. Tibshirani, *An Introduction to Statistical Learning*. New York, NY, USA: Springer, 2013. DOI: 10.1007/978-1-0716-1418-1.
- [58] S. L. Brunton, J. L. Proctor and J. N. Kutz, 'Discovering governing equations from data by sparse identification of nonlinear dynamical systems,' *Proceedings of the National Academy of Sciences of the United States of America*, vol. 113, no. 15, pp. 3932–3937, 2016. DOI: 10.1073/pnas.1517384113.
- [59] S. H. Rudy, S. L. Brunton, J. L. Proctor and J. N. Kutz, 'Data-driven discovery of partial differential equations,' *Science Advances*, vol. 3, no. 4, Apr. 2017, ISSN: 2375-2548. DOI: 10.1126/sciadv.1602614.
- [60] H. Schaeffer, 'Learning partial differential equations via data discovery and sparse optimization,' *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 473, no. 2197, p. 20160446, 2017. DOI: 10.1098/rspa.2016.0446.
- [61] T. H. Stix, *Waves in Plasmas*, 1st ed. American Institute of Physics, 1992, ISBN: 978-0-88318-859-0.
- [62] D. G. Swanson, *Plasma Waves*, 2nd ed. Institute of Physics Publishing, 2003, ISBN: 0-7503-0927-X.
- [63] L. V. Ahlfors, *Complex Analysis*, 3rd ed. New York: McGraw-Hill, Inc., 1979, ch. 4: Complex Integration, pp. 101–172.
- [64] E. R. Ingelsten, 'Quantum optics with giant atoms in 2D structured environments,' M.S. thesis, Chalmers University of Technology, 2023. [Online]. Available: <http://hdl.handle.net/20.500.12380/309263>.

- [65] L. Landau, ‘On the vibrations of the electronic plasma,’ *Journal of Physics, USSR*, vol. 10, p. 27, 25 1946.
- [66] S. P. Gary, *Theory of Space Plasma Microinstabilities*. Cambridge, UK: Cambridge University Press, 1993. DOI: 10.1017/CB09780511551512.
- [67] T. M. O’Neil and J. H. Malmberg, ‘Transition of the dispersion roots from beam-type to Landau-type solutions,’ *The Physics of Fluids*, vol. 11, no. 8, pp. 1754–1760, Aug. 1968, ISSN: 0031-9171. DOI: 10.1063/1.1692190.
- [68] F. I. Parra, *Braginskii fluid equations*, Lecture notes, 2019. [Online]. Available: <https://www-thphys.physics.ox.ac.uk/people/FelixParra/CollisionalPlasmaPhysics/notes/lecIV-braginskii.pdf>.
- [69] G. F. Chew, M. L. Goldberger and F. E. Low, ‘The Boltzmann equation and the one-fluid hydromagnetic equations in the absence of particle collisions,’ *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 236, no. 1204, pp. 112–118, 1956. DOI: 10.1098/rspa.1956.0116.
- [70] P. Hunana, A. Tenerani, G. P. Zank, E. Khomenko, M. L. Goldstein, G. M. Webb, P. S. Cally, M. Collados, M. Velli and L. Adhikari, ‘An introductory guide to fluid models with anisotropic temperatures. Part 1. CGL description and collisionless fluid hierarchy,’ *Journal of Plasma Physics*, vol. 85, no. 6, p. 205 850 602, 2019. DOI: 10.1017/S0022377819000801.
- [71] H. Alfvén, ‘Existence of electromagnetic-hydrodynamic waves,’ *Nature*, vol. 150, pp. 405–406, 1942. DOI: 10.1038/150405d0.
- [72] H. Alfvén, ‘On the existence of electromagnetic-hydrodynamic waves,’ *Arkiv för matematik, astronomi och fysik*, vol. 29B, no. 2, pp. 1–7, 1943.
- [73] J. Egedal, W. Fox, N. Katz, M. Porkolab, M. Øieroset, R. P. Lin, W. Daughton and J. F. Drake, ‘Evidence and theory for trapped electrons in guide field magnetotail reconnection,’ *Journal of Geophysical Research: Space Physics*, vol. 113, no. A12, p. 207, 2008. DOI: 10.1029/2008JA013520.
- [74] A. Le, J. Egedal, W. Daughton, W. Fox and N. Katz, ‘Equations of state for collisionless guide-field reconnection,’ *Physical Review Letters*, vol. 102, p. 085 001, 8 Feb. 2009. DOI: 10.1103/PhysRevLett.102.085001.
- [75] J. Egedal, ‘On a plasma sheath with a small normal magnetic field separating regions of oppositely directed magnetic field,’ *Physics of Plasmas*, vol. 30, no. 11, p. 112 903, Nov. 2023, ISSN: 1070-664X. DOI: 10.1063/5.0170212.
- [76] J. Egedal, ‘The adiabatic 1d kinetic equilibrium of the electron diffusion region during anti-parallel magnetic reconnection,’ *Geophysical Research Letters*, vol. 51, no. 10, e2024GL108895, 2024. DOI: 10.1029/2024GL108895.

- [77] J. Egedal, ‘Analytical approximations for the current sheet adiabatic invariant,’ *Physics of Plasmas*, vol. 32, no. 4, p. 042 101, Apr. 2025, ISSN: 1070-664X. DOI: 10.1063/5.0253738.
- [78] G. W. Hammett and F. W. Perkins, ‘Fluid moment models for Landau damping with application to the ion-temperature-gradient instability,’ *Physical Review Letters*, vol. 64, pp. 3019–3022, 25 Jun. 1990. DOI: 10.1103/PhysRevLett.64.3019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.64.3019>.
- [79] R. D. Hazeltine, ‘Transport theory in the collisionless limit,’ *Physics of Plasmas*, vol. 5, no. 9, pp. 3282–3286, Sep. 1998, ISSN: 1070-664X. DOI: 10.1063/1.872996.
- [80] P. B. Snyder, G. W. Hammett and W. Dorland, ‘Landau fluid models of collisionless magnetohydrodynamics,’ *Physics of Plasmas*, vol. 4, no. 11, pp. 3974–3985, Nov. 1997, ISSN: 1070-664X. DOI: 10.1063/1.872517.
- [81] P. Hunana, A. Tenerani, G. P. Zank, M. L. Goldstein, G. M. Webb, E. Khomenko, M. Collados, P. S. Cally, L. Adhikari and M. Velli, ‘An introductory guide to fluid models with anisotropic temperatures. Part 2. Kinetic theory, Padé approximants and Landau fluid closures,’ *Journal of Plasma Physics*, vol. 85, no. 6, p. 205 850 603, 2019. DOI: 10.1017/S0022377819000850.
- [82] J. J. Ramos, ‘Fluid formalism for collisionless magnetized plasmas,’ *Physics of Plasmas*, vol. 12, no. 5, p. 052 102, Apr. 2005, ISSN: 1070-664X. DOI: 10.1063/1.1884128.
- [83] P. Goswami, T. Passot and P. L. Sulem, ‘A Landau fluid model for warm collisionless plasmas,’ *Physics of Plasmas*, vol. 12, no. 10, p. 102 109, Oct. 2005, ISSN: 1070-664X. DOI: 10.1063/1.2096582.
- [84] T. Passot and P. L. Sulem, ‘Collisionless magnetohydrodynamics with gyrokinetic effects,’ *Physics of Plasmas*, vol. 14, no. 8, p. 082 502, Aug. 2007, ISSN: 1070-664X. DOI: 10.1063/1.2751601.
- [85] T. Passot, P. L. Sulem and P. Hunana, ‘Extending magnetohydrodynamics to the slow dynamics of collisionless plasmas,’ *Physics of Plasmas*, vol. 19, no. 8, p. 082 113, Aug. 2012, ISSN: 1070-664X. DOI: 10.1063/1.4746092.
- [86] P. L. Sulem and T. Passot, ‘Landau fluid closures with nonlinear large-scale finite larmor radius corrections for collisionless plasmas,’ *Journal of Plasma Physics*, vol. 81, no. 1, p. 325 810 103, 2015. DOI: 10.1017/S0022377814000671.
- [87] P. Hunana, G. P. Zank, M. Laurenza, A. Tenerani, G. M. Webb, M. L. Goldstein, M. Velli and L. Adhikari, ‘New closures for more precise modeling of Landau damping in the fluid framework,’ *Phys. Rev. Lett.*, vol. 121, p. 135 101, 13 Sep. 2018. DOI: 10.1103/PhysRevLett.121.135101.
- [88] A. Bravais, *Analyse Mathématique sur les Probabilités des Erreurs de Situation d’un Point*. Paris: Impr. Royale, 1844.

- [89] F. Galton, 'Regression towards mediocrity in hereditary stature,' *The Journal of the Anthropological Institute of Great Britain and Ireland*, vol. 15, pp. 246–263, 1886, ISSN: 09595295, 23972564. Accessed: 16th Mar. 2026. [Online]. Available: <http://www.jstor.org/stable/2841583>.
- [90] K. Pearson, 'VII. Mathematical contributions to the theory of evolution.—III. Regression, heredity, and panmixia,' *Philosophical Transactions of the Royal Society of London, Series A: Containing Papers of a Mathematical or Physical Character*, no. 187, pp. 253–318, Dec. 1896, ISSN: 0264-3952. DOI: 10.1098/rsta.1896.0007.
- [91] J. M. Stanton, 'Galton, Pearson, and the peas: A brief history of linear regression for statistics instructors,' *Journal of Statistics Education*, vol. 9, no. 3, 2001. DOI: 10.1080/10691898.2001.11910537.
- [92] W. S. McCulloch and W. Pitts, 'A logical calculus of the ideas immanent in nervous activity,' *Bulletin of Mathematical Biophysics*, vol. 5, pp. 115–133, 1943. DOI: 10.1007/BF02478259.
- [93] S. Linnainmaa, 'Algoritmin kumulatiivinen pyöristyvirhe yksittäisten pyöristysvirheiden Taylor-kehitemänä,' M.S. thesis, University of Helsinki, 1970. [Online]. Available: <https://people.idsia.ch/~juergen/linnainmaa1970thesis.pdf>.
- [94] P. J. Werbos, 'Applications of advances in nonlinear sensitivity analysis,' in *Proceedings of the 10th IFIP Conference (New York City, USA, August 31–September 4, 1981)*, New York: Springer-Verlag, 1982, pp. 762–770.
- [95] P. J. Werbos, 'Generalization of backpropagation with application to a recurrent gas market model,' *Neural Networks*, vol. 1, no. 4, pp. 339–356, 1988, ISSN: 0893-6080. DOI: 10.1016/0893-6080(88)90007-X.
- [96] D. Rumelhart, G. Hinton and R. Williams, 'Learning representations by back-propagating errors,' *Nature*, vol. 323, pp. 533–536, 1986. DOI: 10.1038/323533a0.
- [97] S. S. Liew, M. Khalil-Hani and R. Bakhteri, 'Bounded activation functions for enhanced training stability of deep neural networks on visual pattern recognition problems,' *Neurocomputing*, vol. 216, pp. 718–734, 2016, ISSN: 0925-2312. DOI: 10.1016/j.neucom.2016.08.037.
- [98] K. Hornik, M. Stinchcombe and H. White, 'Multilayer feedforward networks are universal approximators,' *Neural Networks*, vol. 2, no. 5, pp. 359–366, 1989, ISSN: 0893-6080. DOI: [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0893608089900208>.
- [99] R. Venkatesan and B. Li, *Convolutional Neural Networks in Visual Computing: A Concise Guide*, 1st ed. CRC Press, 2017. DOI: 10.4324/9781315154282.
- [100] V. E. Balas, R. Kumar and R. Srivastava, *Valentina E. Balas, Raghvendra Kumar, Rajshree Srivastava*, 1st ed. Springer Cham, 2019, ISBN: 978-3-030-32643-2. DOI: 10.1007/978-3-030-32644-9.

- [101] A. Daigavane, B. Ravindran and G. Aggarwal, ‘Understanding convolutions on graphs,’ *Distill*, 2021, <https://distill.pub/2021/understanding-gnns>. DOI: 10.23915/distill.00032.
- [102] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser and I. Polosukhin, ‘Attention is all you need,’ in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan and R. Garnett, Eds., vol. 30, Curran Associates, Inc., 2017.
- [103] D. Bahdanau, K. Cho and Y. Bengio, *Neural machine translation by jointly learning to align and translate*, Original version uploaded in 2014., 2016. arXiv: 1409.0473 [cs.CL].
- [104] T. Lin, Y. Wang, X. Liu and X. Qiu, *A survey of transformers*, 2021. arXiv: 2106.04554 [cs.LG].
- [105] S. Islam, H. Elmekki, A. Elsebai, J. Bentahar, N. Drawel, G. Rjoub and W. Pedrycz, *A comprehensive survey on applications of transformers for deep learning tasks*, 2023. arXiv: 2306.07303 [cs.LG].
- [106] C. Ma, B. Zhu, X.-Q. Xu and W. Wang, ‘Machine learning surrogate models for Landau fluid closure,’ *Physics of Plasmas*, vol. 27, no. 4, p. 042502, Apr. 2020, ISSN: 1070-664X. DOI: 10.1063/1.5129158.
- [107] L. Wang, X. Q. Xu, B. Zhu, C. Ma and Y.-a. Lei, ‘Deep learning surrogate model for kinetic Landau-fluid closure with collision,’ *AIP Advances*, vol. 10, no. 7, p. 075108, Jul. 2020, ISSN: 2158-3226. DOI: 10.1063/5.0010917.
- [108] R. Maulik, N. A. Garland, J. W. Burby, X.-Z. Tang and P. Balaprakash, ‘Neural network representability of fully ionized plasma fluid model closures,’ *Physics of Plasmas*, vol. 27, no. 7, p. 072106, Jul. 2020, ISSN: 1070-664X. DOI: 10.1063/5.0006457.
- [109] C. Lamy, B. Dubroca, P. Nicolai, V. Tikhonchuk and J.-L. Feugeas, ‘Modeling of electron nonlocal transport in plasmas using artificial neural networks,’ *Phys. Rev. E*, vol. 105, p. 055201, 5 May 2022. DOI: 10.1103/PhysRevE.105.055201.
- [110] F. Miniati and G. Gregori, ‘Learning transport processes with machine intelligence,’ *Scientific Reports*, vol. 12, no. 1, p. 11709, Jul. 2022, ISSN: 2045-2322. DOI: 10.1038/s41598-022-15416-y.
- [111] Y. Qin, J. Ma, M. Jiang, C. Dong, H. Fu, L. Wang, W. Cheng and Y. Jin, ‘Data-driven modeling of Landau damping by physics-informed neural networks,’ *Physical Review Research*, vol. 5, p. 033079, 3 Aug. 2023. DOI: 10.1103/PhysRevResearch.5.033079.
- [112] N. Barbour, W. Dorland, I. G. Abel and M. Landreman, ‘Machine-learning closure for vlasov-poisson dynamics in fourier-hermite space,’ *arXiv*, 2025, Accepted for publication in the Journal of Plasma Physics. eprint: 2504.13313 (physics.plasm-ph). [Online]. Available: <https://arxiv.org/abs/2504.13313>.

- [113] Z. Huang, C. Dong and L. Wang, ‘Machine-learning heat flux closure for multi-moment fluid modeling of nonlinear Landau damping,’ *Proceedings of the National Academy of Sciences*, vol. 122, no. 11, e2419073122, 2025. DOI: 10.1073/pnas.2419073122.
- [114] M. Luo, C. Heaton, Y. Wang, D. Plummer, M. Fitzgerald, F. Miniati, S. M. Vinko and G. Gregori, *Time-embedded convolutional neural networks for modeling plasma heat transport*, 2025. arXiv: 2509.06088 [physics.plasm-ph].
- [115] D. D. Carvalho, D. R. Ferreira and L. O. Silva, ‘Learning the dynamics of a one-dimensional plasma model with graph neural networks,’ *Machine Learning: Science and Technology*, vol. 5, no. 2, p. 025 048, May 2024. DOI: 10.1088/2632-2153/ad4ba6.
- [116] M. Mlinarević, G. K. Holt and A. Agnello, *Particle-based plasma simulation using a graph neural network*, 2025. arXiv: 2503.00274 [physics.plasm-ph].
- [117] G. Vigot, B. Cuenot, O. Vermorel and M. Bauerheim, ‘Graph neural networks for computational plasma physics on unstructured grids: Application to approximate the Poisson equation for Hall-effect thrusters modeling,’ *Journal of Electric Propulsion*, vol. 4, p. 46, 2025. DOI: 10.1007/s44205-025-00146-w.
- [118] B. Sanchez-Lengeling, E. Reif, A. Pearce and A. B. Wiltschko, ‘A gentle introduction to graph neural networks,’ *Distill*, 2021, <https://distill.pub/2021/gnn-intro>. DOI: 10.23915/distill.00033.
- [119] M. Raissi, P. Perdikaris and G. Karniadakis, ‘Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,’ *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019, ISSN: 0021-9991. DOI: 10.1016/j.jcp.2018.10.045.
- [120] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart and A. Anandkumar, *Fourier neural operator for parametric partial differential equations*, 2021. arXiv: 2010.08895 [cs.LG].
- [121] V. Duruisseaux, J. Kossaifi and A. Anandkumar, *Fourier neural operators explained: A practical perspective*, 2026. arXiv: 2512.01421 [cs.LG].
- [122] Y. Liu, H. Fu, L. Wang, C. Dong, S. Wei, Y. Xi, C. Shang and Y. Qin, ‘Data-driven modeling of electrostatic turbulence by physics-informed fourier neural operator,’ *Machine Learning: Science and Technology*, vol. 6, no. 4, p. 045 050, Nov. 2025. DOI: 10.1088/2632-2153/ae19cd.
- [123] V. Gopakumar, S. Pamela, L. Zanisi, Z. Li, A. Anandkumar and the MAST Team, *Fourier neural operator for plasma modelling*, 2023. arXiv: 2302.06542 [physics.plasm-ph].
- [124] J. Gorard, A. Hakim and J. Juno, *BEACONS: Bounded-error, algebraically-composable neural solvers for partial differential equations*, 2026. arXiv: 2602.14853 [cs.LG].

- [125] M. C. McGrae-Menge, J. R. Pierce, F. Fiuza and E. P. Alves, ‘Embedding physical symmetries into machine-learned reduced plasma physics models via data augmentation,’ *Phys. Rev. Res.*, vol. 8, p. 013 200, 1 Feb. 2026. DOI: 10.1103/9fq1-vpqz.
- [126] J. Moćkus, ‘On Bayesian methods for seeking the extremum,’ in *Optimization Techniques IFIP Technical Conference Novosibirsk, July 1–7, 1974*, G. I. Marchuk, Ed., Berlin, Heidelberg: Springer Berlin Heidelberg, 1975, pp. 400–404, ISBN: 978-3-540-37497-8.
- [127] J. Mockus, ‘The Bayesian approach to global optimization,’ in *System Modeling and Optimization*, R. F. Drenick and F. Kozin, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 1982, pp. 473–481, ISBN: 978-3-540-39459-4.
- [128] R. Garnett, *Bayesian Optimization*. Cambridge University Press, 2023.
- [129] Y. K. Wakabayashi, T. Otsuka, Y. Krockenberger, H. Sawada, Y. Taniyasu and H. Yamamoto, ‘Bayesian optimization with experimental failure for high-throughput materials growth,’ *npj Computational Materials*, vol. 8, p. 180, 2022. DOI: 10.1038/s41524-022-00859-8.
- [130] R. J. Hickman, M. Aldeghi, F. Häse and A. Aspuru-Guzik, ‘Bayesian optimization with known experimental and design constraints for chemistry applications,’ *Digital Discovery*, vol. 1, pp. 732–744, 5 2022. DOI: 10.1039/D2DD00028H.
- [131] X. Wang, Y. Jin, S. Schmitt and M. Olhofer, *Recent advances in Bayesian optimization*, 2022. arXiv: 2206.03301 [cs.LG].
- [132] A. Ferran Pousa, S. Jalas, M. Kirchen, A. Martinez de la Ossa, M. Thévenet, S. Hudson, J. Larson, A. Huebl, J.-L. Vay and R. Lehe, ‘Bayesian optimization of laser-plasma accelerators assisted by reduced physical models,’ *Phys. Rev. Accel. Beams*, vol. 26, p. 084 601, 8 Aug. 2023. DOI: 10.1103/PhysRevAccelBeams.26.084601.
- [133] I. Pusztai, I. Ekmark, H. Bergström, P. Halldestam, P. Jansson, M. Hoppe, O. Vallhagen and T. Fülöp, ‘Bayesian optimization of massive material injection for disruption mitigation in tokamaks,’ *Journal of Plasma Physics*, vol. 89, no. 2, p. 905 890 204, 2023. DOI: 10.1017/S0022377823000193.
- [134] I. Ekmark, M. Hoppe, T. Fülöp, P. Jansson, L. Antonsson, O. Vallhagen and I. Pusztai, ‘Fluid and kinetic studies of tokamak disruptions using Bayesian optimization,’ *Journal of Plasma Physics*, vol. 90, no. 3, p. 905 900 306, 2024. DOI: 10.1017/S0022377824000606.
- [135] J. R. Koza, *Genetic Programming: On the programming of computers by means of natural selection*. MIT Press, 1992.
- [136] J. Bongard and H. Lipson, ‘Automated reverse engineering of nonlinear dynamical systems,’ *Proceedings of the National Academy of Sciences*, vol. 104, no. 24, pp. 9943–9948, 2007. DOI: 10.1073/pnas.0609476104.

- [137] M. Schmidt and H. Lipson, ‘Distilling free-form natural laws from experimental data,’ *Science*, vol. 324, no. 5923, pp. 81–85, 2009. DOI: 10.1126/science.1165893.
- [138] N. Makke and S. Chawla, ‘Interpretable scientific discovery with symbolic regression: A review,’ *Artificial Intelligence Review*, vol. 57, no. 1, p. 2, Jan. 2024, ISSN: 1573-7462. DOI: 10.1007/s10462-023-10622-0.
- [139] R. Dubčáková, ‘Eureqa: Software review,’ *Genetic Programming and Evolvable Machines*, vol. 12, pp. 173–178, 2011. DOI: 10.1007/s10710-010-9124-z.
- [140] W. L. Cava, T. R. Singh, J. Taggart, S. Suri and J. H. Moore, *Learning concise representations for regression by evolving networks of trees*, 2019. arXiv: 1807.00981 [cs.NE].
- [141] W.-X. Wang, R. Yang, Y.-C. Lai, V. Kovanis and C. Grebogi, ‘Predicting catastrophes in nonlinear dynamical systems by compressive sensing,’ *Phys. Rev. Lett.*, vol. 106, p. 154101, 15 Apr. 2011. DOI: 10.1103/PhysRevLett.106.154101.
- [142] E. Candes, J. Romberg and T. Tao, ‘Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,’ *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006. DOI: 10.1109/TIT.2005.862083.
- [143] L. Boninsegna, F. Nüske and C. Clementi, ‘Sparse learning of stochastic dynamical equations,’ *The Journal of Chemical Physics*, vol. 148, no. 24, p. 241723, Mar. 2018, ISSN: 0021-9606. DOI: 10.1063/1.5018409.
- [144] J.-C. Loiseau, B. R. Noack and S. L. Brunton, ‘Sparse reduced-order modelling: Sensor-based dynamics to full-state estimation,’ *Journal of Fluid Mechanics*, vol. 844, pp. 459–490, 2018. DOI: 10.1017/jfm.2018.147.
- [145] N. T. T. Nguyen, G. T. Kenyon and B. Yoon, ‘A regression algorithm for accelerated lattice QCD that exploits sparse inference on the D-Wave quantum annealer,’ *Scientific Reports*, vol. 10, p. 10915, 2020. DOI: 10.1038/s41598-020-67769-x.
- [146] L. Zanna and T. Bolton, ‘Data-driven equation discovery of ocean mesoscale closures,’ *Geophysical Research Letters*, vol. 47, no. 17, e2020GL088376, 2020. DOI: 10.1029/2020GL088376.
- [147] S. Beetham and J. Capecelatro, ‘Formulating turbulence closures using sparse regression with embedded form invariance,’ *Phys. Rev. Fluids*, vol. 5, p. 084611, 8 Aug. 2020. DOI: 10.1103/PhysRevFluids.5.084611.
- [148] M. Dam, M. Brøns, J. Juul Rasmussen, V. Naulin and J. S. Hesthaven, ‘Sparse identification of a predator-prey system from simulation data of a convection model,’ *Physics of Plasmas*, vol. 24, no. 2, p. 022310, Feb. 2017, ISSN: 1070-664X. DOI: 10.1063/1.4977057.

- [149] A. A. Kaptanoglu, K. D. Morgan, C. J. Hansen and S. L. Brunton, ‘Physics-constrained, low-dimensional models for magnetohydrodynamics: First-principles and data-driven approaches,’ *Phys. Rev. E*, vol. 104, p. 015 206, 1 Jul. 2021. DOI: 10.1103/PhysRevE.104.015206.
- [150] E. P. Alves and F. Fiuza, ‘Data-driven discovery of reduced plasma physics models from fully kinetic simulations,’ *Phys. Rev. Res.*, vol. 4, p. 033 192, 3 Sep. 2022. DOI: 10.1103/PhysRevResearch.4.033192.
- [151] A. A. Kaptanoglu, C. Hansen, J. D. Lore, M. Landreman and S. L. Brunton, ‘Sparse regression for plasma physics,’ *Physics of Plasmas*, vol. 30, no. 3, p. 033 906, Mar. 2023, ISSN: 1070-664X. DOI: 10.1063/5.0139039.
- [152] H. Schaeffer and S. G. McCalla, ‘Sparse model selection via integral terms,’ *Phys. Rev. E*, vol. 96, p. 023 302, 2 Aug. 2017. DOI: 10.1103/PhysRevE.96.023302.
- [153] D. A. Messenger and D. M. Bortz, ‘Weak sindy: Galerkin-based data-driven model selection,’ *Multiscale Modeling & Simulation*, vol. 19, no. 3, pp. 1474–1497, 2021. DOI: 10.1137/20M1343166.
- [154] D. A. Messenger and D. M. Bortz, ‘Weak sindy for partial differential equations,’ *Journal of Computational Physics*, vol. 443, p. 110 525, 2021, ISSN: 0021-9991. DOI: 10.1016/j.jcp.2021.110525.
- [155] N. M. Mangan, S. L. Brunton, J. L. Proctor and J. N. Kutz, ‘Inferring biological networks by sparse identification of nonlinear dynamics,’ *IEEE Transactions on Molecular, Biological, and Multi-Scale Communications*, vol. 2, no. 1, pp. 52–63, 2016. DOI: 10.1109/TMBMC.2016.2633265.
- [156] K. Kaheman, J. N. Kutz and S. L. Brunton, ‘SINDy-PI: A robust algorithm for parallel implicit sparse identification of nonlinear dynamics,’ *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 476, p. 20 200 279, 2242 2020. DOI: 10.1098/rspa.2020.0279.
- [157] P. Virtanen et al., ‘SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,’ *Nature Methods*, vol. 17, pp. 261–272, 2020. DOI: 10.1038/s41592-019-0686-2.
- [158] A. Hakim, J. Loverich and U. Shumlak, ‘A high resolution wave propagation scheme for ideal two-fluid plasma equations,’ *Journal of Computational Physics*, vol. 219, no. 1, pp. 418–442, 2006, ISSN: 0021-9991. DOI: 10.1016/j.jcp.2006.03.036.
- [159] A. Hakim, ‘Extended MHD modelling with the ten-moment equations,’ *Journal of Fusion Energy*, no. 27, pp. 36–43, 2008. DOI: 10.1007/s10894-007-9116-z.
- [160] R. D. Hazeltine and S. M. Mahajan, ‘Fluid description of relativistic, magnetized plasma,’ *The Astrophysical Journal*, vol. 567, no. 2, p. 1262, Mar. 2002. DOI: 10.1086/338696.

- [161] E. A. Johnson, *The relativistic Vlasov equation*, 2011. [Online]. Available: <https://www.danlj.org/eaj/math/summaries/relativity/vlasov.pdf>.
- [162] L. Tinti, G. Vujanovic, J. Noronha and U. Heinz, ‘Resummed hydrodynamic expansion for a plasma of particles interacting with fields,’ *Phys. Rev. D*, vol. 99, p. 016009, 1 Jan. 2019. DOI: 10.1103/PhysRevD.99.016009.
- [163] E. R. Most, J. Noronha and A. A. Philippov, ‘Modelling general-relativistic plasmas with collisionless moments and dissipative two-fluid magnetohydrodynamics,’ *Monthly Notices of the Royal Astronomical Society*, vol. 514, no. 4, pp. 4989–5003, May 2022, ISSN: 0035-8711. DOI: 10.1093/mnras/stac1435.
- [164] J. Juno, G. Johnson, A. Philippov, A. Hakim, A. Chernoglazov and S. Zeng, *Modeling of relativistic plasmas with a conservative discontinuous galerkin method*, 2026. arXiv: 2602.17487 [astro-ph.HE].
- [165] I. Svenningsson, E. Yordanova, G. Cozzani, Y. V. Khotyaintsev and M. André, ‘Kinetic generation of whistler waves in the turbulent magnetosheath,’ *Geophysical Research Letters*, vol. 49, no. 15, e2022GL099065, 2022. DOI: 10.1029/2022GL099065.
- [166] I. Svenningsson, E. Yordanova, Y. V. Khotyaintsev, M. André, G. Cozzani and K. Steinvall, ‘Whistler waves in the quasi-parallel and quasi-perpendicular magnetosheath,’ *Journal of Geophysical Research: Space Physics*, vol. 129, no. 6, e2024JA032661, 2024. DOI: 10.1029/2024JA032661.
- [167] J.-H. Li et al., ‘Direct observations of cross-scale wave-particle energy transfer in space plasmas,’ *Science Advances*, vol. 11, no. 6, eadr8227, 2025. DOI: 10.1126/sciadv.adr8227.
- [168] M. F. Ivarsen et al., ‘Characteristic e-region plasma signature of magnetospheric wave-particle interactions,’ *Physical Review Letters*, vol. 134, p. 145201, 14 Apr. 2025. DOI: 10.1103/PhysRevLett.134.145201.
- [169] S. F. Tigik, D. B. Graham and Y. V. Khotyaintsev, ‘Electron-scale energy transfer due to lower hybrid waves during asymmetric reconnection,’ *Journal of Geophysical Research: Space Physics*, vol. 130, no. 4, e2024JA033503, 2025. DOI: 10.1029/2024JA033503.
- [170] J. L. Burch et al., ‘Electron-scale measurements of magnetic reconnection in space,’ *Science*, vol. 352, no. 6290, aaf2939, 2016. DOI: 10.1126/science.aaf2939.
- [171] E. Yordanova et al., ‘Electron scale structures and magnetic reconnection signatures in the turbulent magnetosheath,’ *Geophysical Research Letters*, vol. 43, no. 12, pp. 5969–5978, 2016. DOI: 10.1002/2016GL069191.
- [172] M. Oka, J. Birn, J. Egedal, F. Guo, R. E. Ergun, D. L. Turner, Y. Khotyaintsev, K.-J. Hwang, I. J. Cohen and J. F. Drake, ‘Particle acceleration by magnetic reconnection in geospace,’ *Space Science Reviews*, vol. 219, p. 75, 2023. DOI: 10.1007/s11214-023-01011-8.

- [173] L. Richard, L. Sorriso-Valvo, E. Yordanova, D. B. Graham and Y. V. Khotyaintsev, ‘Turbulence in magnetic reconnection jets from injection to sub-ion scales,’ *Physical Review Letters*, vol. 132, p. 105 201, 10 Mar. 2024. DOI: 10.1103/PhysRevLett.132.105201.
- [174] J. E. Stawarz et al., ‘The interplay between collisionless magnetic reconnection and turbulence,’ *Space Science Reviews*, vol. 220, p. 90, 2024. DOI: 10.1007/s11214-024-01124-8.
- [175] L. Richard, Y. V. Khotyaintsev, C. Norgren, K. Steinvall, D. B. Graham, J. Egedal, A. Vaivads and R. Nakamura, ‘Electron heating by parallel electric fields in magnetotail reconnection,’ *Physical Review Letters*, vol. 134, p. 215 201, 21 May 2025. DOI: 10.1103/PhysRevLett.134.215201.
- [176] Q. M. Lu, B. Lembege, J. B. Tao and S. Wang, ‘Perpendicular electric field in two-dimensional electron phase-holes: A parameter study,’ *Journal of Geophysical Research: Space Physics*, vol. 113, no. A11, 2008. DOI: <https://doi.org/10.1029/2008JA013693>.
- [177] Q. Lu, K. Huang, Y. Guan, S. Lu and R. Wang, ‘Energy dissipation in magnetic islands formed during magnetic reconnection,’ *The Astrophysical Journal*, vol. 954, no. 2, p. 146, Sep. 2023. DOI: 10.3847/1538-4357/acea86.
- [178] S. Wang, R. Wang, Q. Lu, H. Fu and S. Wang, ‘Direct evidence of secondary reconnection inside filamentary currents of magnetic flux ropes during magnetic reconnection,’ *Nature Communications*, vol. 11, no. 1, p. 3964, Aug. 2020, ISSN: 2041-1723. DOI: 10.1038/s41467-020-17803-3.
- [179] P. Grete, D. G. Vlaykov, W. Schmidt and D. R. G. Schleicher, ‘A nonlinear structural subgrid-scale closure for compressible MHD. II. a priori comparison on turbulence simulation data,’ *Physics of Plasmas*, vol. 23, no. 6, p. 062 317, Jun. 2016, ISSN: 1070-664X. DOI: 10.1063/1.4954304.
- [180] M. Nabavi and J. Kim, ‘Optimising subgrid-scale closures for spectral energy transfer in turbulent flows,’ *Journal of Fluid Mechanics*, vol. 982, A18, 2024. DOI: 10.1017/jfm.2024.101.
- [181] K. Jakhar, Y. Guan, R. Mojjani, A. Chattopadhyay and P. Hassanzadeh, ‘Learning closed-form equations for subgrid-scale closures from high-fidelity data: Promises and challenges,’ *Journal of Advances in Modeling Earth Systems*, vol. 16, no. 7, e2023MS003874, 2024. DOI: 10.1029/2023MS003874.
- [182] D. A. St-Onge, M. W. Kunz, J. Squire and A. A. Schekochihin, ‘Fluctuation dynamo in a weakly collisional plasma,’ *Journal of Plasma Physics*, vol. 86, no. 5, p. 905 860 503, 2020. DOI: 10.1017/S0022377820000860.

-
- [183] J. F. Drake, C. Pfrommer, C. S. Reynolds, M. Ruszkowski, M. Swisdak, A. Einarsson, T. Thomas, A. B. Hassam and G. T. Roberg-Clark, ‘Whistler-regulated magnetohydrodynamics: Transport equations for electron thermal conduction in the high-beta intracluster medium of galaxy clusters,’ *The Astrophysical Journal*, vol. 923, no. 2, p. 245, 2021. DOI: 10.3847/1538-4357/ac1ff1.
- [184] Y. V. Khotyaintsev, D. B. Graham, C. Norgren and A. Vaivads, ‘Collisionless magnetic reconnection and waves: Progress review,’ *Frontiers in Astronomy and Space Sciences*, vol. 6, 2019, ISSN: 2296-987X. DOI: 10.3389/fspas.2019.00070.
- [185] K. Fujimoto and R. D. Sydora, ‘Whistler waves associated with magnetic reconnection,’ *Geophysical Research Letters*, vol. 35, no. 19, 2008. DOI: 10.1029/2008GL035201.

Part II

Included papers

**Data-driven discovery of a heat flux closure for
electrostatic plasma phenomena**

E. R. Ingelsten, M. C. McGrae-Menge, E. P. Alves and I. Pusztai

Journal of Plasma Physics **91**(2) E64 (2025)

Data-driven discovery of a heat flux closure for electrostatic plasma phenomena

Emil R. Ingelsten¹ , Madox C. McGrae-Menge², E. Paulo Alves^{2,3} and Istvan Pusztai¹ 

¹Department of Physics, Chalmers University of Technology, Göteborg SE-41296, Sweden

²Department of Physics and Astronomy, University of California, Los Angeles, CA 90095, USA

³Mani L. Bhaumik Institute for Theoretical Physics, University of California at Los Angeles, Los Angeles, CA 90095, USA

Corresponding author: Emil R. Ingelsten, emilraa@chalmers.se

(Received 27 November 2024; revision received 4 March 2025; accepted 4 March 2025)

Progress in understanding multi-scale collisionless plasma phenomena requires employing tools which balance computational efficiency and physics fidelity. Collisionless fluid models are able to resolve spatio-temporal scales that are unfeasible with fully kinetic models. However, constructing such models requires truncating the infinite hierarchy of moment equations and supplying an appropriate closure to approximate the unresolved physics. Data-driven methods have recently begun to see increased application to this end, enabling a systematic approach to constructing closures. Here, we use sparse regression to search for heat flux closures for one-dimensional electrostatic plasma phenomena. We examine OSIRIS particle-in-cell simulation data of Landau-damped Langmuir waves and two-stream instabilities. Sparse regression consistently identifies six terms as physically relevant, together regularly accounting for more than 95 % of the variation in the heat flux. We further quantify the relative importance of these terms under various circumstances and examine their dependence on parameters such as thermal speed and growth/damping rate. The results are discussed in the context of previously known collisionless closures and linear collisionless theory.

Keywords: plasma simulation, plasma dynamics

1. Introduction

Global modelling of multi-scale collisionless plasma phenomena is a long-standing computational challenge. Even with state-of-the-art computing resources, it is often computationally infeasible to resolve the smallest scale, kinetic effects within large-scale, global domains relevant for fusion and astrophysical systems. Thus, progress in understanding these systems is bound to happen through a combination of *ab initio* kinetic modelling and reduced models. In the latter category, collisionless

fluid models have demonstrated their utility in global modelling (Dong *et al.* 2019; TenBarge *et al.* 2019; St-Onge *et al.* 2020; Ng *et al.* 2020) and as part of hybrid (kinetic-fluid) schemes (Shi *et al.* 2021; Arzamasskiy *et al.* 2023; Achikanath Chirakkara *et al.* 2023). For such reduced models to be able to capture the essence of the kinetic physics at play, a systematic approach to constructing accurate fluid closures is called for. A fluid closure relates higher-order fluid quantities (e.g. heat flux) – for which the exact evolution equation is not retained – to lower-order fluid quantities (e.g. density, flow velocity) and fields, allowing the truncation of the otherwise infinite hierarchy of fluid equations. Due to the lack of a universal closure for collisionless dynamics, each closure must be tailored to the phenomena of interest.

In collisional systems, the distribution functions of the particles remain close to local thermodynamic equilibrium, allowing for rigorous construction of closed fluid models (Braginskii 1958; Chapman & Cowling 1991). There is, however, no generally applicable closures for collisionless systems that are characterised by significant departures from Maxwellianity and non-local kinetic phenomena, such as wave-particle interactions. Nevertheless, there are theoretical approaches which distil various aspects of the relevant physics into the form of the closure equations. The closure by Hammett & Perkins (1990) is constructed to capture Landau damping, the Chew-Goldberger-Low (CGL) closure (Chew, Goldberger & Low 1956) involves the (generally anisotropic) pressure in such a way as to conserve adiabatic invariants in collisionless magnetised plasmas, and the closure by Levermore (1996) is based on the principle of maximum entropy, to name a few. However, all of these approaches, while theoretically motivated, have limited scope – such as requiring linearity or exact adiabaticity. These assumptions break down for many problems of interest, for instance, in the presence of turbulence or magnetic reconnection. Furthermore, some variants are numerically difficult to work with due to spatial non-locality. In addition, there are *ad hoc* closures, such as the relaxation closure that drives the pressure tensor towards an isotropic pressure (Wang *et al.* 2015). These have had varying success in reproducing kinetic simulation results, and may contain free parameters that cannot be determined theoretically, and the choice of which can drastically alter dynamics.

An alternative systematic line of action to obtain closures for collisionless plasma systems is to look towards data-driven methods, where the closures are constructed to conform with kinetic simulation data. Neural network-based machine learning is an effective tool to this end (Wang *et al.* 2020; Maulik *et al.* 2020; Qin *et al.* 2023), though it lacks interpretability, which makes it difficult to gain intuition and generalisable understanding from. However, symbolic regression (Makke & Chawla 2024) and sparse regression (SR) (Brunton, Proctor & Kutz 2016; Rudy *et al.* 2017; Schaeffer 2017) methods can be used to infer interpretable and generalisable equations describing a dynamical system. The models thus obtained are parsimonious, lying at the Pareto-front trading between predictive power and model complexity. SR has been used to discover the governing equations of dynamical systems in a broad range of fields previously, but it has only recently been introduced in plasma physics (Dam *et al.* 2017; Kaptanoglu *et al.* 2021b, 2023a; Alves & Fiuza 2022) and there is only a very limited number of attempts to use it for closure discovery. Donaghy & Germaschewski (2023) employ SR to recover collisionless fluid equations and discover a heat flux closure in the strongly nonlinear state of a two-dimensional Harris-sheet reconnection scenario. They do not attempt to interpret the found closure and avoid the linear regime, which is difficult due to noise

at small amplitudes. Combining sparse regression and deep learning neural networks, Cheng *et al.* (2023) recover fluid equations and the local approximation (Sharma *et al.* 2006; Ng *et al.* 2020) of the Hammett–Perkins closure in a one-dimensional linearly Landau-damped Langmuir standing wave set-up.

Here, we employ the SINDy (Sparse Identification of Nonlinear Dynamics) algorithm (Brunton *et al.* 2016; Rudy *et al.* 2017) for sparse regression to obtain a heat flux closure – an expression for the heat flux in terms of lower-order moments – in one-dimensional (1-D) electrostatic plasma scenarios. Specifically, we examine Landau-damped Langmuir waves, and to gain further insights into the identified closure terms and illustrate their more general nature, we also study set-ups exhibiting two-stream instability (Stix 1992) (and the following nonlinear dynamics), ubiquitous in space and astrophysical systems (Khotyaintsev *et al.* 2019). Using particle-in-cell simulation data produced with the OSIRIS code (Fonseca *et al.* 2002, 2008), we search for optimally accurate expressions for the heat flux at each given model complexity (i.e. number of terms in the closure expression). Covering both linear and nonlinear stages, we follow the time evolution of the closure terms across the development of an electrostatic two-stream scenario, from growth through saturation via the formation and merging of phase-space holes. We also elaborate on the parametric dependences of the terms found and quantify their relative importance. The expressions are interpreted in the context of the local Hammett–Perkins closure. Analytically obtained constraints between the various closure terms are also provided to support the regression results and to assist their interpretation.

The rest of the article is organised as follows. In § 2, we describe the sparse regression method employed and the simulation set-up for the systems we study, exhibiting Landau damping and growth. In § 3, we outline and analyse the heat flux model terms identified by SR. More specifically, we describe the results of applying SR to simulations of Landau-damped Langmuir waves and two-stream instabilities in §§ 3.1 and 3.2, respectively. We then examine the relative importance of the various terms found in § 3.3, relate the six terms found most consistently by SR to linear collisionless theory in § 3.4 and finally, in § 3.5, discuss a fundamentally nonlinear seventh term which is also identified as relevant by SR in many cases. We conclude by summarising our results and giving an outlook on future work in § 4. Furthermore, we include Appendix A, containing a derivation of 1-D electrostatic linear collisionless theory from the Vlasov–Maxwell system, as well as the constraints this imposes on the heat flux model found by SR. Finally, in Appendix B, we give a demonstration of how SR works by going through how one can recover the 1-D momentum equation from simulation data.

2. Methods

Our aim is to find approximate, spatio-temporally local analytical expressions for the heat flux in terms of lower-order fluid quantities, such that these expressions capture most of the variation in the heat flux observed in kinetic simulation data. When discovering heat flux closures for a given physical system, we start by performing a kinetic simulation of the system in question using the particle-in-cell (PIC) code OSIRIS. During the simulation, we export diagnostics for all fluid quantities present in the three lowest-order collisionless fluid equations (A.2), namely the number density $n_\sigma = \int d^3\mathbf{v} f_\sigma(\mathbf{v})$, flow velocity $V_\sigma = n_\sigma^{-1} \int d^3\mathbf{v} \mathbf{v} f_\sigma$, mass-normalised pressure tensor $\mathbf{p}_\sigma = \int d^3\mathbf{v} (\mathbf{v} - V_\sigma)^{(2)} f_\sigma$ and mass-normalised heat flux tensor $\mathbf{q}_\sigma = \int d^3\mathbf{v} (\mathbf{v} - V_\sigma)^{(3)} f_\sigma$ for each species σ , as well as electromagnetic field data (electric

field \mathbf{E} and magnetic field \mathbf{B}) at regular time intervals. Here, $f_\sigma(\mathbf{v})$ denotes the distribution function for species σ , and we use notation where $[\mathbf{ab}]_{ij} = a_i b_j$ and $\mathbf{a}^{(2)}$ is shorthand for \mathbf{aa} . For accurate regression results, it is important that the version of OSIRIS used here corrects for the otherwise occurring half-time step shifts between position and momentum data, characteristic of PIC codes using a leap-frog scheme (Boris & Shanny 1972; Hockney & Eastwood 2021). We also post-process our data to correct for staggering of the fields through linear interpolation (see Appendix B for a demonstration of the importance of correcting for such misalignments).

2.1. Sparse regression

In general, the aim of a sparse regression (SR) algorithm is to find an approximate relationship between some target quantity y and a set of M possibly relevant quantities $\{\theta_j\}_{j=1}^M$, while keeping model complexity low. In the version of SINDy we use, which is one of the modified versions of the PDE-FIND algorithm described by Alves & Fiuza (2022), the aim is specifically to approximate y as a linear combination of the θ_j quantities. To accomplish this, we randomly select N small space–time volumes from the simulation domain and integrate both y and all θ_j quantities over these small volumes to reduce noise.¹ We then collect the volume-integrated y and θ_j quantities from all the sampled points in the domain into a vector \mathbf{y} and a matrix Θ defined so that

$$[\mathbf{y}]_i = y \Big|_{\text{pt } i} \quad \text{and} \quad [\Theta]_{ij} = \theta_j \Big|_{\text{pt } i}. \quad (2.1)$$

With these definitions, the task of approximating y becomes a question of finding the coefficient vector ξ that optimally solves the equation

$$\mathbf{y} = \Theta \xi. \quad (2.2)$$

For us, ‘optimally’ means achieving a low mean squared error with as few non-zero terms as possible, maximising not just accuracy but also model simplicity and generalisability. Thus, our cost function looks like

$$C(\xi) = \|\mathbf{y} - \Theta \xi\|^2 + \lambda \|\xi\|_0, \quad (2.3)$$

where $\|\xi\|_0$ denotes the 0-norm of ξ , i.e. the number of non-zero coefficients. The λ hyperparameter is effectively gradually increased from 0, leading to increasingly harsher penalisation of models with many non-zero terms. The end result of this procedure is a sequence of models which are optimally accurate at each given model complexity, sweeping along the Pareto front. To curb overfitting and more easily discern which terms are spurious, we perform 10-fold cross-validation – terms in Θ which are found consistently are more likely to be physical. The efficacy of our

¹This approach, where one reduces the effect of particle noise by integrating over the data with some kernel, is in general known as the weak formulation of SR (Schaeffer & McCalla 2017). Integrating over small space–time volumes ‘without a kernel’, as we do, is a special case of this approach, effectively corresponding to using a space–time box function as a kernel. It should be noted that using smooth test functions, as is done in SPIDER (Gurevich *et al.* 2024) and WSINDy (Messinger & Bortz 2021), provides better accuracy when higher-order derivatives are important.

SR approach is demonstrated in Appendix B through a recovery of the electron momentum equation.

In our case, we seek a heat flux closure for modelling electrostatic plasma phenomena, meaning our y quantities are the elements of the heat flux tensor \mathbf{q}_σ . As for the set of possibly relevant quantities θ_j , in principle, one would want to include all possible expressions involving n_σ , V_σ , \mathbf{p}_σ , \mathbf{E} and \mathbf{B} . In practice, however, this is infeasible, since the space of possible expressions is infinite. As it turns out, even restricting to e.g. arbitrary products of the form

$$\theta_j = n_\sigma^\alpha \prod_k V_{\sigma k}^{\beta_k} E_k^{\gamma_k} B_k^{\delta_k} \prod_l p_{\sigma kl}^{\varepsilon_{kl}}, \tag{2.4}$$

where the exponents α , β_k , γ_k , δ_k and ε_{kl} are non-negative integers summing to \leq some integer s , results in enormous term libraries Θ even when s is relatively small due to the combinatorics involved. Since having a very large term library not only increases computational cost, but also often leads to issues with convergence, choosing a term library with as few superfluous terms as possible is desired.

For the one-dimensional electrostatic plasma problems we consider, where only electron physics is relevant over the time scales of interest, we can first restrict ourselves to considering only n_e , V_{e1} , E_1 , p_{e11} and set $y = q_{e111}$. Since only electrons are relevant, and all vectors and tensors in 1-D have just a single degree of freedom, we can also suppress species and coordinate indices from now on. For convenience, we also normalise to the electron mass m_e , the elementary charge e , the speed of light c and the plasma frequency $\omega_{pe} = \sqrt{\bar{n}_e e^2 / (\varepsilon_0 m_e)}$ at the unperturbed electron density \bar{n}_e . This also normalises distances to the electron inertial length $\delta_e = c / \omega_{pe}$.

To further narrow the range of possible candidate terms, we start with only those terms which are dimensionally consistent with our y variable q , i.e. terms of the form

$$\theta_j = n v_{th}^\alpha V^{3-\alpha} \tag{2.5}$$

for some integer $\alpha \leq 3$, where $v_{th} = \sqrt{T} = \sqrt{p/\bar{n}}$, defining $T = p/\bar{n}$ to be (the 11-component of) the mass-normalised temperature tensor. Inspired by the local approximation (Ng *et al.* 2020) of the Hammett–Perkins closure, which involves a temperature gradient, we extend this initial set of candidate terms to also allow similar ones with first-order spatial derivatives, e.g. $n v_{th} \partial_x(v_{th}) \partial_x(V)$. It should be noted, however, that the presence of the spatial derivative in these additional terms means that the coefficient corresponding to each such term will be dimensional. For instance, the example term mentioned above with two spatial derivatives necessitates a coefficient with a dimensionality of length squared. This in turn suggests a scaling $\sim L^2$ for the coefficient in question, where L is the characteristic length scale for the variation in the quantities involved.

We emphasise that restricting our term library in this way specifically is an arbitrary choice, made to limit the term library size so as to make SR convergence more likely. We start by considering dimensionally consistent terms mainly because models constructed from such terms contain only unitless coefficients, which facilitates generalisability. The restriction to integer α is made for convenience. Our exclusion of terms with higher-order derivatives is chiefly motivated by the fact that their inclusion would lead to difficulties with SR convergence due to the vastly increased term library size. Furthermore, closures constructed from such terms are more difficult

to work with computationally, since even first-order derivatives in the expression for \mathbf{q} yield second-order derivatives in $\nabla \cdot \mathbf{q}$ and thus in the fluid equation system one needs to solve. Importantly, the function space we have restricted ourselves to seems sufficient to model q accurately, as we shall see.

2.2. Simulation set-up

In all of our simulations, we kinetically model an electron–proton plasma in one spatial and three² velocity dimensions in the centre-of-mass (CoM) frame, with physical mass ratio, a spatial resolution of $\Delta x = 10^{-3} \delta_e$ and periodic boundary conditions. Since our simulations are all performed in the initial CoM frame and ion flow velocities remain negligibly small, every instance of an electron flow velocity V below can be thought of as $V - v_{\text{CoM}}$, with a spatial average value of 0. Here, v_{CoM} is the velocity of the centre of mass. This quantity is invariant under Galilean transformations, just like n and v_{th} , meaning that all terms in our term library are frame-independent with this interpretation of V . This is very much desirable since q , the quantity we are seeking to model, is a Galilean-invariant quantity.³ For numerical stability, we consistently use a simulation-internal time step slightly smaller than the spatial resolution: $9.5 \times 10^{-4} \omega_{\text{pe}}^{-1}$. To limit the amount of data output as diagnostics, we save the state of the simulation only once every 100 time steps, thus our regression analysis uses an effective temporal resolution of $\Delta t = 0.095 \omega_{\text{pe}}^{-1}$.

2.2.1. Landau-damped Langmuir waves

When studying Landau-damped Langmuir waves, we initialise the plasma as a Maxwell distribution with various non-relativistic thermal speeds $v_{\text{th}} \sim 0.01c$ using a domain size of $0.256 \delta_e$ with 10^5 (10^4) electrons (ions) per cell. To excite Langmuir waves, we then apply and smoothly turn off an external sinusoidal \mathbf{E} -field perturbation propagating in the $+x$ or $-x$ direction with wavenumber $|k| = 4\pi / (0.256 \delta_e)$ and a frequency ω_r matching that of the analytic Langmuir mode. More specifically, this is done by using a single-cycle sine squared envelope, reaching maximum amplitude at $\omega_{\text{pe}} t = 3$ and being fully turned off at $\omega_{\text{pe}} t = 6$. The values of v_{th} considered, along with corresponding $|k| \lambda_{D,e}$ values (where $\lambda_{D,e}$ is the electron Debye length), as well as frequencies and growth rates of the resulting Langmuir waves, are shown in [table 1](#).

After the external forcing is removed, the system is left to evolve self-consistently, with the resulting Langmuir waves decaying due to Landau damping – initially exponentially, with only linear processes involved, as can be seen in [figure 1\(a\)](#). The PDE-FIND algorithm is then applied to find a closure for q during the timeframe of length Δt_L where decay is judged to be exponential (e.g. $6.0 < \omega_{\text{pe}} t < 21.0$ for initial $v_{\text{th}} \leq 0.01c$ – see also [figure 1a](#), where this time range is highlighted in red). In total, $\sim 6\%$ of this space–time range is randomly sampled per cross-validation fold, of which there are 10. The values of Δt_L for all values of v_{th} considered are listed in [table 1](#), together with the estimated bounce times t_b for trapped electrons. Note that for the four lower thermal speeds considered, decay is exponential for roughly

²As noted above, however, we only expect the components along the single modelled spatial dimension to be of importance, meaning we are for most intents and purposes treating our simulation as 1DIV.

³An alternative approach to ensuring Galilean or Lorentz invariance (or some other symmetry of the system) is the augmentation of simulation data through application of transformations of the corresponding type before performing SR (McGrae-Menge *et al.* 2023).

Initial $v_{th}/(0.01c)$:	0.8	0.9	1	1.1	1.25	1.5
Resulting $ k \lambda_{D,e}$:	0.393	0.442	0.491	0.540	0.614	0.736
ω_r/ω_{pe} ($\pm 0.06^*$):	1.24	1.31	1.36	1.42	1.50	1.67
Resulting $\omega_r/(kv_{th})$:	3.16	2.97	2.77	2.63	2.44	2.27
γ/ω_{pe} :	-0.0692	-0.109	-0.150	-0.195	-0.291	-0.429
$\omega_{pe}\Delta t_L$:	15.0	15.0	15.0	14.5	8.3	6.4
$\omega_{pe}t_b$:	7.30	7.66	8.10	8.61	9.56	11.3

TABLE 1. Values of v_{th} used when studying Landau-damped Langmuir waves, the frequency ω_r and the (negative) growth rate γ of the excited Langmuir wave and the duration Δt_L of the period of exponential decay, over which SR is applied, as well as the estimated bounce time t_b for trapped electrons in each case. We also list the values of $|k|\lambda_{D,e}$ and $\omega_r/(kv_{th})$ resulting from the other parameters. The frequencies ω_r are calculated via Jacobsen interpolation (Jacobsen & Kootsookos 2007) of the peaks in the discrete Fourier transform (DFT) spectrum for the \mathbf{E} -field, with uncertainty (*) corresponding to half the DFT bin size. The growth rate γ is calculated via linear regression on logarithmised data of the average \mathbf{E} -field energy density over the period of exponential decay. The estimated bounce time is calculated as $t_b = \sqrt{m_e/(e|k|E_{rms})}$, where $E_{rms} = \sqrt{\langle E^2 \rangle}$ is the spatial root-mean-square average of the \mathbf{E} -field magnitude at the point in time when the external drive is switched off.




FIGURE 1. Evolution of the spatially averaged \mathbf{E} -field energy density $\langle \frac{1}{2}\epsilon_0 E^2 \rangle$ over time (a) in the Landau damping case where $v_{th} = 0.01c$ and (b) in the two-stream case where $n_b/n_e = 0.1$, normalised to the rest energy of an electron and the unperturbed electron number density \bar{n} . The linear decay/growth phase is highlighted in red.

twice the bounce time, meaning that we may expect slight nonlinear trapping effects towards the end of the sampled time window. At higher values of v_{th} , however, these effects are drowned out by numerical noise present at the low perturbation amplitudes reached towards the end of exponential decay. Indeed, this is the reason for Δt_L being lower than t_b for the two highest v_{th} values considered – in these cases, the decay is so rapid that the perturbation disappears in the numerical noise before trapping effects become visible.

n_b/n_e :	0.01	0.02	0.05	0.1	0.2	0.4	0.5
$v_{\text{th}}/(0.01c)$:	0.373	0.422	0.503	0.678	0.860	1.03	1.05
$\delta_e k_{\text{char}}$:	-60.8	-59.9	-60.2	-62.3	-64.3	-67.6	-67.4
Resulting $ k_{\text{char}} \lambda_{D,e}$:	0.227	0.253	0.303	0.422	0.553	0.696	0.708
$\omega_r/\omega_{\text{pe}} (\pm 0.03^*)$:	1.00	0.96	0.88	0.80	0.66	0.19	0
$\gamma/\omega_{\text{pe}}$:	0.0515	0.0837	0.137	0.184	0.234	0.273	0.276
$\omega_{\text{pe}}\Delta t_L$:	31.3	23.7	20.9	19.0	13.3	12.3	11.9
$\omega_{\text{pe}}t_b$:	5.95	4.95	3.54	2.94	2.54	2.28	2.29

TABLE 2. Values of n_b/n_e examined when studying two-stream instabilities and the corresponding values of v_{th} , as well as the frequency ω_r , growth rate γ and characteristic wavenumber k_{char} for the excited perturbation (also listing $|k_{\text{char}}|\lambda_{D,e}$), together with the duration Δt_L of exponential growth and the bounce time t_b . The same methods are used to calculate ω_r and γ as in the Landau damping cases, described in table 1. The values of k_{char} are calculated by using a weighted average over the DFT spectrum at the end of linear growth, using the absolute value of the Fourier amplitude squared as the weighting. The minus signs signify propagation towards $-x$. Similarly to what was done in the Landau damping cases, we estimate $t_b = \sqrt{m_e/(e|k_{\text{char}}|E_{\text{rms}})}$, with both k_{char} and E_{rms} in this case being evaluated at the time of maximum average \mathbf{E} -field energy density. Note that while k_{char} at this time is slightly different than the values listed in this table, the difference is marginal ($\sim 1\%$).

2.2.2. Two-stream instabilities

Apart from studying Landau damping, we also examine a set-up exhibiting Landau growth, namely a two-stream unstable plasma. In this case, we choose a larger domain size of $2.048\delta_e$ to limit the effects of the periodic boundary conditions employed. We initialise the ions in equilibrium, with the electrons split into counterflowing equal temperature Maxwellian populations – a core population with density n_c and flow velocity V_c , and a beam population with density n_b and flow velocity V_b . The thermal speed for each population individually is $u_{\text{th}} = 3.16 \times 10^{-3}c$. We vary n_b as a fraction of the total electron density n_e , keeping the relative velocity $V_{\text{rel}} = V_c - V_b$ constant at $0.02c$ and staying in the zero-current frame by enforcing $n_c V_c + n_b V_b = 0$. Specifically, we consider the range of values for n_b/n_e listed in table 2. In these simulations, we use 10^4 (200) electrons (ions) per cell.

Note that the non-zero relative velocity V_{rel} between the populations means that the combined electron population had a thermal speed $v_{\text{th}} = \sqrt{p/n} > u_{\text{th}}$. More specifically, this combined thermal speed is $v_{\text{th}} = [u_{\text{th}}^2 + (n_b/n_e)(1 - n_b/n_e)V_{\text{rel}}^2]^{1/2}$, with a maximum of $v_{\text{th}} \approx 1.05 \times 10^{-2}c$ for $n_b/n_e = 0.5$.

The counterflowing electron populations drive wave growth via inverse Landau damping. Similarly to the decay of the Langmuir waves above, this growth is initially exponential. It eventually saturates, however (as can be seen in figure 1b), leading to the formation of phase-space electron holes. With the data from these simulations, SR is performed (a) over the linear part of the growth phase, and (b) over small time slices of length $1.9\omega_{\text{pe}}^{-1}$ covering both the growth phase and the saturated phase to study how the closure coefficients evolve over time, as illustrated in figure 3 – in both cases for all values of n_b/n_e listed in table 2. Part (a) here is very much analogous to what was done for the Landau damping case. For example, in the case where $n_b/n_e = 0.1$, the time range sampled is $19.0 < \omega_{\text{pe}}t < 38.0$, highlighted in red in figure 1(b), meaning $\omega_{\text{pe}}\Delta t_L = 19.0$. In part (b), the time slices are centred on

time steps $t = 1.9 \times \{1, 2, 3, \dots\} \omega_{pe}^{-1}$, covering the entire simulation domain up to the last such time step which is $> 0.95 \omega_{pe}^{-1}$ from the end of the simulation, so that every time slice falls entirely within the domain of the simulation. In both of these cases, each of the 10 cross-validation folds randomly samples $\sim 2.5\%$ of the data in the space–time ranges of interest.

Since we are now dealing with exponential growth rather than exponential decay, the influence of noise towards the end of the linear part of the process is significantly decreased compared with the situation in the Landau damping simulations. This is also clearly visible in the stronger relationship between Δt_L and t_b in these simulations – we consistently have $\Delta t_L \sim 5.5 t_b$, meaning growth is exponential for a little over five times the bounce time. This suggests that there is a high probability of nonlinear trapping-related effects being present to some extent in the data towards the latter half of the sampled time range.

The excited perturbations in this case are more broad-spectrum than the Langmuir waves examined above, necessitating the introduction of characteristic wavenumbers k_{char} , calculated as outlined in the caption of [table 2](#). The temporal spectrum is dominated by a single frequency peak however, at the value of ω_r listed in [table 2](#). Like in [table 1](#), we also, for convenience, show $|k_{char}| \lambda_{D,e}$.

For simplicity and to more easily compare our results with those from the Landau-damping case, we only consider the combined electron species rather than treating the counter-streaming populations separately when performing SR. That is, all of our closure models are models of the total electron heat flux, and our term library is constructed from fluid quantities relating to the entire electron population. We note, however, that for the purposes of modelling two-stream unstable systems in fluid codes, it is likely more practical to treat the two electron populations as separate species, with their own respective closures. Preliminary results suggest that applying SR when using such an approach also yields broadly similar closure terms as those identified here, in appropriately chosen reference frames. To avoid diverting our focus, we leave a more detailed analysis of two-stream instability along these lines, with separate closures for the beam and core populations, outside the scope of this article.

3. Results

For both of the set-ups we considered, SR yields very similar results. In both cases, a six-term model $q = q_{even} + q_{odd}$ was found, where

$$\begin{cases} q_{even} = A_1 n v_{th}^2 V + A_2 v_{th}^3 \partial_x n + A_3 n v_{th}^2 \partial_x v_{th}, \\ q_{odd} = A_4 + A_5 n v_{th}^3 + A_6 n v_{th}^2 \partial_x V. \end{cases} \quad (3.1)$$

The split into q_{even} and q_{odd} is based on the dependence on the propagation of the perturbations involved. While the coefficients in front of the q_{even} terms are independent of propagation direction (and thus ‘even in k ’), the q_{odd} coefficients switch sign if the propagation direction is reversed (and are thus ‘odd in k ’). This also means that if there is no wave propagation, or when oppositely propagating waves are of similar amplitudes, such as in a standing-wave scenario, all q_{odd} coefficients go to zero. In several cases, an additional term $\propto n v_{th} V^2$ is found. As this term mostly appears for low $|\gamma|$ – specifically towards the end of linear growth or decay processes when $|\gamma|$ is decreasing – it appears to help capture weak nonlinear trapping effects.

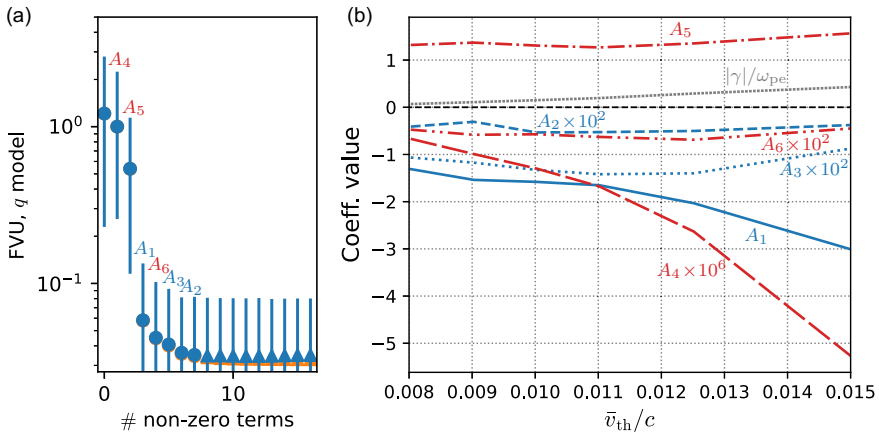


FIGURE 2. Results from the Landau damping simulations. (a) FVU of the successive closures found for $v_{th}/c = 0.01$ (where $\gamma = -0.150$), with each new term indicated by its coefficient in (3.1). Circles, unlike triangles, denote consistently found terms and blue (orange) marker colour corresponds to performance on testing (training) data. Note that while FVU is typically smaller than unity, it is higher than unity in certain situations. For example, this is the case for the 0-term model, which is identically zero, meaning $\hat{y}_i = 0$ for all i , while the mean of the y data is $\bar{y} \neq 0$. (b) Dependence of the closure coefficients on v_{th} , for a wave propagating in the $+x$ direction. Blue (red) lines/symbols indicate q_{even} (q_{odd}) coefficients. For comparison, $|\gamma|/\omega_{pe}$ is plotted as a grey tightly dotted line.

We note that, due to the presence of spatial derivatives in terms 2, 3 and 6, the corresponding coefficients are dimensional, having units of length. When plotting them, they are implicitly given in units of δ_e , though often re-scaled by a factor of 10^2 for readability. In other words, a label $A_2 \times 10^2$ should read $10^2 A_2 / \delta_e$. Similarly, A_4 has the units of mass-normalised heat flux, or number density times velocity cubed, and is implicitly given in units of $\bar{n}c^3$. Thus, a curve labelled $A_4 \times 10^6$ shows $10^6 A_4 / (\bar{n}c^3)$.

We consistently quantify the error of the various models found by SR using the fraction of variance unexplained (FVU), defined as $\text{FVU} = \sum_i (y_i - \hat{y}_i)^2 / \sum_i (y_i - \bar{y})^2$, where y_i is the value of the y quantity at the i th sampling point, \hat{y}_i is the y -value predicted by the model at that point, $\bar{y} = \frac{1}{N} \sum_i y_i$ is the mean y -value and the sums run over the N samples. As stated in § 2, the y -quantity of interest to us is the total electron heat flux $q = q_{e111}$.

3.1. Landau-damped Langmuir waves

In the simulations of Landau-damped Langmuir waves, SR found the six-term closure in (3.1) consistently, with an FVU of 2%–7%. The FVU increases with higher initial v_{th} and corresponding stronger damping. In simulations of standing waves (i.e. the sum of oppositely propagating waves), only the q_{even} terms are found, consistent with the lack of a preferred direction.

We note that in some cases, additional terms are sometimes found consistently. For example, the unannotated seventh circle marker in figure 2(a) corresponds to a term $\propto nV^2 \partial_x V$. In particular, at low $|\gamma|$, where nonlinear trapping effects are

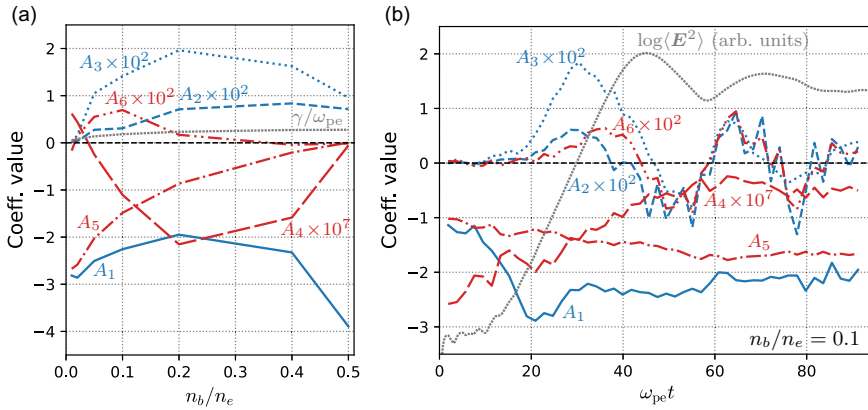


FIGURE 3. Results from the two-stream instability simulations. (a) Dependence of the growth phase closure coefficients on n_b/n_e . For comparison, γ/ω_{pe} is plotted as a grey tightly dotted line. (b) Evolution of the coefficients over time in the case where $n_b/n_e = 0.1$. Blue (red) lines indicate q_{even} (q_{odd}) coefficients. Here, the grey tightly dotted line is the logarithm of the spatially averaged E -field energy density in arbitrary units, recognisable from figure 1(b).

expected to be slightly more important, a term $\propto nv_{\text{th}}V^2$ is found, likely helping capture these weak nonlinear effects.

Since we are examining an electrostatic 1-D setting affected by Landau damping, one might expect the closure to be similar to the local approximation (Sharma *et al.* 2006; Ng *et al.* 2020) of the Hammett–Perkins closure (Hammett & Perkins 1990): $q \sim -(\chi/|k_{\text{char}}|)nv_{\text{th}}^2\partial_x v_{\text{th}}$, where k_{char} is a characteristic wavenumber. The value one should choose for the heat diffusivity χ depends on which values of $\omega/(kv_{\text{th}})$ are of interest, with the original Hammett–Perkins paper focusing on regimes relevant for the ion temperature gradient (ITG) instability. In our closure, the A_3 term can be identified as playing this role. Sparse regression finds v_{th} -dependent coefficients, see figure 2(b). For most coefficients, however, this v_{th} -dependence is significantly weaker than the v_{th} -dependence of γ over the examined range of values – the major exception being the constant term A_4 , which does not affect $\partial_x q$, the quantity we are ultimately in need of a closure for. There is one coefficient with non-negligible v_{th} -dependence which does affect $\partial_x q$, however – namely, A_1 . Interestingly, we find that the improvement in predictive power caused by the introduction of this $A_1nv_{\text{th}}^2V$ term into the model is significantly larger than the one coming from the Hammett–Perkins-like $A_3nv_{\text{th}}^2\partial_x v_{\text{th}}$ term, as seen in figure 2(a).

3.2. Two-stream instabilities

As noted above, broadly speaking, the same six-term model is found for two-stream instabilities as for Landau-damped Langmuir waves, at an even lower FVU than in the Landau damping simulations of between 0.5% and 5%. Furthermore, the term $\propto nv_{\text{th}}V^2$ is again found in cases with low $|\gamma|$ – in fact, it appears for all $n_b/n_e < 0.4$ (but is left out of figure 3 for readability). Also, some coefficients are close to zero at certain n_b/n_e values – see figure 3(a). Such near-zero coefficients are generally not found consistently. For example, when $n_b = 0.5n_e = n_c$, the lack of a preferred direction forces all q_{odd} coefficients to zero, just like for a Landau-damped standing wave. The sub-1% FVU values are achieved in the middle of the examined

beam density range, where amplitudes are relatively small but growth is still rapid enough for trapping effects to be mostly negligible throughout the linear growth phase.

Both in these set-ups and the Landau damping ones examined above, the growth (or decay) phase value of A_3 tends to be roughly twice A_2 . An exact factor of 2 would correspond to having a single term $\propto v_{\text{th}} \partial_x p = v_{\text{th}}^3 \partial_x n + 2n v_{\text{th}}^2 \partial_x v_{\text{th}}$, i.e. a pressure gradient driven contribution to the heat flux, whereas the A_2 and A_3 terms, on their own, correspond to contributions due to density and temperature gradient driven heat flux contributions, respectively. Notably, while A_2 and A_3 are negative for Landau damping they are positive in the growth phase of the instability, where inverse Landau damping occurs. Interestingly, the same holds for A_6 in most cases examined here, despite the fact that the wave propagation is towards $-x$ here, whereas it is towards $+x$ for the Langmuir waves examined in figure 2. As we will see in § 3.4, this can be explained quite well by the constraints imposed on the coefficients from linear collisionless theory. Being a k -odd term, the A_6 term represents a contribution to the heat flux coming from the pressure times the rate of change of flow velocity in the direction of local wave propagation.

So far, we have considered model coefficients obtained only for the growth phase of the instability. Now, we will consider both the growth phase and the saturated phase of the simulation, with coefficients obtained in a time-resolved manner. We find that while the same model terms are sufficient to accurately approximate the heat flux throughout the simulation,⁴ some of the coefficients vary significantly across these phases. Overall, A_3 is very well correlated with the instantaneous growth rate, as is A_2 and A_6 . This is seen in figure 3(b), but is even more apparent if one plots the three coefficients in question against the instantaneous growth rate $\gamma(t) = \frac{1}{2} \partial_t \ln \langle \mathbf{E}^2 \rangle$ itself – see figure 4. What is interesting is that while the amplitude of the oscillation in A_3 decreases along with that of the oscillation in $\gamma(t)$, the other two growth-related coefficients exhibit a far smaller change in oscillation amplitude from the growth phase to the saturated phase. It is also notable that while A_2 and A_6 are practically equal during the saturated phase, they are desynchronised during the growth phase, with A_2 reaching its growth-phase maximum earlier than A_6 . Note, however, that while A_2 and A_6 are out of phase in this way during the growth phase for all examined values of n_b/n_e (disregarding the symmetric set-up with $n_b/n_e = 0.5$, where $A_6 = 0$), the exact nature of their relationship varies depending on n_b/n_e , as seen comparing the two cases shown in figure 4. In general, they synchronise earlier in simulations with higher beam density. Furthermore, the fact that the oscillation amplitudes of A_2 and A_6 are approximately equal only holds when $n_b/n_e \sim 0.1$, as one might suspect from the fact that $A_6 \rightarrow 0$ as $n_b/n_e \rightarrow 0.5$ by virtue of being k -odd.

Compared with the growth-related terms, A_1 and A_5 (as well as A_4) vary relatively slowly over time – especially A_5 . Overall, the value of A_1 fits quite well with the heuristic $A_1 \sim -3 + \frac{1}{2} \text{sgn}(k) A_5$ obtained from linear theory in the limit where $|\gamma| \ll \omega_r \sim \omega_{\text{pe}} \sim |k| \bar{v}_{\text{th}}$ (see § 3.4), the actual values in this case being $|\gamma| < 0.2 \omega_{\text{pe}}$, $\omega_r = 0.80 \omega_{\text{pe}}$ and $|k| \bar{v}_{\text{th}} \sim 0.4 \omega_{\text{pe}}$. The one notable exception to this is the very early

⁴This is not necessarily what one would expect. In general, the model terms – or even the modelling approach – might need to be adapted to the various phases of the system's evolution. To inform one's choice of model, and to quantify the complexity of the training data, which affects the difficulty of recovering/discovering terms, the Shannon information entropy metric (Kaptanoglu *et al.* 2023; Vasey *et al.* 2025b) can be used.

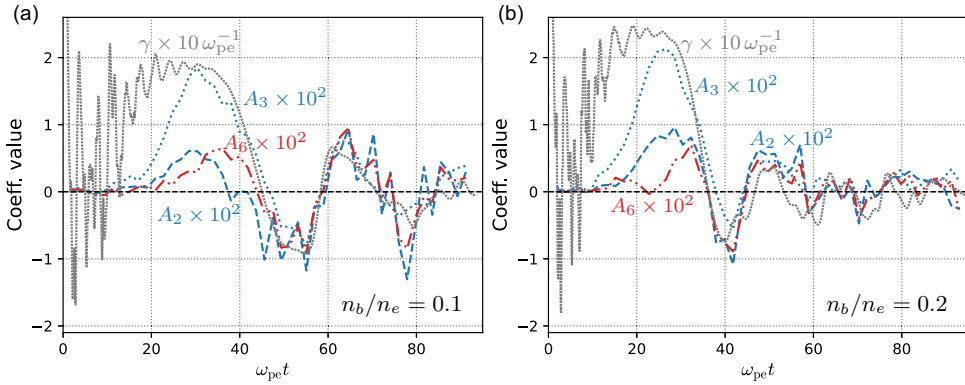


FIGURE 4. Three growth-/damping-related coefficients A_2 , A_3 and A_6 , compared with the instantaneous growth rate (a) in the case where $n_b/n_e = 0.1$, also shown in figure 3(b), and in the case where $n_b/n_e = 0.2$. As we can see, A_3 is very nearly precisely proportional to γ , while the other two coefficients are also strongly correlated with it.

growth phase, where the prevalence of high- k noise means that the physics involved is very far from this limit. In fact, in the later parts of the saturated phase, the heuristic performs better than expected, given the fact that the merging of electron holes should decrease the characteristic wavenumber $|k|$ as time goes on. There also appears to be a contribution from the growth rate γ on top of this, leading to slight oscillations in A_1 over time, consistent with the constraints imposed by linear theory as outlined in § 3.4.

When it comes to giving a physical interpretation of these non-growth-correlated terms, A_4 can be thought of as a global heat flux induced by passing waves. As seen in § 3.4, this term is beyond the purview of linear theory – in fact, the exact mechanism giving rise to this contribution is unclear. However, the A_4 term does not affect the divergence of the heat flux, which is what one is ultimately seeking to model when creating a closure for the pressure equation. The A_1 term is transparently a product of pressure and flow velocity, and it provides either the most or the second most important contribution to q , depending on n_b/n_e . Its appearance in our closure might be related to the fact that $\{V\mathbf{p}\}$ is part of the expression relating the energy flux $\mathbf{Q} = \int d^3\mathbf{v} \mathbf{v}^{(3)} f$ to the heat flux \mathbf{q} : $\mathbf{Q} = n\mathbf{V}^{(3)} + 3\{V\mathbf{p}\} + \mathbf{q}$. Thus, having a term in our q model equal to specifically $-3\{V\mathbf{p}\}$ (i.e. having $A_1 = -3$, which is quite a typical value found by SR) would signify a cancelling of this term in the energy flux. In the 1-D pressure equation, having an A_1 term as part of q similarly leads to partial cancellations. Specifically, writing the rest of q (i.e. q excluding the A_1 term) as q_r , the equation reduces to

$$\partial_t p + (1 + A_1) V \partial_x p + (3 + A_1) p \partial_x V + \partial_x q_r = 0, \tag{3.2}$$

so that the value $A_1 = -1$ would cancel the second term and $A_1 = -3$ would cancel the third term. The final of the three most important terms in our q model (see § 3.3), i.e. the A_5 term, is k -odd. Thus, it is clearly related to the wave propagation direction. Furthermore, in set-ups like ours which are in the CoM frame, V oscillates around zero while n and v_{th} have positive equilibrium values. As discussed in § 3.4, this means that all of the terms in our six-term model contribute to q at first-order in perturbation theory. While there are many terms in our term library

that do this, there are actually only two which contribute to q at zeroth-order in perturbation theory due to how our term library is constructed – precisely the A_4 and A_5 terms. This may be related to why they are identified by SR as being useful for modelling q .

The fact that some of the coefficients correlate with growth (or decay) – and as such, differ significantly between the growth and the saturated phases of the instability – means that we should not expect to find a closure with fixed coefficients which is accurate throughout both phases with respect to the contribution from these terms. To capture both phases, one would need coefficients which are informed about the phase, through e.g. a volume-averaged electric-to-thermal energy ratio. We do not aim to provide such closures here. However, as we shall see in the next section, the growth-related terms contribute relatively little to the accuracy of the model compared with the A_1 , A_4 and A_5 terms, meaning the closure we obtain in the growth phase remains quite accurate even in the saturated phase.

3.3. Quantifying the importance of terms: δ_{FVU}

To get a better sense of the circumstances under which each term in our closure is important, let us quantify their individual contributions by how much their exclusion increases the FVU of the closure, a measure we will refer to as δ_{FVU} . Doing this for the various time slices examined in [figure 3\(b\)](#), we get [figure 5](#).

The two by far most important terms are the two terms with order unity coefficients, i.e. A_5 and A_1 . The next most important term by δ_{FVU} is the constant term A_4 , although, as noted previously, this term is not very relevant to the accuracy of the closure since it has no impact on $\partial_x q$. Among the growth-related terms, A_3 is overall the most important by some margin, while A_6 and A_2 are the least important terms – A_6 , in general, being slightly more important than A_2 in this case. Interestingly, A_2 is important mostly in the first half of a linear growth or decay process, while A_6 mostly matters during the latter half. This agrees very well with what one might guess from solely looking at the sizes of the coefficients in question in [figure 3\(b\)](#). While it is not obvious that this should be the case, it is reasonable from the perspective that if the best fit for a coefficient is zero at some point in time, its importance is necessarily also zero. Since the achieved FVU is at best $\sim 5 \times 10^{-3}$, any term with $\delta_{FVU} \lesssim 10^{-4}$ can be safely assumed to be irrelevant at our level of accuracy for describing the physics during that time range.

The fact that the most important terms, A_1 and A_5 , vary quite slowly means that regardless of on which time range we perform the regression, we should expect the resulting closure to work quite well over the entire simulated time range. And indeed, this is what one finds if one plots the six-term model $q = q_{\text{even}} + q_{\text{odd}}$ with coefficient values from (e.g.) the growth phase over the entire space–time domain of the simulation and compares it to a plot of the actual q output by OSIRIS – see [figure 6](#).

This is especially promising since one of the primary use cases envisioned for these kinds of closures is sub-grid scale modelling within a larger simulation, where the instability occurs on a very short time scale compared with the overall time scales of interest. In such a situation, modelling the saturated phase ‘end state’ where $\gamma \approx 0$ is the most important. The fact that there is no growth on average means that the six-term model can be reduced to a three-term model with only A_1 , A_4 and A_5 – and, of course, providing a value for A_4 is unnecessary if one is only interested in solving the fluid equations, since A_4 does not affect $\partial_x q$.




FIGURE 5. Importance of the six terms found by SR as they vary over time in the two-stream unstable set-up with $n_b/n_e = 0.1$, as measured by δ_{FVU} .

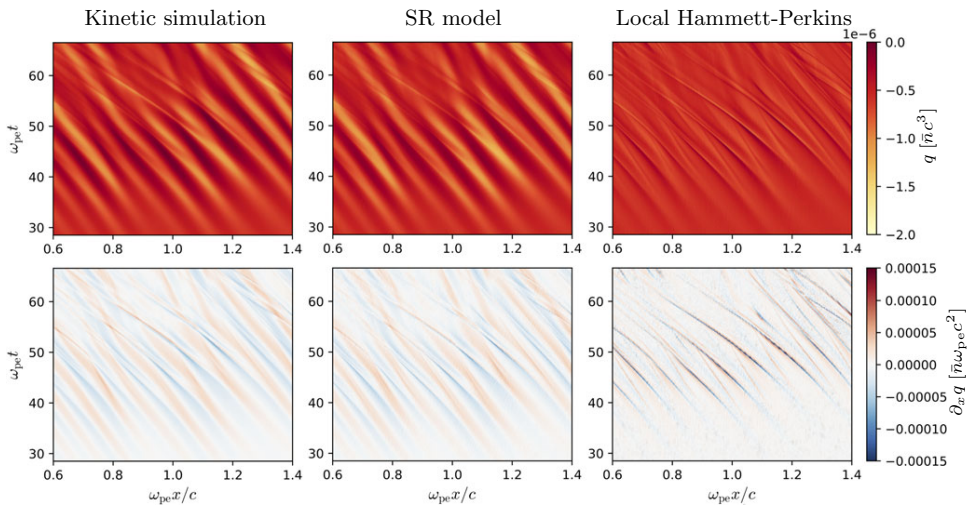


FIGURE 6. A comparison of the OSIRIS q data (top left) from the $n_b/n_e = 0.1$ simulation with our six-term SR model (top middle) and a local Hammett–Perkins model equivalent to keeping only A_3 and A_4 (top right), as well as the resulting $\partial_x q$ (bottom row). Both models are trained solely on data from the linear growth phase, corresponding to $19.0 < \omega_{pe}t < 38.0$. The six-term model performs very well even in the saturated regime, corresponding to $\omega_{pe}t \gtrsim 40$. All mass-normalised heat fluxes, like the A_4 coefficient, are given in units of $\bar{n}c^3$, and spatial derivatives of such quantities are given in units of $\bar{n}\omega_{pe}c^2$.

In general, using a six-term model trained solely on growth phase data like in figure 6 tends to yield FVU $\sim 1\%$ during the growth phase and FVU $\sim 5\%–10\%$ in the saturated phase, while using a six-term model trained on data from the saturated phase where there is no net wave growth (or, more-or-less equivalently, a three-term model with only the A_1 , A_4 and A_5 terms) yields FVU $\sim 5\%–10\%$ in

the growth phase and $FVU \sim 2\% - 5\%$ in the saturated phase. Thus, our six-term (or indeed three-term) model seems to be largely sufficient for modelling the saturated phase, despite the presence of nonlinear phenomena like particle trapping and soliton-like phase-space electron holes.

3.4. Comparison with constraints from linear theory

During linear decay or growth, the plasma should be well described by linear collisionless theory. As explained in [Appendix A](#), this gives us two predictions relating the different closure coefficients, namely

$$\begin{cases} kA_6 = \beta + \left[-kA_2\omega_r + \frac{1}{2}kA_3(1 + \Phi_-)\omega_r - \frac{1}{2}A_5(1 + 3\Phi_+)\gamma \right] \frac{k\bar{v}_{th}}{|\omega|^2}, \\ A_1 = -3 - \alpha + \left[-kA_2\gamma + \frac{1}{2}kA_3(1 + \Phi_+)\gamma + \frac{1}{2}A_5(1 + 3\Phi_-)\omega_r \right] \frac{k\bar{v}_{th}}{|\omega|^2}, \end{cases} \quad (3.3)$$

where we are using shorthand notation

$$\alpha = \frac{\omega_{pe}^2 + \gamma^2 - \omega_r^2}{k^2\bar{v}_{th}^2}, \quad \beta = \frac{2\omega_r\gamma}{k^2\bar{v}_{th}^2}, \quad \Phi_{\pm} = \frac{\omega_{pe}^2 \pm |\omega|^2}{k^2\bar{v}_{th}^2}. \quad (3.4)$$

Let us first consider some limiting cases. First, take a marginally stable perturbation with $\gamma \rightarrow 0$, where (3.3) simplifies to

$$\begin{cases} A_6 = \left[-A_2 + \frac{1}{2}A_3 \left(1 + \frac{\omega_{pe}^2 - \omega^2}{k^2\bar{v}_{th}^2} \right) \right] \frac{k\bar{v}_{th}}{\omega}, \\ A_1 = -3 - \frac{\omega_{pe}^2 - \omega^2}{k^2\bar{v}_{th}^2} + \frac{1}{2}A_5 \left(1 + 3\frac{\omega_{pe}^2 - \omega^2}{k^2\bar{v}_{th}^2} \right) \frac{k\bar{v}_{th}}{\omega}, \end{cases} \quad (3.5)$$

with $\omega = \omega_r$ real. We can immediately see that one solution of the first of these equations is $A_2 = A_3 = A_6 = 0$, in agreement with the relationship $A_{2,3,6} \sim \gamma$ found via SR. The second equation is less straightforward to interpret. If $\omega \sim \omega_{pe} \sim |k|\bar{v}_{th}$, like in most of our simulations, we expect $A_1 \sim -3 + \frac{1}{2} \text{sgn}(k)A_5$ – and this should hold as a rule of thumb even when γ is non-zero but small compared with ω_r . Qualitatively, this agrees decently with our results, even those from the growth phase where $\gamma > 0$ (but in most cases $< \omega_r$). Generally, both A_1 and A_5 are order unity, and A_1 is shifted a bit upwards from -3 in [figures 2](#) and [3](#) for all set-ups we consider except the symmetric two-stream set-up, matching the fact that $\text{sgn}(kA_5) = +1$. That this case should disagree with our rule of thumb is not surprising, since it has $\omega_r \approx 0$, while γ is finite.

The growth rate is truly negligible mainly during parts of the saturated phase of our two-stream simulations. However, the saturated phase is generally dominated by nonlinear physics, thus linear theory should not be expected to give accurate predictions. Therefore, we restrict our comparison with (3.5) to the start at the saturated phase, before the created electron holes start to merge. Specifically, let us examine the time around when the peak average \mathbf{E} -field energy is reached, marked in [figure 7\(a\)](#) for the case where $n_b/n_e = 0.1$. Performing SR over the region where the instantaneous growth rate satisfies $|\gamma(t)|/\omega_{pe} < 0.02$ near this peak for each two-stream simulation and comparing the SR value of A_1 with the value predicted by (3.5) yields [figure 7\(b\)](#). As we can see, the agreement is good for weaker beam strengths, but becomes less accurate when approaching $n_b/n_e = 0.5$, where the physics involved is more nonlinear by virtue of the larger perturbation amplitudes.

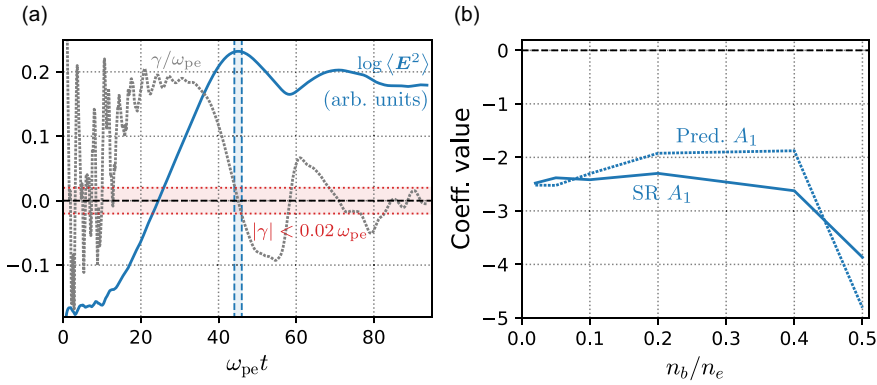


FIGURE 7. (a) Time range around peak E -field energy density where $|\gamma|/\omega_{pe} < 0.02$ for $n_b/n_e = 0.1$. (b) Values of A_1 found by SR in the two-stream simulations during the equivalent time range for all examined values of n_b/n_e , compared with the linear theory prediction at $\gamma = 0$ given by (3.5), inserting the values of A_5 found by SR.

If we instead consider the limit $\omega_r \rightarrow 0$, corresponding to a non-oscillatory – but possibly growing or decaying – perturbation, we get

$$\begin{cases} kA_6 = -\frac{1}{2}A_5 \left(1 + 3 \frac{\omega_{pe}^2 + \gamma^2}{k^2 \bar{v}_{th}^2} \right) \frac{k\bar{v}_{th}}{\gamma}, \\ A_1 = -3 - \frac{\omega_{pe}^2 + \gamma^2}{k^2 \bar{v}_{th}^2} + \left[-kA_2 + \frac{1}{2}kA_3 \left(1 + \frac{\omega_{pe}^2 + \gamma^2}{k^2 \bar{v}_{th}^2} \right) \right] \frac{k\bar{v}_{th}}{\gamma}. \end{cases} \quad (3.6)$$

Similar to the case with $\gamma \rightarrow 0$, the first equation allows for a solution where $A_6 = A_5 = 0$, consistent with the lack of wave propagation – in fact, A_4 can also be set to zero. As for the second equation, if we insert inferred parameter values from the symmetric two-stream unstable set-up,

$$\begin{cases} \gamma = 0.276 \omega_{pe}, \\ \bar{v}_{th} = 1.05 \times 10^{-2} c, \\ k \approx k_{char} = -67.4 \delta_e^{-1}, \end{cases} \quad (3.7)$$

we get

$$\frac{\omega_{pe}^2 + \gamma^2}{k^2 \bar{v}_{th}^2} \approx 2.15 \quad \text{and} \quad \frac{k\bar{v}_{th}}{\gamma} \approx -2.56, \quad (3.8)$$

giving a prediction of

$$A_1 \approx -5.15 + 2.56(kA_2 - 1.58kA_3). \quad (3.9)$$

The coefficient values found by SR in this case are

$$\begin{cases} A_1 = -3.90, \\ A_2 = 7.18 \times 10^{-3} \delta_e, \\ A_3 = 9.56 \times 10^{-3} \delta_e, \end{cases} \quad (3.10)$$

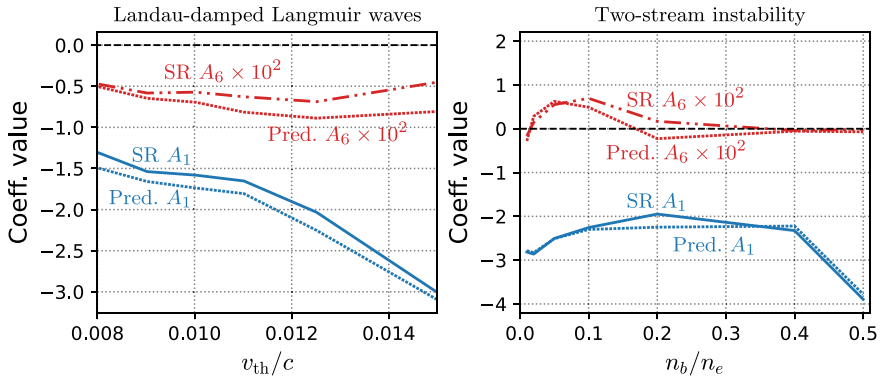


FIGURE 8. Values of A_1 and A_6 found by SR during the growth phase compared with the linear theory prediction given by (3.3), inserting the values of the other coefficients found by SR. The plot on the left contains the results from the Landau-damped Langmuir wave simulations, while the plot on the right contains the results from those with two-stream unstable set-ups.

and if we insert our values for A_2 and A_3 into the approximate expression for A_1 , we get $A_1 = -3.79$, which is reasonably accurate considering the quite broad k spectrum.

Performing a similar comparison between the values of A_1 and A_6 found by SR during linear decay/growth, and those predicted by (3.3) for all simulations, given the other coefficients as input, yields figure 8. The agreement is decent for both the Landau damping simulations and the two-stream instability ones. Interestingly, the two-stream instability simulations agree even better with linear theory, despite their more broad-spectrum nature. This is likely due to a combination of several factors. The instantaneous decay rate in the Landau damping case oscillates throughout the decay, which means that assigning the A_2 , A_3 and A_6 coefficients a single value for the entire decay phase is less accurate than doing the same for the growth phase in the two-stream simulations. In addition, the space–time domain is larger in the two-stream simulations, yielding higher-resolution DFT spectra.

3.5. A term $\propto nv_{th}V^2$

The terms making up the six-term model are in some cases not the only ones found consistently by SR. In particular, a term $\propto nv_{th}V^2$ appears consistently in regressions over the growth phase when $|\gamma|$ is small – at $\bar{v}_{th}/c < 0.01$ corresponding to $|\gamma|/\omega_{pe} \ll 0.15$ for the Landau damping simulations and at $n_b/n_e < 0.4$, i.e. $|\gamma| < 0.27\omega_{pe}$, in the two-stream simulations. Like the terms in q_{odd} , this term is k -odd, switching sign depending on the propagation direction of the waves. Notably, the cases where the term shows up are precisely those where one would expect a k -odd trapping-related term to show up, corresponding to high $\Delta t_L/t_b$ and a clear wave propagation direction (i.e. high $|\omega/k|$).

Examining how the importance of this term evolves over time, as measured by δ_{FVU} , we consistently find that it is almost as important as the A_2 and A_6 terms. At a few instances during the simulations, however, it is significantly more important than these two terms and sometimes even more important than A_3 . It is never as important as the A_1 and A_5 terms, however. Consistently, the periods where it

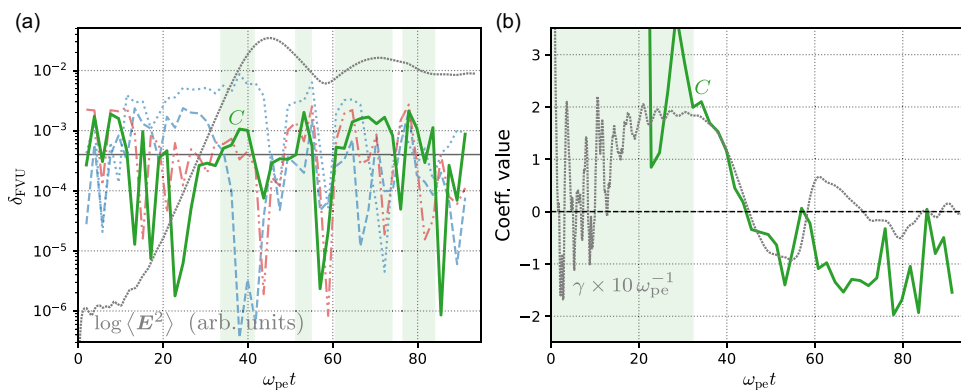


FIGURE 9. Variation over time of (a) the importance of the braking term $Cnv_{th}V^2$ as measured by δ_{FVU} and (b) the coefficient value C itself, for the two-stream unstable set-up with $n_b/n_e = 0.1$. In panel (a), we also show the δ_{FVU} values corresponding to the A_2 , A_3 and A_6 terms for reference, illustrating how the braking term is of similar importance. We highlight the time ranges where the δ_{FVU} value is above the semi-arbitrary cutoff 4×10^{-4} , largely corresponding to decreasing $|\gamma(t)|$. We also show the time evolution of the E -field energy in arbitrary units. In panel (b), we show the instantaneous growth rate $\gamma(t)$, illustrating its correlation with C . In the highlighted region, the prevalence of high-frequency noise and rapid oscillation in the growth rate causes the best-fit value of C to oscillate wildly. To improve legibility, we thus only show C at $\omega_{pe}t > 20$.

is of higher importance occur towards the end of linear processes, when $|\gamma(t)|$ is decreasing, as can be seen in figure 9(a), which shows the case $n_b/n_e = 0.1$. Because of this, it might be reasonable to refer to this term as a kind of ‘braking term’, working to damp ongoing growth or decay.

As for the value of the coefficient itself (which we can call C), it correlates well with $\gamma(t)$ starting in the latter half of the growth phase. This can be seen in figure 9(b), where the time evolution of the C coefficient in the case where $n_b/n_e = 0.1$ is shown. Unlike terms A_2 , A_3 and A_6 , however, it only changes sign once for this value of n_b/n_e , going from positive in the growth phase to negative in the saturated phase. However, its magnitude continues to oscillate along with the growth-related terms even in the saturated phase. The high-frequency noise (and oscillation in γ) in the beginning of the growth phase causes the best fit value to oscillate wildly during this time span, and because of this, we plot C only after these oscillations start to die down. At other values of n_b/n_e , the time evolution of C is similar, but the average value of the coefficient in the saturated phase, around which it oscillates, varies. More specifically, both the average value and the oscillation amplitude of C seem to grow in absolute value as n_b/n_e increases, until the term becomes unimportant at $n_b/n_e \rightarrow 0.5$ like the other k -odd terms.

This behaviour can be partly explained by the fact that $nv_{th}V^2$ is second-order in r since $V \sim rv_{ph}$, while both n and v_{th} have non-zero unperturbed values. More specifically, because of its second-order nature, we expect this term to matter only when the processes of interest deviate significantly from linearity.

4. Discussion and conclusions

Using methods from across the spectrum between physics fidelity on one hand and numerical tractability on the other is vital in exploring and understanding collisionless multi-scale plasma systems. In this context, collisionless fluid models play an important role, allowing global, long-time scale modelling of systems where this is not feasible with kinetic simulations. However, fluid models require closures to capture essential unresolved kinetic physics relevant to the plasma phenomena in question. Furthermore, theoretical closures are often derived by making idealised assumptions such as linearity or adiabatic invariance, which in many cases are broken by the dynamics of the system, motivating the use of data-driven approaches.

We have performed a theoretical and numerical study of how sparse regression can be used to discover local heat flux closures in 1-D electrostatic plasmas, examining Landau-damped Langmuir waves as well as two-stream instabilities. To ensure our inferred closures are able to capture kinetic effects, we generate our data using first-principles kinetic simulations – specifically, the OSIRIS code. The closures identified by sparse regression regularly account for more than 95% of the variation in the heat flux, while remaining limited in complexity. Thus, we demonstrate the utility of an SR-based approach to systematic closure discovery.

As noted in §2, the high accuracy of the models found by SR suggests that our term library is large enough to capture the majority of the physics at play. To test whether there is room for improvement by exploring a larger function space (while having the closure remain local), one could employ neural network models and see how they perform compared with our SR models. No such analysis is included in this paper, however, since our models already reach FVU values of a few percent.

In addition to a local approximation of the Hammett–Perkins closure, we consistently find several additional closure terms in both scenarios. Notably, the three overall most important terms – often accounting for over 90% of the variation in q – do not include the local Hammett–Perkins term. As one of these is a constant term, most of the variation in the heat flux divergence is captured by only two terms in the q model: one $\propto nv_{\text{th}}^2 V$ and one $\propto nv_{\text{th}}^3$. We further describe how the accuracy of the closure can be further improved by adding three more terms. These include a term $\propto nv_{\text{th}}^2 \partial_x v_{\text{th}}$, the further approximation of the Hammett–Perkins closure, along with terms $\propto v_{\text{th}}^3 \partial_x n$ and $\propto nv_{\text{th}}^2 \partial_x V$. These terms are closely connected to the growth or decay of waves, their best-fit coefficients being approximately proportional to the growth rate.

Among these six terms, three are independent of propagation direction and three change sign depending on propagation direction (and are thus ‘ k -even’ and ‘ k -odd’, respectively). In a three-dimensional (3-D) setting, this likely corresponds to the absence or presence of a unit vector in the wave propagation direction in the tensorial expression for the closure terms, so that terms with k -odd expressions would appear as e.g. $\propto \{\hat{k}\Omega\}$ for some two-tensor Ω . This dependence on propagation direction for the k -odd terms also means that they can only exist when wave propagation or beam asymmetry breaks isotropy.

Having reached these results, we compared the closure coefficients with predictions from linear collisionless theory, overall with quite good agreement. The two constraints imposed by linear theory give relationships between the coefficients of the various terms found by SR and the wave parameters ω_r , γ and k , as well as the plasma frequency ω_{pe} and ambient thermal speed \bar{v}_{th} of the plasma. Not entirely unexpectedly, the appearance of frequency and wavenumber in these constraints

suggests that fully capturing wave–plasma interactions requires a spatio-temporally non-local model.

Local approximations can still very much be used, however, provided one knows which parameter regimes are likely to be most relevant for the physics – cf. the local approximation of the originally non-local Hammett–Perkins closure, which needs to be supplied with a characteristic wavenumber and heat diffusivity. Already from the two constraints given by linear theory, various heuristics can be extracted depending on the parameter regime of interest. Ultimately, all parameter dependencies should be made explicit and absorbed in the closure terms, so that the coefficients are parameter-independent, giving the closure as wide a range of applicability as possible. Further elucidating the parameter dependencies of the closure coefficients is left for future work.

After examining these six terms, which are all zeroth- or first-order in the perturbation amplitude $r = \tilde{n}/\bar{n}$, we also studied one additional term with non-negligible influence on q . This term, $\propto nv_{\text{th}}V^2$, is second-order in r , and was found to mostly be important when growth or decay is slowing down. The importance of each of these seven terms was then investigated by examining their respective contributions to lowering the fraction of variance unexplained, or FVU, of the model.

Despite the high accuracy of the models described in this work, expanding the term library used by SR does merit some further investigation. Apart from more complicated expressions involving n , V and v_{th} , the \mathbf{E} and \mathbf{B} fields are also of interest – especially for higher-dimensional non-electrostatic set-ups. However, if the relevant physics is 2-D or 3-D, the number of relevant components of the vectors and tensors involved increases as well, necessitating even more careful selection of which terms to include in the SR term library. There is also reason to explore alternative algorithms for sparse regression such as SINDy-PI (Kaheman, Kutz & Brunton 2020), which generalises PDE-FIND to allow for implicit expressions for the quantity of interest y . It should further be noted that using SR in Fourier space to directly identify non-local closures like the one originally proposed by Hammett and Perkins for Landau damping warrants more investigation.

As outlined in § 1, exploring ways of systematically discovering fluid closures is chiefly motivated by the enormous computational complexity of accurately modelling multi-scale processes in plasmas. In particular, one of the main envisioned use cases for closures of the type presented in this paper is sub-grid scale modelling of rapid, small-scale processes within larger simulations. When modelling instabilities in this way, the far future limit of the saturated regime (relative to the time scales of the instability) is the most important regime to model correctly.

Finally, evaluating the performance of – and fully benefiting from – data-driven closures of this type requires a flexible implementation of closure terms in collisionless fluid solvers, such as the 10-moment solver of Gkeyll (Hakim, Loverich & Shumlak 2006; Hakim 2008), and comparing the results with equivalent kinetic simulations. Flexible closure prescription would also be a requirement for e.g. the robust – and challenging – approach demonstrated by Joglekar & Thomas (2023). In that work, the free parameters of a fluid closure are learned by neural network models, and a differentiable fluid solver is used, enabling the calculation of the loss function gradient with respect to the neural network weights. The loss function in this case quantifies the difference between the long time predictions of physics observables as calculated by a kinetic and the fluid solver. The complexity of this training process limits the number of free parameters. We envision that the SR approach we explored here can inform a good (and interpretable) starting point for closure

parametrisation. The parametric dependences of the coefficients of the model terms may then be fine-tuned with a differentiable fluid solver.

It should be noted that implementing fluid closures in existing fluid codes is far from trivial, due to the possibility of unphysical instabilities arising unless the closure is chosen with this aspect in mind; indeed, arbitrarily small errors in model coefficients may lead to numerical instability. Enforcing long-time boundedness of discovered models is possible in systems with quadratic nonlinearities (Kaptanoglu *et al.* 2021a), but it is difficult more generally. In fact, many closures (e.g. the various existing *ad hoc* relaxation closures (Wang *et al.* 2015; Ng *et al.* 2015)) are in part used because of their ability to ‘diffuse’ anisotropies, reducing the complexity of dynamics and increasing the stability of the simulation. In our context, the signs of the closure coefficients (in particular, that of the Hammett–Perkins-like term A_3) found in the Landau damping scenario are such as to increase entropy and thus provide stability. However, they correspond to an instability in the two-stream unstable scenario. It should be emphasised that having such fundamentally instability-driving terms in our closure in this case is necessary to accurately model the growth phase solely because we are modelling this physically unstable scenario using only a single electron species. Treating the counter-streaming populations as separate fluid species may be applied to resolve this shortcoming, which is the subject of ongoing investigations.

Notably, however, closures found by SR such as those described in this paper hold an advantage when it comes to ensuring stability as compared with those based on e.g. neural networks. This is due to their interpretability and low complexity, which makes it possible to study the properties of the closures analytically – something which is generally not feasible for neural networks.

Acknowledgements

The authors are grateful to K. Steinvall, D. Graham and T. Fülöp for fruitful discussions. The computations used the OSIRIS particle-in-cell simulation code, and were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement No. 2022–06725.

Editor Nuno Loureiro thanks the referees for their advice in evaluating this article.

Funding

The work was supported by the Knut and Alice Wallenberg foundation (Dnr. 2022.0087) and the Swedish Research Council (Dnr. 2021-03943). This work was also supported by the National Science Foundation Grants No. PHY-2018087 and PHY-2018089.

Declaration of interests

Competing interests: The authors declare none.

Appendix A. Constraints from linear collisionless theory

In this section, we construct constraint relations between terms appearing in a local expression for the heat flux, which then can be applied to the closure terms found using SR.

Our starting point is the Vlasov–Maxwell system, i.e.

$$\begin{cases} \partial_t f_\sigma + \mathbf{v} \cdot \nabla f_\sigma + \frac{q_\sigma}{m_\sigma} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_\sigma = 0, \\ \nabla \cdot \mathbf{E} = \frac{1}{\epsilon_0} \sum_\sigma q_\sigma n_\sigma, \\ \nabla \cdot \mathbf{B} = 0, \\ \nabla \times \mathbf{E} = -\partial_t \mathbf{B}, \\ \nabla \times \mathbf{B} = c^{-2} \partial_t \mathbf{E} + \mu_0 \sum_\sigma q_\sigma n_\sigma \mathbf{V}_\sigma, \end{cases} \quad (\text{A.1})$$

which is the most accurate self-consistent continuum description of collisionless plasmas. However, since this coupled system of partial differential equations (PDEs) is very expensive to solve over large domains while retaining high resolution, it is often necessary to simplify it. Most relevant for us is the fact that one can integrate the Vlasov equation over velocity space to instead get the fluid equations (Grad 1949; Levermore 1996). Truncating these after the pressure equation yields the system sometimes referred to as the 10-moment model (Wang *et al.* 2015),

$$\begin{cases} \partial_t n_\sigma + \nabla \cdot (n_\sigma \mathbf{V}_\sigma) = 0, \\ n_\sigma (\partial_t + \mathbf{V}_\sigma \cdot \nabla) \mathbf{V}_\sigma + \nabla \cdot \mathbf{p}_\sigma = \frac{q_\sigma}{m_\sigma} n_\sigma (\mathbf{E} + \mathbf{V}_\sigma \times \mathbf{B}), \\ \partial_t \mathbf{p}_\sigma + \nabla \cdot (\mathbf{V}_\sigma \mathbf{p}_\sigma) + 2 \{ \mathbf{p}_\sigma \cdot \nabla \mathbf{V}_\sigma \} + \nabla \cdot \mathbf{q}_\sigma = \frac{2q_\sigma}{m_\sigma} \{ \mathbf{p}_\sigma \times \mathbf{B} \}, \end{cases} \quad (\text{A.2})$$

which needs to be closed by supplying an additional expression for \mathbf{q}_σ (or $\nabla \cdot \mathbf{q}_\sigma$) in terms of the lower moments. Here, $\{ \cdot \}$ denotes symmetrisation, so that e.g.

$$\{ \mathbf{a} \mathbf{b}^{(2)} \} = \frac{1}{3} (\mathbf{a} \mathbf{b}^{(2)} + \mathbf{b} \mathbf{a} \mathbf{b} + \mathbf{b}^{(2)} \mathbf{a}), \quad (\text{A.3})$$

and $\mathbf{p}_\sigma \times \mathbf{B}$ should be interpreted as the two-tensor with elements $[\mathbf{p}_\sigma \times \mathbf{B}]_{ij} = \epsilon_{jkl} p_{\sigma ik} B_l$.

In 1-D electron–proton set-ups like those of interest to us, where only the electron dynamics are important, the 10-moment fluid model (sans closure) together with Maxwell’s equations simplifies to

$$\begin{cases} n (\partial_t + V \partial_x) V + \partial_x p = -\frac{e}{m_e} n E, \\ (\partial_t + V \partial_x) p + 3p \partial_x V + \partial_x q = 0, \\ \partial_x E = \frac{e}{\epsilon_0} (\bar{n} - n), \\ \partial_t E = \frac{e}{\epsilon_0} n V. \end{cases} \quad (\text{A.4})$$

Note that the two remaining Maxwell’s equations imply the continuity equation. Here, we have taken the ions to be immobile with number density equal to the average electron density \bar{n} to ensure quasi-neutrality. Now, let us consider a small wave-like perturbation around equilibrium in the CoM frame, i.e.

$$\begin{cases} n = \bar{n} + \tilde{n} e^{i(kx - \omega t)}, \\ V = \tilde{V} e^{i(kx - \omega t)}, \\ p = n v_{\text{th}}^2, \quad v_{\text{th}} = \bar{v}_{\text{th}} + \tilde{v}_{\text{th}} e^{i(kx - \omega t)}, \\ q = \bar{q} + \tilde{q} e^{i(kx - \omega t)}, \\ E = \tilde{E} e^{i(kx - \omega t)}. \end{cases} \quad (\text{A.5})$$

Here, we assume the wavenumber k to be real, but the frequency $\omega = \omega_r + i\gamma$ is allowed to be complex, γ being the growth rate. Of course, one could equally well set $\omega = \omega_r - i\gamma$, taking γ as the decay rate.

Inserting ansatz (A.5) into (A.4) and keeping only terms up to first-order in the perturbations, we get the relations making up 1-D linear collisionless theory:

$$\begin{cases} -i\omega\tilde{V} + ik(\tilde{n}\tilde{v}_{th}^2 + 2\tilde{n}\tilde{v}_{th}\tilde{v}_{th}) = -\frac{e}{m_e}\tilde{n}\tilde{E}, \\ -i\omega(\tilde{n}\tilde{v}_{th}^2 + 2\tilde{n}\tilde{v}_{th}\tilde{v}_{th}) + 3ik\tilde{n}\tilde{v}_{th}^2\tilde{V} + ik\tilde{q} = 0, \\ ik\tilde{E} = -\frac{e}{\epsilon_0}\tilde{n}, \\ -i\omega\tilde{E} = \frac{e}{\epsilon_0}\tilde{n}\tilde{V}, \end{cases} \tag{A.6}$$

or equivalently

$$\begin{cases} \tilde{E} = i\frac{e}{k\epsilon_0}r\tilde{n} \\ \tilde{V} = rv_{ph}, \\ \tilde{v}_{th} = -\frac{1}{2}\left(1 + \frac{\omega_{pe}^2 - \omega^2}{k^2\tilde{v}_{th}^2}\right)r\tilde{v}_{th}, \\ \tilde{q} = -\left(3 + \frac{\omega_{pe}^2 - \omega^2}{k^2\tilde{v}_{th}^2}\right)r\tilde{n}\tilde{v}_{th}^2v_{ph}. \end{cases} \tag{A.7}$$

To make the notation neater, we have introduced the (complex) phase velocity $v_{ph} = \omega/k$ and the shorthand notation r for \tilde{n}/\tilde{n} , quantifying the amplitude of the perturbation. We have also introduced the electron plasma frequency $\omega_{pe} = \sqrt{e^2\tilde{n}/m_e\epsilon_0}$. Note that when the ion dynamics is negligible, any heat flux closure would need to agree with the expression for \tilde{q} to first-order in r to be viable for modelling weak wave-like perturbations.

A.1 Evaluating the six-term closure

For a closure to be consistent with theory, the heat flux perturbation amplitude \tilde{q} given by the closure must agree with the final part of (A.7), i.e.

$$\tilde{q} = -\left(3 + \frac{\omega_{pe}^2 - \omega^2}{k^2\tilde{v}_{th}^2}\right)r\tilde{n}\tilde{v}_{th}^2v_{ph}, \tag{A.8}$$

to first-order in r . To see whether this is indeed the case, let us calculate the contribution to \tilde{q} from each of the terms in the six-term model. For the terms in q_{even} , we get

$$\begin{cases} A_1(\tilde{n}v_{th}^2V + 2nv_{th}\tilde{v}_{th}V + nv_{th}^2\tilde{V}) & = A_1r\tilde{n}\tilde{v}_{th}^2v_{ph}, \\ ikA_2(3\tilde{n}v_{th}^2\tilde{v}_{th} + \tilde{n}v_{th}^3) & = ikA_2r\tilde{n}\tilde{v}_{th}^3, \\ ikA_3(\tilde{n}v_{th}^2\tilde{v}_{th} + 2nv_{th}\tilde{v}_{th}^2 + nv_{th}^2\tilde{v}_{th}) & = -\frac{1}{2}ikA_3\left(1 + \frac{\omega_{pe}^2 - \omega^2}{k^2\tilde{v}_{th}^2}\right)r\tilde{n}\tilde{v}_{th}^3, \end{cases} \tag{A.9}$$

keeping only terms up to first-order in r . In other words,

$$\tilde{q}_{\text{even}} = -\left[-A_1 + ik\left(\frac{1}{2}A_3 - A_2\right)\frac{\tilde{v}_{th}}{v_{ph}} + \frac{1}{2}ikA_3\frac{\tilde{v}_{th}}{v_{ph}}\frac{\omega_{pe}^2 - \omega^2}{k^2\tilde{v}_{th}^2}\right]r\tilde{n}\tilde{v}_{th}^2v_{ph}. \tag{A.10}$$

As for the terms in q_{odd} , the constant term A_4 does not contribute to \tilde{q} , but the contribution from the other two terms is non-zero:

$$\begin{cases} A_5 (\tilde{n}v_{\text{th}}^3 + 3nv_{\text{th}}^2\tilde{v}_{\text{th}}) & = -\frac{1}{2}A_5 \left(1 + 3\frac{\omega_{\text{pe}}^2 - \omega^2}{k^2\tilde{v}_{\text{th}}^2}\right) r\tilde{n}\tilde{v}_{\text{th}}^3, \\ ikA_6 (\tilde{n}v_{\text{th}}^2\tilde{V} + 2nv_{\text{th}}\tilde{v}_{\text{th}}\tilde{V} + nv_{\text{th}}^2\tilde{V}) & = ikA_6 r\tilde{n}\tilde{v}_{\text{th}}^2 v_{\text{ph}}. \end{cases} \quad (\text{A.11})$$

Thus, the contribution from q_{odd} is

$$\tilde{q}_{\text{odd}} = - \left[\frac{1}{2}A_5 \frac{\tilde{v}_{\text{th}}}{v_{\text{ph}}} - ikA_6 + \frac{3}{2}A_5 \frac{\tilde{v}_{\text{th}}}{v_{\text{ph}}} \frac{\omega_{\text{pe}}^2 - \omega^2}{k^2\tilde{v}_{\text{th}}^2} \right] r\tilde{n}\tilde{v}_{\text{th}}^2 v_{\text{ph}}. \quad (\text{A.12})$$

Defining

$$\Phi(\omega, k) = \frac{\omega_{\text{pe}}^2 - \omega^2}{k^2\tilde{v}_{\text{th}}^2} \quad (\text{A.13})$$

as well as

$$\alpha = \text{Re}\Phi = \frac{\omega_{\text{pe}}^2 + \gamma^2 - \omega_r^2}{k^2\tilde{v}_{\text{th}}^2}, \quad \beta = -\text{Im}\Phi = \frac{2\omega_r\gamma}{k^2\tilde{v}_{\text{th}}^2}, \quad (\text{A.14})$$

demanding (A.8) hold for our closure is equivalent to demanding

$$(3 + \Phi)\omega = -(A_1 + ikA_6)\omega + \left[\frac{1}{2}A_5 + ik\left(\frac{1}{2}A_3 - A_2\right) \right] k\tilde{v}_{\text{th}} + \frac{1}{2}(3A_5 + ikA_3) k\tilde{v}_{\text{th}}\Phi, \quad (\text{A.15})$$

or equivalently

$$\begin{cases} (3 + \alpha)\omega_r + \beta\gamma = -A_1\omega_r + kA_6\gamma + \frac{1}{2} [kA_3\beta + A_5(1 + 3\alpha)] k\tilde{v}_{\text{th}}, \\ -(3 + \alpha)\gamma + \beta\omega_r = A_1\gamma + kA_6\omega_r + \left[kA_2 - \frac{1}{2}kA_3(1 + \alpha) + \frac{3}{2}A_5\beta \right] k\tilde{v}_{\text{th}}. \end{cases} \quad (\text{A.16})$$

If we now define

$$\Phi_{\pm} = \frac{\omega_{\text{pe}}^2 \pm |\omega|^2}{k^2\tilde{v}_{\text{th}}^2} \quad (\text{A.17})$$

and solve for A_1 and kA_6 , the result can be simplified to the form of (3.3), giving us two constraints on the coefficients which we can check.

Appendix B. Demonstration: recovery of the momentum equation

While the main use of sparse regression in this paper is to discover unknown approximate relations between the heat flux and lower-order fluid quantities, it is useful to verify that our workflow is able to identify known exact relations that the simulation data must obey. To illustrate that this is indeed the case, we here show

Model	$-E$	$-\partial_x T$	$-T \partial_x \ln \frac{n}{\bar{n}}$	$-V \partial_x V$	$\partial_x E$
1.	0.739				
2.	1.234	1.138			
3.	0.985	0.977	1.006		
4.	0.999	0.998	0.996	0.997	
5. (uncorrected E)	1.000	0.998	0.991	0.991	5.00×10^{-4}

TABLE 3. Optimal coefficient values for the terms in the various models found by sparse regression to approximate the momentum equation as given in (B.1). The numbering of the models is the same as in figure 10. The theoretical value of almost all coefficients multiplying the terms shown in the first row is 1. The only exception is the non-physical $\partial_x E$ term, which is only found consistently when the staggering of the E field data is not corrected for (model 5). This coefficient should have the value 5×10^{-4} to provide a first-order correction for the half-cell size shift of the field data.

how one can recover the density-normalised version of the 1-D momentum equation for electrons,

$$\partial_t V = -V \partial_x V - T \partial_x \ln \frac{n}{\bar{n}} - \partial_x T - \frac{e}{m_e} E, \quad (\text{B.1})$$

from our two-stream simulation data. Note that the arbitrary normalisation of n to the unperturbed total electron density \bar{n} does not affect the logarithmic derivative. In the sparse regression, we now set $\partial_t V$ as our target variable y and use a term library including all products of up to two terms from the set consisting of n , V , T and E , their spatial derivatives and a constant term.

We work with the $n_b/n_e = 0.1$ two-stream instability simulation, sampling $\sim 2.5\%$ of the data in each of the 10 cross-validation folds. This time, we take samples from almost the entire time domain, and not just e.g. the growth phase. The sparse regression algorithm yields the sequence of models for $\partial_t V$ shown in figure 10(a).

As shown by the marker shapes – indicating whether models contain consistently the same terms (circles) or different ones (triangles) – SR correctly recovers the momentum equation and finds no further terms consistently. That the model is complete at four terms is also supported by the fact that the accuracy plateaus after this point, at an FVU of $\sim 3.26 \times 10^{-4}$. This signifies that $\sim 99.97\%$ of the variation in $\partial_t V$ can be explained by the 1-D momentum equation as given in (B.1). SR also gives us the coefficients of the terms in the equation with an error of less than half a percent; these are listed in table 3.

Next, we illustrate the importance of appropriately aligning the E -field data with the fluid data in space. Misalignment between particle and field data naturally occurs in PIC schemes using a staggered Yee grid (Yee 1966), where the E - and B -field nodes appear in cell edges and cell faces, respectively, while the particle data are often cell-centred.⁵ This discrepancy then propagates through to any fluid and electromagnetic field data exported from the simulation.

⁵In addition, the leap-frog type time integration scheme results in momentum and position information available half-time step apart from which, unless corrected for (as in the version of OSIRIS we use), may also lead to a reduced accuracy. Note that, unlike the spatial misalignment, the temporal one cannot be corrected by post-processing; the correction needs to be made before the fluid quantities are computed by the simulation code.




FIGURE 10. Sequence of models found by sparse regression to approximate $\partial_t V$, recovering the momentum equation as given in (B.1). The consistently found models are labelled by the new consistently found term which is added to the model compared with the next most simple one. The coefficients for all terms in each model are provided in table 3 (with the same numbering of models as here). Circles, unlike triangles, denote consistently found terms, and blue (orange) marker colour corresponds to testing (training) data (note that they overlap almost completely here). The two panels show the results (a) when using E -field data which has been corrected for the half-cell grid shift with respect to the fluid quantities, which finds only the expected terms, and (b) when using uncorrected E -field data, which also finds a non-physical $\partial_x E$ term adding the correction back in.

In figure 10(b), we show the result of performing SR on the same simulation data without correcting for the misalignment of the electric field data. The first four consistently found terms are the same as before, but now this four-term model is less accurate, with an FVU $> 10^{-3}$. A fifth term $\propto \partial_x E$ is also consistently identified, however, and if retained, increases the accuracy to match that achieved with corrected E -field data. This term corresponds to performing a first-order Taylor expansion to evaluate E half a grid cell further towards $-x$, since $E|_{x-\frac{1}{2}\Delta x} = E|_x - \frac{1}{2}\Delta x \partial_x E|_x$. Specifically, if we insert our spatial resolution $\Delta x = 10^{-3} \delta_e$, we find that

$$E_{\text{aligned}} = E_{\text{shifted}} - 5 \times 10^{-4} \delta_e \partial_x E_{\text{shifted}}, \tag{B.2}$$

corresponding exactly to the correction introduced by SR, since δ_e is our unit of length.

REFERENCES

ACHIKANATH CHIRAKKARA, R., SETA, A., FEDERRATH, C. & KUNZ, M.W. 2023 Critical magnetic Reynolds number of the turbulent dynamo in collisionless plasmas. *Mon. Not. R. Astron. Soc.* **528** (1), 937–953.

ALVES, E.P. & FIUZA, F. 2022 Data-driven discovery of reduced plasma physics models from fully kinetic simulations. *Phys. Rev. Res.* **4** (3), 033192.

ARZAMASSKIY, L., KUNZ, M.W., SQUIRE, J., QUATAERT, E. & SCHEKOCIHIN, A.A. 2023 Kinetic turbulence in collisionless high- β plasmas. *Phys. Rev. X* **13**, 021014.

BORIS, J.P. & SHANNY, R.A. 1972 *Proceedings [of The] 4th Conference on Numerical Simulation of Plasmas*, Naval Research Laboratory.

Braginskii, S.I. 1958 Transport phenomena in a completely ionized two-temperature plasma. *Sov. Phys. JETP* **6** (33), 358–369.

- BRUNTON, S.L., PROCTOR, J.L. & KUTZ, J.N. 2016 Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl Acad. Sci.* **113** (15), 3932–3937.
- CHAPMAN, S. & COWLING, T.G. 1991 *The Mathematical Theory of Non-Uniform Gases: An Account of the Kinetic Theory of Viscosity, Thermal Conduction and Diffusion in Gases*. Cambridge University Press.
- CHENG, W., FU, H., WANG, L., DONG, C., JIN, Y., JIANG, M., MA, J., QIN, Y. & LIU, K. 2023 Data-driven, multi-moment fluid modeling of Landau damping. *Comput. Phys. Commun.* **282**, 108538.
- CHEW, G.F., GOLDBERGER, M.L. & LOW, F.E. 1956 The Boltzmann equation and the one-fluid hydro-magnetic equations in the absence of particle collisions. *Proc. R. Soc. Lond. Series A. Math. Phys. Sci.* **236** (1204), 112–118.
- DAM, M., BRØNS, M., JUUL RASMUSSEN, J., NAULIN, V. & HESTHAVEN, J.S. 2017 Sparse identification of a predator-prey system from simulation data of a convection model. *Phys. Plasmas* **24** (2), 022310.
- DONAGHY, J. & GERMASCHEWSKI, K. 2023 In search of a data-driven symbolic multi-fluid ten-moment model closure. *J. Plasma Phys.* **89** (1), 895890105.
- DONG, C., WANG, L., HAKIM, A., BHATTACHARJEE, A., SLAVIN, J.A., DIBRACCIO, G.A. & GERMASCHEWSKI, K. 2019 Global ten-moment multifluid simulations of the solar wind interaction with mercury: from the planetary conducting core to the dynamic magnetosphere. *Geophys. Res. Lett.* **46** (21), 11584–11596.
- FONSECA, R.A., MARTINS, S.F., SILVA, L.O., TONGE, J.W., TSUNG, F.S. & MORI, W.B. 2008 One-to-one direct modeling of experiments and astrophysical scenarios: pushing the envelope on kinetic plasma simulations. *Plasma Phys. Control. Fusion* **50** (12), 124034.
- FONSECA, R.A. *et al.* 2002 OSIRIS: a three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators. In *Computational Science – ICCS. 2002*, (ed. Sloot, P.M.A., Hoekstra, A.G., Tan, C.J.K. & Dongarra, J.J.), pp. 342–351. Springer.
- GRAD, H. 1949 On the kinetic theory of rarefied gases. *Commun. Pure Appl. Maths* **2** (4), 331–407.
- GUREVICH, D.R., GOLDEN, M.R., REINBOLD, P.A.K. & GRIGORIEV, R.O. 2024 Learning fluid physics from highly turbulent data using sparse physics-informed discovery of empirical relations (SPIDER). *J. Fluid Mech.* **996**, A25.
- HAKIM, A. 2008 Extended MHD modelling with the ten-moment equations. *J. Fusion Energy* **27** (1–2), 36–43.
- HAKIM, A., LOVERICH, J. & SHUMLAK, U. 2006 A high resolution wave propagation scheme for ideal two-fluid plasma equations. *J. Comput. Phys.* **219** (1), 418–442.
- HAMMETT, G.W. & PERKINS, F.W. 1990 Fluid moment models for Landau damping with application to the ion-temperature-gradient instability. *Phys. Rev. Lett.* **64** (25), 3019–3022.
- HOCKNEY, R.W. & EASTWOOD, J.W. 2021 *Computer Simulation Using Particles*. CRC Press.
- JACOBSEN, E. & KOOTSOOKOS, P. 2007 Fast, accurate frequency estimators [DSP tips & tricks]. *IEEE Signal Process. Mag.* **24** (3), 123–125.
- JOGLEKAR, A.S. & THOMAS, A.G.R. 2023 Machine learning of hidden variables in multiscale fluid simulation. *Mach. Learn.: Sci. Technol.* **4** (3), 035049.
- KAHEMAN, K., KUTZ, J.N. & BRUNTON, S.L. 2020 SINDy-PI: a robust algorithm for parallel implicit sparse identification of nonlinear dynamics. *Proc. R. Soc. A: Math. Phys. Engng Sci.* **476** (2242), 20200279.
- KAPTANOGLU, A.A., CALLAHAM, J.L., ARAVKIN, A., HANSEN, C.J. & BRUNTON, S.L. 2021a Promoting global stability in data-driven models of quadratic nonlinear dynamics. *Phys. Rev. Fluids* **6** (9), 094401.
- KAPTANOGLU, A.A., HANSEN, C., LORE, J.D., LANDREMAN, M. & BRUNTON, S.L. 2023 A sparse regression for plasma physics. *Phys. Plasmas* **30** (3), 033906.
- KAPTANOGLU, A.A., MORGAN, K.D., HANSEN, C.J. & BRUNTON, S.L. 2021b Physics-constrained, low-dimensional models for magnetohydrodynamics: first-principles and data-driven approaches. *Phys. Rev. E* **104** (1), 015206.
- KAPTANOGLU, A.A., ZHANG, L., NICOLAOU, Z.G., FASEL, U. & BRUNTON, S.L. 2023b Benchmarking sparse system identification with low-dimensional chaos. *Nonlinear Dyn.* **111** (14), 13143–13164.

- KHOTYAINITSEV, Y.V., GRAHAM, D.B., NORNGREN, C. & VAIVADS, A. 2019 Collisionless magnetic reconnection and waves: progress review. *Front. Astron. Space Sci.* **6**, 70.
- LEVERMORE, C.D. 1996 Moment closure hierarchies for kinetic theories. *J. Stat. Phys.* **83** (5), 1021–1065.
- MAKKE, N. & CHAWLA, S. 2024 Interpretable scientific discovery with symbolic regression: a review. *Artif. Intell. Rev.* **57** (1), 2.
- MAULIK, R., GARLAND, N.A., BURBY, J.W., TANG, X.-Z. & BALAPRAKASH, P. 2020 Neural network representability of fully ionized plasma fluid model closures. *Phys. Plasmas* **27** (7), 072106.
- MCGRAE-MENGE, M. C., PIERCE, J. R., FIUZA, F. & ALVES, E. P. 2023 Embedding Lorentz covariance in data-driven reduced plasma models from fully kinetic simulations. In *65th Annual Meeting of the APS Division of Plasma Physics Meeting Abstracts, APS Meeting Abstracts*, vol. **2023**, p. UO04.14.
- MESSINGER, D.A. & BORTZ, D.M. 2021 Weak SINDy: galerkin-based data-driven model selection. *Multiscale Model. Simul.* **19** (3), 1474–1497.
- NG, J., HAKIM, A., WANG, L. & BHATTACHARJEE, A. 2020 An improved ten-moment closure for reconnection and instabilities. *Phys. Plasmas* **27** (8), 082106.
- NG, J., HUANG, Y.-M., HAKIM, A., BHATTACHARJEE, A., STANIER, A., DAUGHTON, W., WANG, L. & GERMASCHEWSKI, K. 2015 The island coalescence problem: scaling of reconnection in extended fluid models including higher-order moments. *Phys. Plasmas* **22** (11), 112104.
- QIN, Y., MA, J., JIANG, M., DONG, C., FU, H., WANG, L., CHENG, W. & JIN, Y. 2023 Data-driven modeling of Landau damping by physics-informed neural networks. *Phys. Rev. Res.* **5** (3), 033079.
- RUDY, S.H., BRUNTON, S.L., PROCTOR, J.L. & KUTZ, J.N. 2017 Data-driven discovery of partial differential equations. *Sci. Adv.* **3** (4), e1602614.
- SCHAEFFER, H. 2017 Learning partial differential equations via data discovery and sparse optimization. *Proc. R. Soc. A: Math. Phys. Engng Sci.* **473** (2197), 20160446.
- SCHAEFFER, H. & MCCALLA, S.G. 2017 Sparse model selection via integral terms. *Phys. Rev. E* **96** (2), 023302.
- SHARMA, P., HAMMETT, G.W., QUATAERT, E. & STONE, J.M. 2006 Shearing box simulations of the MRI in a collisionless plasma. *Astrophys. J.* **637** (2), 952–967.
- SHI, F., LIN, Y., WANG, X., WANG, B. & NISHIMURA, Y. 2021 3-D global hybrid simulations of magnetospheric response to foreshock processes. *Earth Planets Space* **73** (1), 138.
- ST-ONGE, D.A., KUNZ, M.W., SQUIRE, J. & SCHEKOCIHIN, A.A. 2020 Fluctuation dynamo in a weakly collisional plasma. *J. Plasma Phys.* **86** (5), 905860503.
- STIX, T.H. 1992 *Waves in Plasmas*. 1st edn. American Institute of Physics.
- TENBARGE, J.M., NG, J., JUNO, J., WANG, L., HAKIM, A.H. & BHATTACHARJEE, A. 2019 An extended MHD study of the 16 October 2015 MMS diffusion region crossing. *J. Geophys. Res.: Space Phys.* **124** (11), 8474–8487.
- VASEY, G., MESSENGER, D., BORTZ, D., CHRISTLIEB, A. & O'SHEA, B. 2025 Influence of initial conditions on data-driven model identification and information entropy for ideal MHD problems. *J. Comput. Phys.* **524**, 113719.
- WANG, L., HAKIM, A.H., BHATTACHARJEE, A. & GERMASCHEWSKI, K. 2015 Comparison of multi-fluid moment models with particle-in-cell simulations of collisionless magnetic reconnection. *Phys. Plasmas* **22** (1), 012108.
- WANG, L., XU, X., ZHU, B., MA, C. & LEI, Y.-A. 2020 Deep learning surrogate model for kinetic Landau-fluid closure with collision. *AIP Adv.* **10** (7), 075108.
- YEE, K. 1966 Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media. *IEEE Trans. Antennas Propag.* **14** (3), 302–307.

**Data-driven multi-species heat flux closures for
two-stream-unstable plasmas with nonlinear
sparse regression**

E. R. Ingelsten, M. C. McGrae-Menge, E. P. Alves and I. Pusztai

Accepted for publication in *Physics of Plasmas*.

Data-driven multi-species heat flux closures for two-stream-unstable plasmas with nonlinear sparse regression

Emil R. Ingelsten,^{1, a)} Madox C. McGrae-Menge,² E. Paulo Alves,^{2, 3} and Istvan Pusztai¹

¹⁾*Department of Physics, Chalmers University of Technology, Göteborg, SE-41296, Sweden*

²⁾*Department of Physics and Astronomy, University of California, Los Angeles, CA 90095, USA*

³⁾*Mani L. Bhaumik Institute for Theoretical Physics, University of California at Los Angeles, Los Angeles, CA 90095, USA*

(Dated: 25 March 2026)

The dual aims of accuracy and computational efficiency in computational plasma physics lend themselves well to the use of fluid models. The first of these goals, however, is only satisfied for such models insofar as the utilized closure can capture the neglected kinetic physics—something which has proven challenging for multi-scale collisionless processes. In a recent article [E. R. Ingelsten et al. (2025) *J. Plasma Phys.* **91** E64], we used the data-driven method of sparse regression to discover a novel heat flux closure for electrostatic phenomena. Here, we generalize the six-term closure model found in that work from single- to multi-species modeling. Using data from OSIRIS particle-in-cell simulations over a range of initial conditions, we then demonstrate how the unknown coefficients in front of the three most important terms in the closure can be estimated from box-averaged fluid quantities. Both neural networks and a newly developed framework for nonlinear sparse regression are showcased. The resulting models predict the heat flux for each species with a typical accuracy of 80–90% and regularly account for 85–95% of the rate of change in the pressure. The models are also compared with results from multi-species linear collisionless theory.

I. INTRODUCTION

Fast and accurate plasma modeling has been a central goal of numerical plasma physics since its inception, but achieving both simultaneously is rare. Fully kinetic methods offer high accuracy but are too computationally expensive for inherently multi-scale, three-dimensional problems. Fluid models are therefore often used—either on their own^{1–4} or as part of a kinetic-fluid hybrid approach^{5–7} offering much greater efficiency with only limited loss of accuracy. The reliability of fluid-based approaches, however, depends on the closure’s ability to capture neglected kinetic effects—a task which for many systems is highly challenging.

For collisional plasmas, where particles generally remain close to local thermal equilibrium, one can derive closures rigorously from first principles^{8–12}. The same is generally not true in the collisionless case, however—there are no known generally applicable closures for collisionless systems. Nonlinear and nonlocal kinetic processes have proven particularly challenging to model well without fully resolving the underlying kinetic physics. Such processes are nevertheless ubiquitous in many space and astrophysical contexts—turbulence^{13–17}, electron holes^{18–22}, wave-particle interactions^{14, 23–26} and magnetic reconnection^{16, 17, 26–33} being prominent examples^{34, 35}. In other words, there is strong reason to investigate potential avenues to improved closure construction, so as to capture the net effects of small-scale kinetic phenomena without in-detail modeling.

To this end, several approaches have been explored. Traditionally, theory has been the starting point, employing simplifying assumptions to derive collisionless closures valid in restricted limits. For instance, the Hammett–Perkins closure^{36, 37}, which models Landau damping, assumes linearity, the Chew–Goldberger–Low (CGL) closure³⁸ assumes adiabaticity and full magnetization, and the Le closure^{39, 40}, interpolating between the CGL and isothermal closures, assumes magnetization as well as an electron thermal speed much larger than the Alfvén speed. These assumptions often fail in many collisionless regimes of interest, such as magnetic reconnection, turbulence, and instability saturation. Nevertheless, theoretically derived closures remain widely used—sometimes even beyond their regimes of validity, as their well-defined physical basis clarifies which processes are accurately represented and which are not. Ad-hoc closures such as the relaxation closure⁴¹, which forces the pressure tensor toward isotropy, have also been used, with mixed success compared to kinetic simulations.

A different approach, originating in the field of machine learning, is data-driven closure development. With the rapid growth of this field, numerous methods have emerged. Neural networks^{42–44} are perhaps the most well-known, not least through their role in generative artificial intelligence. Another method of interest is symbolic regression^{45–48}, which typically uses evolutionary algorithms to construct parsimonious models by encoding constituent coefficients, dependent variables and elementary functions in a graph structure. Most central to this work, however, is sparse regression^{49–53}, which more expressly focuses model space search along the Pareto front between sparsity and accuracy for a pre-defined func-

^{a)}Electronic mail: emilraa@chalmers.se

tion library, so as to reduce computational complexity and improve scalability. The latter two methods generally yield more interpretable models than neural network-based approaches, due to the black-box nature of neural networks—though often at the cost of a slight decrease in expressive power. Given the ubiquity of data-driven model construction in science and technology, all of these methods have been applied across many disciplines^{54–75}.

Recent examples of neural network-based approaches to closure construction include Refs. 76–80. In Ref. 76, a reservoir computing architecture was used to construct a closure for a one-dimensional electrostatic Vlasov-Poisson system expressed in a pseudo-spectral Fourier-Hermite basis. Similarly, Ref. 77 used Fourier neural operators to model nonlinear Landau damping. In Ref. 78, on the other hand, time-embedded convolutional neural networks were used to model the heat flux and the ratio between the mean free path length and the characteristic length scale of electron temperature variations. Sparse regression has also started to see increased use in plasma physics over the past few years^{81–85}, but for closure discovery specifically there has only been limited exploration of its utility thus far. In Ref. 86, it was used to recover the collisionless moment equations and discover a heat flux closure for the strongly nonlinear dynamics involved in two-dimensional Harris-sheet reconnection. The linear regime was left out of the analysis, however, and no interpretation of the closure was given. Similarly, sparse regression was used in Ref. 87 in conjunction with deep neural networks to recover the fluid equations as well as the local approximation^{3,88} of the Hammett-Perkins closure for a one-dimensional setup exhibiting linear Landau damping of Langmuir waves. More recently, the SINDy algorithm for sparse regression^{52,53} was used in Ref. 89 to discover a six-term heat flux closure capable of describing the Landau-damping of Langmuir waves and two-stream instability—both in the linear and nonlinear regimes, including the local approximation of the Hammett-Perkins closure as one of the terms. Two important questions were left unanswered in that work, however.

Firstly, only fluid models with a single combined electron species were considered. This would be impractical for simulation of two-stream-unstable setups, since the instability must then be imposed manually through the closure unless one includes higher-order fluid moments (which would entail increased closure complexity and could potentially lead to complications with regards to ensuring hyperbolicity⁹⁰). This raised the question of whether the six-term closure could be generalized to multi-species electron fluids. Secondly, the closure coefficients were free parameters found to vary with plasma conditions. For full fluid-code implementation, it would thus be desirable to estimate these from other fluid quantities.

In this article, we address both of these questions. We extend the six-term closure to multiple species and examine how its coefficients vary with plasma conditions. We

then show that the three most important coefficients can be accurately predicted from box-averaged fluid quantities using neural networks. With the aim of obtaining more interpretable expressions, we further adapt the SINDy sparse regression framework to perform nonlinear sparse regression with rational functions, achieving neural network-level accuracy. Finally, we describe how one can leverage the Bernstein polynomial basis to avoid poles on the domain of interest at only a limited cost in expressive power.

This article is structured as follows: Following a short reflective note on the scope and assessment of data-driven closures, in Sec. II we describe the sparse regression and neural network-based methodologies we employ, as well as the simulation setup used to generate our two-stream instability dataset. In Sec. III, we first describe how the previously identified six-term heat flux model can be generalized to multi-species setups, and then illustrate how the three most important terms can be estimated from box-averaged fluid quantities, discussing the results in the context of linear collisionless theory. Finally, we summarize our results, discuss avenues towards further improvements and give an outlook towards future work in Sec. IV. We also include Appendix A, describing multi-species collisionless theory and the constraints it places on heat flux closures of our type, Appendix B, containing a detailed overview of the relationship between separate- and combined-species fluid quantities, as well as Appendix C, recapitulating the single-species closure described in Ref. 89 and discussing the physical interpretation of the included terms.

A. On the Objectives and Assessment of Data-Driven Fluid Closures

Data-driven collisionless closures can enhance the physical fidelity of fluid simulations by effectively capturing certain kinetic processes in regimes where analytical treatments are difficult (e.g., strongly nonlinear dynamics), as well as by representing underresolved phenomena in large-scale simulations. Owing to the enormous diversity and strong parameter dependence of collisionless processes, any such closure – whether data-driven or otherwise – can realistically capture only a limited set of mechanisms over a finite region of parameter space. Guidance from kinetic theory is therefore essential in identifying the dominant microphysical processes relevant to a given environment. For example, closures designed for high-beta astrophysical plasma turbulence may focus on gyro-scale, pressure-anisotropy-driven instabilities whose large-scale effects resemble pitch-angle-scattering collisions, thereby breaking adiabatic invariance in systems that would otherwise evolve nearly ideally.

The ultimate test of a data-driven closure is the degree of agreement between simulations produced by its implementation in a fluid solver and corresponding results from kinetic simulations. Certain aspects of closure

performance can be quantified straightforwardly, such as accuracy relative to kinetic simulations over short or long timescales (using appropriate error metrics, see footnote⁹¹) and numerical complexity, ensuring that the fluid model retains a computational advantage over kinetic approaches. Other desirable properties are more qualitative or essentially binary, including numerical stability and the satisfaction of fundamental physical requirements such as conservation laws and symmetries.

More broadly, an effective closure should ideally capture parametric dependencies across wide regions of parameter space and, if possible, represent multiple physical phenomena. Achieving such generalization beyond the training regime remains notoriously challenging. In this context, interpretability – although loosely defined – can be valuable by enabling analytical insight into the behavior of the closure. The value of a data-driven model can also lie in its ability to inform analytic theory, for example by identifying which variables or terms are most important for a given process, or by indicating how much nonlocal information is required to achieve accurate predictions. Finally, the practical feasibility of constructing a closure also depends on the ability to generate a sufficient amount of training data at reasonable cost.

II. METHODS

When constructing heat flux closures, we start by performing particle-in-cell (PIC) simulations of the modeled system—in this case a two-stream-unstable electron-proton plasma (see Sec. II C)—using the OSIRIS code^{92,93}. From the kinetic simulation data we then export the fluid quantities relevant to us, namely the number density $n_\sigma = \int d\mathbf{v} f_\sigma$, the flow velocity $\mathbf{V}_\sigma = n_\sigma^{-1} \int d\mathbf{v} \mathbf{v} f_\sigma$, the mass-normalized pressure $\mathbf{p}_\sigma = \int d\mathbf{v} (\mathbf{v} - \mathbf{V}_\sigma)^{(2)} f_\sigma$ and the mass-normalized heat flux $\mathbf{q}_\sigma = \int d\mathbf{v} (\mathbf{v} - \mathbf{V}_\sigma)^{(3)} f_\sigma$, as well as electric and magnetic field data (\mathbf{E} and \mathbf{B}), at regular time intervals. In these expressions, $f_\sigma(\mathbf{v})$ is the distribution function for species σ , and we use notation where $[\mathbf{ab}]_{ij} = a_i b_j$ and $\mathbf{a}^{(2)}$ is shorthand for \mathbf{aa} . Note also that we for convenience omit the spatiotemporal dependence of the distribution function, fluid quantities and electromagnetic field in this notation. Furthermore, since the PIC simulations here are one-dimensional in position space, we will henceforth drop the vector and tensor notation for fluid and electromagnetic field quantities unless otherwise noted, using p_σ and E to mean $p_{\sigma,11}$ and E_1 . Additionally, the processes we are modeling occur at electron timescales. Thus, unless otherwise noted, we will only refer to electron fluid quantities, and fluid quantities without explicit species labeling, like n , should be interpreted as referring to the combined electron species. We treat the counterstreaming electron populations separately, each characterized by its density flow velocity and temperature, and refer to the one with higher (lower) density as the core

(beam) population.

As in Ref. 89, we use a version of OSIRIS which corrects for the temporal staggering of particle position and velocity information when generating the data used for the data-driven closure models in this work^{94,95}. We additionally correct for spatial grid staggering via linear interpolation (see Appendix B of Ref. 89 for a more detailed discussion of this).

Having conducted such simulations with a range of initial conditions, we use a version of the SINDy sparse regression algorithm, as described in Sec. II A 1, to find optimally accurate and maximally sparse heat flux closures valid in the parameter regimes in question. Reducing the discovered models for the beam and core electron heat fluxes to the three most important terms, we proceed to use neural networks as outlined in Sec. II B to show that it is possible to predict the three most important free closure parameters at quite high accuracy from the box-averaged beam density \bar{n}_b , relative flow velocity $\bar{V}_{\text{rel}} = \bar{V}_c - \bar{V}_b$, beam thermal speed $\bar{v}_{\text{th},b}$ and core thermal speed $\bar{v}_{\text{th},c}$. Here we have also introduced the species labels $\sigma = b$ and $\sigma = c$ for the beam and core electron species, respectively. For a more detailed description of the two-species treatment of the electrons in our PIC simulations, as well as how the simulations were set up in general, see Sec. II C. Finally, we use nonlinear sparse regression described in Secs. II A 2 and II A 3 to discover rational models for both beam, core and combined electron model coefficients $A_{1,4,5}$ in terms of these four box-averaged fluid quantities.

A. Sparse regression

Sparse regression (*SR*) refers to a large and growing family of function fitting algorithms, all with the common goal of finding models for a target variable y which are simultaneously optimized for both *accuracy* and *model sparsity*. Optimizing for sparsity helps limit overfitting, identifies which parts of the model in question are most important and yields more interpretable results. Unlike some alternative methods with similar aims like symbolic regression, sparse regression performs this optimization by picking terms from a pre-defined term library $\{\theta_j\}$. This enables faster optimization, at the cost of leaving the selection of the term library, which can be highly nontrivial, up to the user.

Of particular interest are algorithms based on SINDy, which applies sparse regression to identify the governing differential equations of a system⁵². For partial differential equations, this approach is often referred to as PDE-FIND^{53,96}. These methods start by optimizing a “maximalist” model including all terms in the library. Terms are then iteratively pruned: the least important term is removed, the model is re-optimized, and this process is then repeated until only one term remains. This yields a sequence of increasingly sparse models that approximately trace the Pareto front between accuracy and

sparsity.

1. SINDy

In the original SINDy algorithm, as in most data-driven approaches, one first needs to acquire a dataset of target quantity values $\{y_i\}$ —e.g., corresponding to measurements of y at different points in space and time, along with a corresponding dataset of θ_j values $\{\Theta_{ij}\}$ from the same points in spacetime. The succession of more and more sparse y models is then constructed through linear combination of the terms in the term library: each model is defined according to $\hat{y} = \sum_j \theta_j \xi_j = \boldsymbol{\theta} \cdot \boldsymbol{\xi}$ with increasingly sparse coefficient vectors $\boldsymbol{\xi}$. Consistently, we will in this section use i for the index of a sample or datapoint within a dataset, and use j, k and l for selecting or summing over terms in the term library.

Model optimization at each level of sparsity is performed through linear least-squares regression, finding the coefficient vector $\boldsymbol{\xi}^*$ which minimizes the least-squares cost function

$$C(\boldsymbol{\xi}) = \|\mathbf{y} - \boldsymbol{\Theta}\boldsymbol{\xi}\|^2 = \sum_i \left(y_i - \sum_j \Theta_{ij} \xi_j \right)^2, \quad (1)$$

where $\boldsymbol{\Theta}$ is the term library matrix with elements equal to Θ_{ij} , and $\hat{\mathbf{y}} = \boldsymbol{\Theta}\boldsymbol{\xi}$ is the model vector, containing the value of the model at each point in the dataset.

Having found this optimal coefficient vector, the version of SINDy we use in this work uses sequentially thresholded least-squares (STLS) regression, as proposed in the original SINDy article⁵²—see also e.g. Refs. 83, 85, and 89—to determine which term should be considered least important (and thus be removed before the next round of optimization). Specifically, the deleted term is the one whose coefficient has the smallest absolute value when using a term library normalized according to term variance (i.e., with $\Theta_{ij} \rightarrow \Theta_{ij}/\text{var}_i(\Theta_{ij})$). Alternatively, one may view this as working with an arbitrary (unnormalized) term library and discarding the term with index

$$j^{\text{discarded}} = \underset{j}{\text{argmin}} \left[\text{var}_i(\Theta_{ij}) |\xi_j^*| \right]. \quad (2)$$

This term is then excluded from all following models. Since constant terms have zero variance, they are not normalized in this way—instead, both the term itself and its “variance” normalization factor are set to unity when such terms are included. To protect against overfitting, we further perform 10-fold cross-validation (i.e., dividing data into 10 groups, each serving as a test set once, with the others as training data, ensuring all data points are used for validation).

Similarly to Ref. 89, this version of SINDy is essentially what we use when searching for heat flux closures. When treating the counter-streaming populations as separate species σ , we specifically take $y = q_\sigma$, and use a term

library analogous to the one we used in the single-species case, with one modification: We consider terms of the form

$$\theta_j = n_\sigma v_{\text{th},\sigma}^\alpha (V_\sigma - \bar{V}_\sigma)^{3-\alpha}, \quad (3)$$

for integer $\alpha \leq 3$ and with $v_{\text{th}} = \sqrt{T} = \sqrt{p/n}$, where $T = p/n$ is the (11-component of the) mass-normalized temperature tensor. We also include terms containing first-order derivatives of these quantities, e.g. $\partial_x(n_\sigma v_{\text{th},\sigma} \partial_x(v_{\text{th},\sigma})(V_\sigma - \bar{V}_\sigma)$. All library terms are thus manifestly invariant under Galilean transformations, which is desirable, since the quantity we are attempting to model, q_σ , is also Galilean-invariant. If this were not the case, Galilean invariance would have to be ensured through other means, e.g. by enlarging the dataset via the application of random Galilean boosts, as described in Ref. 85. Note also that the replacement $V \rightarrow V_\sigma - \bar{V}_\sigma$ here compared to the single-species case is consistent with the fact that we are working in the combined-species center-of-mass frame, meaning the combined-species V variable should really be considered shorthand for $V - \bar{V}$. That this really is the correct choice when generalizing to the multi-species case can also be motivated from multi-species linear theory—see Appendix A.

When the y and θ_j quantities are given on a spacetime grid, as is the case for $y = q_\sigma$, we additionally perform spatiotemporal averaging on the raw simulation data around each sampled point in spacetime to calculate our y_i and Θ_{ij} quantities. Specifically, we average over a spacetime box 5 by 5 datapoints in size centered on the spacetime point in question. This procedure, which reduces noise and helps with SR convergence, may be considered a special case of weak SINDy⁹⁷, with an integration kernel equal to a spacetime box function. A smoother integration kernel, as implemented in the WSINDy and SPIDER sparse regression frameworks of Refs. 98 and 99, may be preferred for term libraries featuring higher than first-order derivatives.

However, we do not perform any such averaging procedure when we use SR to estimate the closure coefficients $A_{1,4,5}$ from box-averaged fluid quantities, since the $A_{1,4,5}$ coefficients in the dataset are already effectively averaged over 20 time steps (and the entire spatial domain), due to being calculated via sparse regression over such time slices in the first place.

When performing this secondary sparse regression step to create a three-term heat flux closure usable without supplying ad-hoc values for the closure coefficients, it quickly becomes apparent that using SR with a naive term library consisting of multivariate polynomials in $\{\bar{n}_b/\bar{n}, \bar{V}_{\text{rel}}/c, \bar{v}_{\text{th},b}/c, \bar{v}_{\text{th},c}/c\}$ (where c denotes the speed of light) is insufficient to adequately model $A_{1,4,5}$ at reasonable term library sizes. However, since neural networks with the same four inputs, designed and trained as described in Sec. IIB, are able to express $A_{1,4,5}$ in terms of these box-averaged quantities quite accurately

(see Tab. I), the main issue is clearly that a polynomial term library has insufficient expressive power.

Inspired by the high efficiency of Padé approximants compared to truncated Taylor expansions for achieving high accuracy at relatively low complexity, we look towards a rational function-based term library for a solution to this problem. However, such a library is difficult to implement in regular SINDy, since the method is based on *linear* combination of library terms, which is not well suited to handle the addition or removal of terms in the denominator of the rational model. There are extensions to SINDy which attempt to solve this problem without sacrificing the linearity of the least-squares optimization lying at the core of SINDy, such as implicit-SINDy¹⁰⁰ and SINDy-PI¹⁰¹. However, such approaches require a low level of “noise” in the y and θ_j data, i.e. that there exist a near-perfect model reachable with the used term library, which is not the case for us. In our case, some closure coefficients have minimum FVU error rates on the order of $\sim 20\%$, judging by the performance of our neural network models (described in Sec. II B). Thus, we need a different solution.

The FVU, or the *Fraction of Variance Unexplained*, is the metric we use to quantify the error for a model \hat{y} of a quantity y over a given dataset, and is defined as

$$\text{FVU}[\hat{y}] = \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}. \quad (4)$$

Here, y_i is the true and $[\hat{y}]_i = \hat{y}_i$ is the predicted value of quantity y at the i^{th} datapoint, while \bar{y} is the mean of y over the dataset.

2. Nonlinear sparse regression

When seeking to model $A_{1,4,5}$ in terms of \bar{n}_b/\bar{n} , \bar{V}_{rel}/c , $\bar{v}_{\text{th},b}/c$ and $\bar{v}_{\text{th},c}/c$, we opted to replace the linear SINDy model $\hat{y} = \Theta \xi$ with the nonlinear model calculating \hat{y} via

$$[\hat{y}]_i = \frac{\sum_{j=0}^{M-1} \Theta_{ij} \xi_j}{1 + \sum_{k=1}^{M-1} \Theta_{ik} \zeta_k}. \quad (5)$$

Here, we use a convention where the index $j = 0$ corresponds to the constant term and the library has M terms in total. This is effectively the same as considering models $[\hat{y}]_i = [\Theta \xi]_i / [\Theta \zeta]_i$ with the zeroth component of ζ set to one to make each model correspond uniquely to a specific combined coefficient vector $\xi' = (\xi, \zeta)$ by removing the degree of freedom corresponding to simply multiplying both ξ and ζ with the same constant. Note that this means that our ζ has length $M - 1$, while ξ has length M .

We want to use these nonlinear models to express $A_{1,4,5}$ as rational functions in terms of $\{\bar{n}_b/\bar{n}, \bar{V}_{\text{rel}}/c, \bar{v}_{\text{th},b}/c, \bar{v}_{\text{th},c}/c\}$, as mentioned above. To curb overfitting and encourage SR convergence,

we limit our term library $\{\theta_j\}$ to containing only multivariate monomials of the four dependent variables of order 0, 1 and 2, so that $M = 3^4 = 81$, giving us a combined coefficient vector ξ' of length $2M - 1 = 161$. Our term library thus includes e.g. $\bar{n}_\sigma \bar{V}_{\text{rel}}^2 / (\bar{n} c^2)$, but not $\bar{n}_\sigma \bar{V}_{\text{rel}}^3 / (\bar{n} c^3)$. Whether such a library is sufficient can then be judged by comparing the FVUs reached to those of the neural network models. We additionally restrict our dataset to only those points where the optimal three-term q_σ model FVU is below 30%, since the accuracy of our $A_{1,4,5}$ models is largely unimportant if the resulting q_σ model would be inaccurate regardless.

To now perform sparse regression and find the optimal ξ' vector at each complexity, we can still use the same least-squares cost function as before, except with our new nonlinear \hat{y} :

$$C(\xi') = \|\varepsilon(\xi, \zeta)\|^2 = \sum_i \left(y_i - \frac{\sum_j \Theta_{ij} \xi_j}{1 + \sum_k \Theta_{ik} \zeta_k} \right)^2. \quad (6)$$

Here, we have also introduced the notation $\varepsilon(\xi, \zeta)$ for the error vector $y - \hat{y}$.

Since this function depends nonlinearly on ζ , we can no longer find the globally optimal model via *linear* least-squares regression. Instead, we must perform *nonlinear* regression to optimize our least-squares cost function. To this end, we use SciPy’s `optimize.basinhopping()` algorithm¹⁰², which lets us exploit analytically computable gradients to speed up optimization:

$$\begin{cases} \partial_{\xi_j} C = -2 \sum_i \frac{\varepsilon_i \Theta_{ij}}{1 + \sum_{k=1}^{M-1} \Theta_{ik} \zeta_k} \\ \partial_{\zeta_k} C = 2 \sum_i \frac{\varepsilon_i \Theta_{ik} \sum_{j=0}^{M-1} \Theta_{ij} \xi_j}{\left(1 + \sum_{l=1}^{M-1} \Theta_{il} \zeta_l\right)^2}. \end{cases} \quad (7)$$

These expressions—and the cost function itself—are then compiled with Numba’s `@njit` decorator¹⁰³ to speed up execution.

Nonlinear global optimization of this type is generally far slower than linear regression, which uses an analytical expression to calculate the optimum. Additionally, finding the global optimum is not guaranteed in the nonlinear case. This is not solely negative, however, as it helps limit overfitting. Another key difference between the `basinhopping()` algorithm and the linear regression used in regular SINDy is that `basinhopping()` is non-deterministic. This means that one can never know with certainty how close the optimizer is to finding the global optimum (or indeed if it has already been reached). There is simply a tradeoff between spending more computational time with the possibility of improving FVU further, or terminating the optimization at the current optimum.

In our case, we used `basinhopping()` with an L-BFGS-B minimizer and an optimization temperature of $T = \min(4, C_{\text{min}})$, where C_{min} is the minimal cost reached thus far. These settings were sufficient to reach accuracy

equivalent to that of neural networks designed according to Sec. II B. Apart from the fact that the linear regression step at each complexity is replaced by nonlinear optimization using `basinhopping()`, our nonlinear sparse regression (NLSR) algorithm is identical to the version of SINDy we use for discovering models for q_σ .

3. Avoiding poles with Bernstein polynomials

While the overall FVUs reached with these kinds of rational models may be neural network-equivalent, the models in question are nevertheless unsuitable for use in fluid codes, due to the possibility of poles within the parameter domain where the model is to be used. Furthermore, simply expanding the training dataset is not a satisfactory solution to the problem: besides making model training more computationally expensive, this approach alone can never guarantee a complete absence of poles.

Instead, we can ensure there are no poles in the domain of interest by limiting the range of allowed values for the denominator coefficients $\{\zeta_k\}$. A naive way to accomplish this for positive inputs is to simply restrict the coefficients so that $\zeta_k \geq 0$ for all k . However, as discussed in Ref. 104, this restricts the coefficients (and thus the space of models which are reachable) more than is necessary—one can do better by applying this same restriction in the so-called Bernstein polynomial basis.

In the single-variable case, the Bernstein basis for polynomials of degree n are defined as

$$B_{n,k}(x) = \binom{n}{k} x^k (1-x)^{n-k}, \quad (8)$$

with k running from 0 to n . Note that $\sum_k B_{n,k}(x) = 1$ for all n , since the resulting expression is just the binomial expansion of $1^n = (x+1-x)^n$. Each basis function can be thought of as “governing” its own section of the closed interval $[0, 1]$, of width $1/(n+1)$: The maximum of the k^{th} basis function on $[0, 1]$ lies at $x = \frac{k}{n}$ and decreases relatively quickly outside its corresponding interval section, reaching zero at $x = 0$ and $x = 1$ (except the 0th and n^{th} basis functions, which are only zero at the endpoints opposite to those they govern—i.e., $B_{n,0}(1) = B_{n,n}(0) = 0$, while $B_{n,0}(0) = B_{n,n}(1) = 1$).

Just like the monomial basis we are used to, with basis functions $M_{n,k} = x^k$, the Bernstein basis can be used to express any polynomial $p_n(x)$ of degree n . In other words, for any coefficient vector $\zeta_M \in \mathbb{R}^n$, we can find a vector $\zeta_B \in \mathbb{R}^n$ such that

$$p_n(x) = \sum_k \zeta_{M,k} x^k = \sum_k \zeta_{B,k} B_{n,k}(x) \quad (9)$$

performing a change of basis into the Bernstein basis.

Importantly, the Bernstein basis functions are all strictly positive on $(0, 1)$, just like the monomial basis functions, meaning we can ensure a polynomial is strictly

positive on $[0, 1]$ by demanding all of its Bernstein-basis coefficients be nonnegative (with at least one being nonzero). In fact, as shown in Ref. 104, any polynomial with solely nonnegative monomial basis coefficients is guaranteed to also only have nonnegative coefficients in the Bernstein basis. However, the same does not hold in reverse—a polynomial whose Bernstein basis coefficients are all nonnegative may still have negative coefficients in the monomial basis. In other words, the space of polynomials with nonnegative Bernstein coefficients is strictly larger than the space of polynomials with nonnegative monomial coefficients, the former including the latter as a subspace.

In the case of interest to us, where $n = 2$, we get the three basis functions

$$\begin{cases} B_{2,0}(x) = (1-x)^2 \\ B_{2,1}(x) = 2x(1-x) \\ B_{2,2}(x) = x^2. \end{cases} \quad (10)$$

Expanding out these expressions, we can deduce that the change-of-basis matrix Ω_B satisfying $\zeta_M = \Omega_B \zeta_B$ is

$$\Omega_B = \begin{bmatrix} 1 & 0 & 0 \\ -2 & 2 & 0 \\ 1 & -2 & 1 \end{bmatrix}, \text{ with inverse } \Omega_B^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & \frac{1}{2} & 0 \\ 1 & 1 & 1 \end{bmatrix}. \quad (11)$$

The presence of negative numbers in Ω_B but not in Ω_B^{-1} illustrates the fact that if ζ_M is entirely nonnegative, the same holds for ζ_B , while the reverse is not necessarily true, in this special case where $n = 2$.

Generalizing this to four variables is relatively straightforward. Defining $\mathbf{x} = \left[\frac{\bar{n}_b}{\bar{n}}, \frac{\bar{V}_{\text{rel}}}{c}, \frac{\bar{v}_{\text{th},b}}{c}, \frac{\bar{v}_{\text{th},c}}{c} \right]$, we find that just like the multivariate monomial basis functions in this case are simply

$$M_k(\mathbf{x}) = \prod_i x_i^{a_{i,k}}, \quad (12)$$

with $a_{i,k} \in \{0, 1, 2\}$, multivariate Bernstein basis functions can be defined via

$$B_k(\mathbf{x}) = \prod_i B_{2,a_{i,k}}(x_i). \quad (13)$$

Note that with our choice of x_i quantities, ensuring that there are no poles on $[0, 1]$ for each x_i is indeed sufficient for all applications of interest, since $\bar{n}_b \leq \bar{n}$ and $\bar{V}_{\text{rel}}, \bar{v}_{\text{th},b}, \bar{v}_{\text{th},c} < c$, with one caveat: Special care needs to be taken when handling \bar{V}_{rel}/c , since it is a signed quantity. There are several ways one could address this, including e.g. switching to using $x_2 = (\bar{V}_{\text{rel}}/c + 1)/2$, but in our case, since we have knowledge about how A_1 , A_4 and A_5 ought to transform under reflections (which is equivalent to switching the sign of \bar{V}_{rel}/c in this context), we can simply take $x_2 = |\bar{V}_{\text{rel}}/c|$ and handle the sign of \bar{V}_{rel}/c manually. With this modified \mathbf{x} , where x_2 is restricted to be nonnegative (yielding a similarly modified

term library Θ), we can define

$$\begin{cases} [\hat{A}_1]_i = \frac{[\Theta\xi]_i}{[\Theta_B\tilde{\zeta}_B]_i} \\ [\hat{A}_{4,5}]_i = \text{sgn}(\bar{V}_{\text{rel}}) \frac{[\Theta\xi]_i}{[\Theta_B\tilde{\zeta}_B]_i}, \end{cases} \quad (14)$$

with Θ_B being the Bernstein basis term library. If we want to remove a degree of freedom in the denominator to preserve model uniqueness like before, while still requiring the maximally sparse denominator to correspond to the constant polynomial $p(\mathbf{x}) = 1$, we can use the fact that this constant model corresponds to $\tilde{\zeta}_B = \mathbf{1}$, and note that the basis function $B_0(\mathbf{x}) = \prod_i B_{2,0}(x_i)$ controls the value of the denominator at the origin, just like in the monomial basis: Fixing $\tilde{\zeta}_B = \mathbf{1} + [0, \zeta_B]$, so that ζ_B has $M - 1 = 80$ elements, accomplishes this task.

With this definition, setting an increasing number of elements in ζ_B to zero approaches the constant polynomial, as desired. The sufficient demand on ζ_B to eliminate all poles on the domain of interest also changes to demanding $\zeta_{B,k} \geq -1$ for all k . Our final model thus in principle looks like

$$\begin{cases} [\hat{A}_1]_i = \frac{[\Theta\xi]_i}{1 + [\tilde{\Theta}_B\zeta_B]_i} \\ [\hat{A}_{4,5}]_i = \text{sgn}(\bar{V}_{\text{rel}}) \frac{[\Theta\xi]_i}{1 + [\tilde{\Theta}_B\zeta_B]_i}, \end{cases} \quad (15)$$

where $\tilde{\Theta}_B$ denotes the Bernstein-basis term library with the 0th column removed.

However, the $\hat{A}_{4,5}$ models found by sparse regression sometimes have issues due to near-poles at $\bar{V}_{\text{rel}} = 0$. The poor ability of SR to predict behavior in this parameter regime is likely related to the fact that conditions in such cases may be more similar to $\bar{V}_{\text{rel}} > 0$ in certain space-time regions and more similar to $\bar{V}_{\text{rel}} < 0$ in others. To resolve this, we replace the step-like sign function with a smoothly clamped version, tuned to minimally affect performance outside of the problematic small- \bar{V}_{rel} regime. Specifically, we make the replacement

$$\text{sgn}(\bar{V}_{\text{rel}}) \rightarrow \text{sgn}(\bar{V}_{\text{rel}}) \exp\left[-\left(\frac{v_{\text{ref}}}{\bar{V}_{\text{rel}}}\right)^6\right] \quad (16)$$

with $v_{\text{ref}} = 5 \times 10^{-3} c$. It is likely that performance may be further improved by predicting $A_{1,4,5}$ from locally rather than globally averaged data, and/or working in Fourier space to more easily capture wave effects, but computational complexity would in both cases likely increase as a result. Further investigation of such modifications is left outside the scope of this work.

B. Neural networks

To establish a baseline for accuracy and ascertain whether a predictive model is feasible at all,

we also leverage more flexible—yet uninterpretable—machine learning tools in the form of deep neural networks¹⁰⁵ to learn the dependence of $A_{1,4,5}$ upon $\mathbf{x} = [\bar{n}_b/\bar{n}, \bar{V}_{\text{rel}}/c, \bar{v}_{\text{th},b}/c, \bar{v}_{\text{th},c}/c]$. Specifically, we use multi-layer perceptrons (MLPs): fully connected, feedforward neural networks, which have been shown to be universal function approximators¹⁰⁶. These MLPs are organized into a sequence of $L + 1$ layers. The input layer $f^{l=0}$ receives the four plasma quantities contained in \mathbf{x} , performs an affine transformation into a space with dimensionality equal to the number of nodes N per layer and then applies a nonlinear function σ element-wise, i.e.

$$f^0(\mathbf{x}) = \sigma(\mathbf{W}^0\mathbf{x} + \mathbf{b}^0). \quad (17)$$

The $N \times 4$ matrix \mathbf{W}^0 and the $N \times 1$ column vector \mathbf{b}^0 are the *weight* and *bias* matrices of the input layer, respectively, and the nonlinear function σ is often referred to as the *activation function*.

Each subsequent layer l of the MLP is characterized by the application of a new affine transformation, parametrized by a different weight matrix \mathbf{W}^l and bias matrix \mathbf{b}^l . For the hidden layers ($1 \leq l \leq L - 1$) making up the middle of the MLP, \mathbf{W}^l has dimensionality $N \times N$, while \mathbf{b}^l is still $N \times 1$. These unknown parameters are optimized using stochastic gradient descent during training on our dataset. Following the application of this affine transformation, σ is applied again, once for every layer. The output of each hidden layer is thus given recursively by

$$f^l(\mathbf{x}) = \sigma(\mathbf{W}^l f^{l-1}(\mathbf{x}) + \mathbf{b}^l). \quad (18)$$

The output layer ($l = L$) needs to output a scalar, thus \mathbf{W}^L is $1 \times N$ and \mathbf{b}^L is 1×1 . In order to allow the final output to have absolute values greater than one, the final layer omits the application of σ , and its output is instead simply given by

$$f^L(\mathbf{x}) = \mathbf{W}^L f^{L-1}(\mathbf{x}) + \mathbf{b}^L, \quad (19)$$

The output of this recursively defined function constitutes the output of the MLP.

Being universal function approximators, MLPs cover the space of all functions well. This is not necessarily the case for more traditional function fitting approaches, such as the ones used in sparse regression. However, a significant drawback to using MLPs—like with other similar neural network-based approaches—is that the resulting models are highly opaque. This makes it difficult to gain physical intuition from them, and hard to judge how likely they are to extrapolate well into regimes beyond the training dataset. To mitigate the latter issue, one can utilize cross-validation, just like we do in our implementations of sparse regression discussed above. Even then, however, well-behaved extrapolation can be difficult to ensure.

We trained one MLP for each combination of electron species (beam, core and combined) and SR coefficient




FIG. 1: Log-scale plot of training data FVU vs iteration for the neural network model estimating the combined-species A_1 coefficient (blue), along with a 51-iteration moving average (black). The 51-iteration average FVU attained at the end of training, 21%, is highlighted with a red horizontal line.

of interest (A_1 , A_4 and A_5) to get an approximate lower bound on what FVU scores are possible to reach with our model in terms of $\{n_b/n, V_{\text{rel}}/c, v_{\text{th},b}/c, v_{\text{th},c}/c\}$. Each MLP used three hidden layers of $N = 25$ nodes, with hyperbolic tangent activation functions: $[\sigma(\mathbf{a})]_i = \tanh(a_i)$. The total number of free parameters in each MLP is thus¹⁰⁷ 2101. These were then split into 10 batches and trained for 4000 epochs using the AdamW algorithm with mean-squared cost functions and tuned learning rates on the order of $\sim 10^{-3}$ (see Fig. 1 and Tab. I), as implemented in the PyTorch package^{108,109}.

The input space employed here, consisting of only four box-averaged fluid quantities, limits the achievable accuracy of the closure model. This limitation is reflected in the residual error of the neural network regression, as illustrated in Fig. 1. In the present work, this modeling error dominates over numerical inaccuracies arising from finite simulation resolution.

C. Simulations and dataset

We use 1D3V physical mass ratio electron-proton PIC simulations in the OSIRIS code to generate the data required for SR-based closure discovery. Consistently, we initialize the plasma as consisting of three populations, each in internal thermal equilibrium: a stationary zero-temperature ion (proton) population, a beam electron population and a core electron population. All simulations are performed in the center-of-momentum (CoM) frame, where the initial flow velocity V of the combined

electron fluid, as well as the net electric current, is zero. Since ion dynamics largely only occur over time scales significantly longer than the ones considered, V remains negligible throughout all simulations, and ion fluid quantities can without issue be excluded from our analysis.

Our dataset in total consists of the output from 76 simulations with different initial conditions, varying the four degrees of freedom available to us, namely the beam density n_b , the relative flow velocity $V_{\text{rel}} = V_c - V_b$ between the two electron populations and the thermal speeds $v_{\text{th},b,c}$ for the two electron species. As described in, e.g., Ref. 110, the electron/electron beam instability occurs only at sufficiently high n_b/n , $|V_{\text{rel}}|/v_{\text{th},c}$ and $v_{\text{th},c}/v_{\text{th},b}$, with the dependence on beam-core temperature differential being relatively weak. For example, using a very rough approximation, instability requires $|V_{\text{rel}}|/v_{\text{th},c} \gtrsim 2 + \lg(v_{\text{th},b}/v_{\text{th},c}) - \lg(n_b/n)$ for $10^{-1} \lesssim v_{\text{th},b}/v_{\text{th},c} \lesssim 10$ and with n_b/n in the ranges of interest to us, as can be seen in Fig. 3.9 in Ref. 110.

Of the 76 simulations, 47 constitute parameter sweeps in each of the four initialization parameters, over the range where instability occurs and relativistic effects are weak, centered on $n_b = 0.1 \bar{n}$, $V_{\text{rel}} = 0.02c$ and $v_{\text{th},b} = v_{\text{th},c} = \sqrt{10^{-5} c^2} \approx 3.16 \times 10^{-3} c$. Specifically,

$$\begin{cases} n_b \in [0.01, 0.5] \bar{n}, \\ V_{\text{rel}} \in [0.012, 0.1] c, \\ v_{\text{th},b}, v_{\text{th},c} \in [5 \times 10^{-4}, 8 \times 10^{-3}] c \end{cases} \quad (20)$$

are considered. The remaining 29 simulations use random parameters satisfying the same constraints, i.e. exhibiting instability and being at most weakly relativistic. Parameters were selected from uniform distributions over

$$\begin{cases} n_b \leq 0.5 \bar{n}, \\ V_{\text{rel}} \leq 0.08 c, \\ v_{\text{th},b}, v_{\text{th},c} \leq 8 \times 10^{-3} c. \end{cases} \quad (21)$$

Additionally, we limit ourselves to cases where

$$\frac{n_b}{\bar{n}} \frac{|V_{\text{rel}}|^3}{v_{\text{th},b}^3} \gtrsim 1, \quad (22)$$

since cases where this expression is less than unity are either stable or correspond to Langmuir beam instability rather than electron/electron beam instability¹¹⁰.

For all simulations we use a box size $L = 2.048 \delta_e$ with periodic boundary conditions and spatial resolution $\Delta x = 10^{-3} \delta_e$, where $\delta_e = c/\omega_{\text{pe}}$ is the electron inertial length, $\omega_{\text{pe}} = \sqrt{\bar{n} e^2 / m_e \epsilon_0}$ being the electron plasma frequency (in this expression, e is the elementary charge, m_e is the electron mass and ϵ_0 is the permittivity of the vacuum). To ensure numerical stability, the temporal resolution is selected to be slightly higher than the spatial one in natural units, corresponding to a time step of $9.5 \times 10^{-4} \omega_{\text{pe}}^{-1}$. To limit data output, we only record the state of the simulation every 100 time steps, meaning our sparse regression dataset has a temporal resolution of

$\Delta t = 0.095 \omega_{pe}^{-1}$. Most simulations have a total duration of $100 \omega_{pe}^{-1}$, while simulations with weaker growth rates instead use a duration of $200 \omega_{pe}^{-1}$, to ensure both the growth phase and saturated phase of the instability are captured well. We consistently use 10^4 electrons per cell of each species (i.e., 2×10^4 electrons per cell in total), and 200 ions per cell.

III. RESULTS

The previously discovered six-term single-species heat flux closure can be generalized in a straightforward manner to handle separate-species modeling, in accordance with what one may expect from linear theory, as discussed in Appendix A and Sec. II A 1. Furthermore, the three most important unknown coefficients $A_{1,4,5}$ can be predicted with reasonable accuracy from the four box-averaged fluid quantities $\{n_b, V_{rel}, v_{th,b}, v_{th,c}\}$, normalized to the total number density and the speed of light, respectively. Specifically, neural network models explain 77–97% of the variation in the coefficients, while NLSR-discovered rational models have comparable but generally slightly lower accuracies, ranging from 72% to 94% in the monomial-basis case and from 68% to 96% in the Bernstein-basis case, depending on the species and coefficient considered. The resulting heat flux models regularly account for 80–90% of the variation in q_σ , translating into a typical accuracy of 85–95% when predicting $\partial_t p_\sigma$. Notably, the error when predicting $\partial_t p_\sigma$ is consistently more than halved (typically smaller by a factor of ~ 4) compared to using a naive $\hat{q}_\sigma = 0$ model.

A. Multi-species closures

We first note some differences between the time evolution of separate- and combined-species fluid quantities. As shown in Appendix B, n and nV are simple sums of the single-species quantities. Since the combined quantities follow linear theory during the growth phase, the

same can be expected for the separate-species quantities. To first order in perturbation theory, the same holds for V , being the ratio of nV and n .

This is indeed what we observe. To quantify the variation in a quantity X over the simulation box at a certain point in time, let us use the spatial variance

$$\text{var}_x X = \frac{\Delta x}{L} \sum_i (X_i - \bar{X})^2. \quad (23)$$

For a sinusoidally varying quantity described by $X = \bar{X} + \delta_X \sin[kx + \varphi_X]$ at some point in time, this reduces to $\delta_X^2/2$ in the limit of small Δx . With $X = E$, for example, this corresponds to the average E -field energy density, since $\bar{E} = 0$. In Fig. 2ab, we plot the spatial variance of n_σ and V_σ for each species (beam, core and combined), together with that of E (in arbitrary units) to compare the perturbation growth for each quantity over the course of the simulations. As we can see, all seven quantities increase in tandem with the same growth rate $\gamma_E = \frac{1}{2} \partial_t \ln \text{var}_x E$ once above their respective noise floors, in accordance with single-mode linear theory.

The same is not true for p_σ and q_σ . Even though the combined-species p and q perturbations continue to grow at a rate $\gamma(t) \approx \gamma_E(t)$, the separate-species $p_{b,c}$ and $q_{b,c}$ do not increase at the same rate. Specifically, in most of the simulations considered, including the one depicted in Fig. 2, the perturbation in q_b grows at a weakly accelerating rate during the growth phase, while the q_c perturbation remains negligible until the end of linear growth, after which its growth rapidly accelerates, quickly reaching similar orders of magnitude as the beam perturbation. In simulations with very cold beams, however, the situation is reversed, the perturbation in q_c instead growing with an accelerating rate and the perturbation in q_b staying negligible before growing very quickly during saturation.

This nonlinear behavior is due to the fact that the combined-species p and q are not simple sums of the single-species p_σ and q_σ , respectively. In fact, the combined-species p and especially q are dominated by the extra terms in Eqs. (B9) and (B13) during linear growth:

$$\begin{cases} p \sim \sum_\sigma \bar{n}_\sigma \bar{v}_{th,\sigma}^2 + \frac{n_b n_c}{n} (V_b - V_c)^2 \\ q \sim \sum_\sigma \bar{q}_\sigma + \frac{n_b n_c}{n} (V_b - V_c) \left[3(T_b - T_c) - \frac{n_b - n_c}{n} (V_b - V_c)^2 \right] \approx \\ \approx \frac{n_b n_c}{n} (V_b - V_c) \left[3(T_b - T_c) - \frac{n_b - n_c}{n} (V_b - V_c)^2 \right]. \end{cases} \quad (24)$$

For q in particular, having non-negligible

$$\frac{\sum_\sigma \text{var}_x q_\sigma}{\text{var}_x q} \quad (25)$$

correlates very strongly with the presence of nonlinear effects. We can see this clearly if we plot this ratio and the corresponding quantity for the pressures, as in Fig. 3. Notably, the beam and core heat fluxes at the end of

linear growth seemingly mainly work to either decrease or increase the amplitude of the peaks in the extra q term.




FIG. 2: The time evolution of the spatial variance in (a) n_σ , (b) V_σ , (c) p_σ and (d) q_σ , with beam quantities in dashed blue, core quantities in dotted orange and combined-species quantities in dash-dotted green, along with $\text{var}_x E$ in solid black, rescaled to fit the plotted ranges in each panel. In the depicted simulation, an initial condition of $\{n_b, V_{\text{rel}}, v_{\text{th},b}, v_{\text{th},c}\} = \{0.17 \bar{n}, 3.4 \times 10^{-2} c, 5.0 \times 10^{-3} c, 6.4 \times 10^{-3} c\}$ was used. The linear growth phase, defined as the time range of linear E -field perturbation growth, is highlighted in red.

The nonlinear behavior of q_σ in the growth phase leads to some interesting complications when it comes to finding a heat flux closure. In particular, since the subspecies-internal heat fluxes are for the most part negligible compared to the extra term in q during linear growth, taking $q_\sigma = 0$ is a viable closure for this part of the simulation. On the other hand, in the saturated phase, q_c in particular is important to model well, and q_b is also non-negligible in some parts of the simulation. The challenge is thus to find a closure able to stitch the behavior in the two phases together.

1. The multi-species closures found by SR

As outlined in Appendix A, building q_σ closures which behave analogously at first order in perturbation theory to the original single-species closure requires the A_1 term to be generalized according to

$$A_1 n v_{\text{th}}^2 V \rightarrow A_1 n_\sigma v_{\text{th},\sigma}^2 (V_\sigma - \bar{V}_\sigma), \quad (26)$$

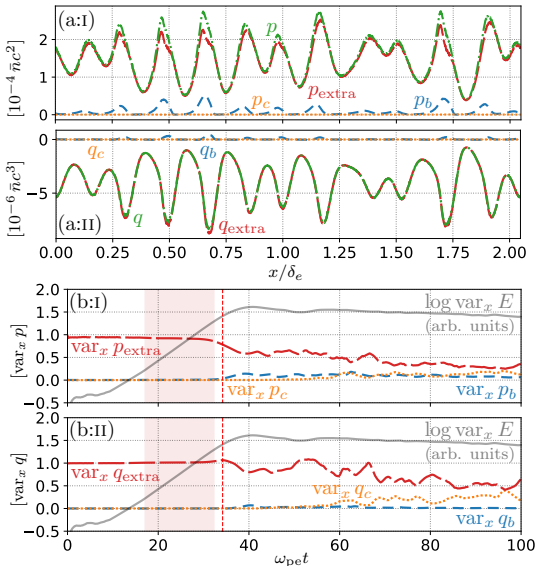


FIG. 3: Snapshots of (a:I) p_σ and (a:II) q_σ at the end of linear growth, with beam in dashed blue, core in dotted orange, combined-species in dash-dotted green and the “extra term” in long dashed red. Additionally, we show the time evolution of the spatial variance in p_b and q_b (dashed blue), p_c and q_c (dotted orange) and the extra p and q terms (long dashed red), normalized to the combined-species variances $\text{var}_x p$ and $\text{var}_x q$. Pressure quantities are shown in (b:I) and heat flux quantities in (b:II). The vertical dashed red lines in these plots mark the time of the snapshot panels. For reference we also show the time evolution of the logarithmic E -field variance in gray in arbitrary units, and highlight the linear growth phase in red. As before, the depicted case uses the initial condition $\{n_b, V_{\text{rel}}, v_{\text{th},b}, v_{\text{th},c}\} = \{0.17 \bar{n}, 3.4 \times 10^{-2} c, 5.0 \times 10^{-3} c, 6.4 \times 10^{-3} c\}$.

with all other terms in the combined-species six-term closure simply being amended by adding a σ subscript to every fluid quantity, so that the heat flux of each subspecies is modeled according to $q_\sigma = q_{\text{even},\sigma} + q_{\text{odd},\sigma}$ with

$$\begin{cases} q_{\text{even},\sigma} = A_1 n_\sigma v_{\text{th},\sigma}^2 (V_\sigma - \bar{V}_\sigma) + A_2 v_{\text{th},\sigma}^3 \partial_x n_\sigma + \\ \quad + A_3 n_\sigma v_{\text{th},\sigma}^2 \partial_x v_{\text{th},\sigma} \\ q_{\text{odd},\sigma} = A_4 + A_5 n_\sigma v_{\text{th},\sigma}^3 + A_6 n_\sigma v_{\text{th},\sigma}^2 \partial_x V_\sigma. \end{cases} \quad (27)$$

And indeed, this is what is found by SR using a term library analogous to the one used in Ref. 89 for the combined-species case, consisting of dimensionally consistent terms constructed from products of n_σ , $V_\sigma - \bar{V}_\sigma$ and $v_{\text{th},\sigma}$, as well as first-order spatial derivatives of these quantities. Since $\partial_x (V_\sigma - \bar{V}_\sigma) = \partial_x V_\sigma$, such a library still allows us to identify the A_6 term without issue. For com-

pletteness the single-species 6-term closure is reproduced in Appendix C, where the various terms are also physically interpreted.

A naive term library using V_σ instead of $V_\sigma - \bar{V}_\sigma$ in principle still allows SR to find this closure, provided one runs the sparse regression algorithm on narrow enough time slices of simulation data that \bar{V}_σ is approximately constant over each time slice. In practice, however, the fact that \bar{V}_σ does vary over each time slice, combined with the difficulty of identifying an extra term which is correlated with a term already included in the closure means that SR convergence often fails in such cases. As one might expect, SR usually does find the closure in Eq. (27) with such a library in situations where V_σ is sufficiently small—sans the \bar{V}_σ correction.

Similarly, expanding the term library to also include terms with a mix of quantities from the two species, like $n_b v_{th,b} v_{th,c}^2$, generally also leads to convergence issues due to the vastly increased library size and correlations between the perturbations in the beam and core fluid quantities. When SR does converge for such libraries, no additional closure terms to those appearing in Eq. (27) are found consistently.

There are some more significant differences from the combined-species case, however. In a single-species treatment, the terms corresponding to A_2 , A_3 and A_6 are strongly correlated with the instantaneous growth rate of the instability γ_E , as illustrated in Fig. 4 of Ref. 89. When modeling the subspecies separately, however, these coefficients are significantly less transparently related to γ_E . There are two reasons for this. The first is that the two subspecies interact with each other through the electric field. As discussed in Appendix A, this can be understood in terms of each subspecies having its own, complex, plasma frequency $\omega_{p\sigma}$, which depends on the ratio between the (complex) relative density perturbation $r_\sigma = \tilde{n}_\sigma/\bar{n}_\sigma$ for the species in question and the analogous quantity $r = \tilde{n}/\bar{n}$ for the combined species, in accordance with Eq. (A1): $\omega_{p\sigma} = \omega_{pe} \sqrt{r}/r_\sigma$. The possible presence of an imaginary part to this quantity makes the constraints derived from linear theory, Eq. (A6), significantly more complicated than in the combined-species case, and among other things implies that $A_2 = A_3 = A_6 = 0$ is no longer a solution at $\gamma = 0$ for arbitrary A_5 . Instead, setting all “growth-related” terms equal to zero at $\gamma = 0$ is only possible for $A_5 = 2\omega_r^2/3k\bar{v}_{th,\sigma}$. However, as we shall see, this is actually quite a good estimate of A_5 in the saturated phase, where $\gamma \sim 0$, so this is not the main reason why the correlation between $A_{2,3,6}$ and γ_E breaks down in the separate-species case.

The main reason is instead that $A_{2,3,6}$ for each species correlates with the growth rate γ_q of the heat flux perturbation, which is not equal to γ_E in the separate-species case, as discussed above. This correlation, though less transparent than in the combined-species case, can be seen clearly in Fig. 4. The fact that the correlation is with γ_q rather than γ_E is problematic from a closure construction perspective, however, since γ_q for the sub-

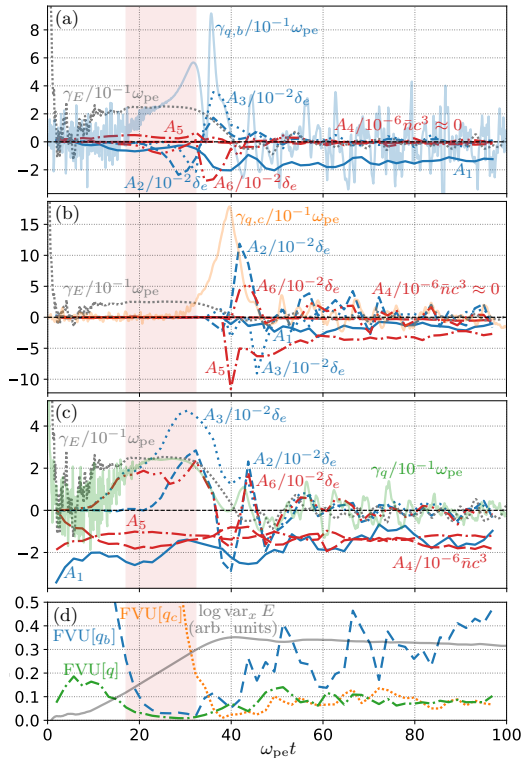


FIG. 4: Time evolution of 6-term closure coefficients compared to growth rates γ_E and γ_q for the (a) beam species, (b) core species and (c) combined species, together with (d) the FVU over time for the discovered models of q_b (dashed blue), q_c (dotted orange) and q (dash-dotted green). In the former three panels, q_{even} (q_{odd}) coefficients are shown in blue (red), while γ_E is shown in tightly dotted gray and γ_q is shown in faded blue (orange, green) for the beam (core, combined) cases, respectively. In the latter panel, $\log \text{var}_x E$ in arbitrary units is additionally shown in gray for reference. In every subplot, the linear growth phase is highlighted in red.

species cannot be predicted from lower-order moments, but rather depends directly on the unknown quantity q_σ . In principle, it might be possible to predict γ_q to some extent from γ_E together with γ_p , and then in turn use this to predict the growth-related coefficients. In this work, however, we have elected to sidestep this issue by restricting ourselves to the three-term model containing only the three most important terms, i.e.

$$q_\sigma = A_1 n_\sigma v_{th,\sigma}^2 (V_\sigma - \bar{V}_\sigma) + A_4 + A_5 n_\sigma v_{th,\sigma}^3. \quad (28)$$

This decision is also motivated by the fact that includ-

ing any instantaneous growth rate $\gamma(t)$ in a closure necessitates very careful consideration of the exact implementation if one wants to avoid introducing unphysical instabilities.

We note that the FVU for both the six-term and three-term q_σ models are generally somewhat higher than the FVU for the corresponding combined-species q model, as can be seen in Fig. 4d in the six-term case. One needs to be careful not to draw the wrong conclusion from this, however. While it is true that q can most often be modeled more accurately than q_σ for either subspecies, modeling both species separately with imperfect q_σ models still captures the physics better than modeling the electron population as a single fluid species. This is because q in large part is determined by the extra term in Eq. (B13), which is “automatically” modeled exactly when using two electron fluids.

Restricting to the three-term model, we can solve the linear theory constraints for A_1 and A_5 explicitly. The resulting expressions, listed in Eq. (A), are not very practical to work with in our case, however. The main reason for this is that some quantities which need to be provided, like ω'_σ , are poorly defined for perturbations which deviate significantly from pure sinusoidality. A proper treatment would thus require working in Fourier space. In the low- γ limit, however, the expressions simplify considerably, becoming $A_1 = -8/3$ and $A_5 = 2\omega'_r/3k\bar{v}_{th,\sigma}$. This limit is of particular interest to us, since it describes the saturated phase, which is the regime where q_b and q_c are consistently non-negligible compared to q . As we can see in Fig. 5, these simplified expressions indeed agree fairly well with the values for A_1 and A_5 found by SR post-saturation, despite the presence of nonlinear effects. The A_5 coefficient in particular is predicted quite accurately, despite the fact that the perturbations in the simulation are multi-modal in nature—the faded lines shown in the figure use characteristic values for ω_r and k , approximately corresponding to the peaks in the temporal and spatial Fourier spectra. Specifically, the characteristic (angular) wavenumber is computed at each timestep by taking a weighted average over the spatial fast Fourier transform of the E -field, using the Fourier magnitude squared as the weight at each wavenumber k . For the characteristic (angular) frequency, a similar weighted averaging procedure is utilized, but with temporal continuous wavelet transform (CWT) magnitudes used instead of Fourier magnitudes to retain temporal resolution. This latter computation utilizes the Scalegram module for CWT data analysis¹¹¹, based on the PyWavelets library¹¹². The predicted value of A_1 generally lies further from the SR value, but is correct as an order-of-magnitude estimate. Since the linear prediction in this case is simply a constant value, this is not entirely unexpected. Furthermore, in the growth phase, where γ is nonzero, both coefficients—unsurprisingly—deviate significantly from their predicted values at $\gamma = 0$ in the separate-species cases. In the case depicted in Fig. 5, $\gamma_E \approx 0.25\omega_{pe}$ during linear growth.

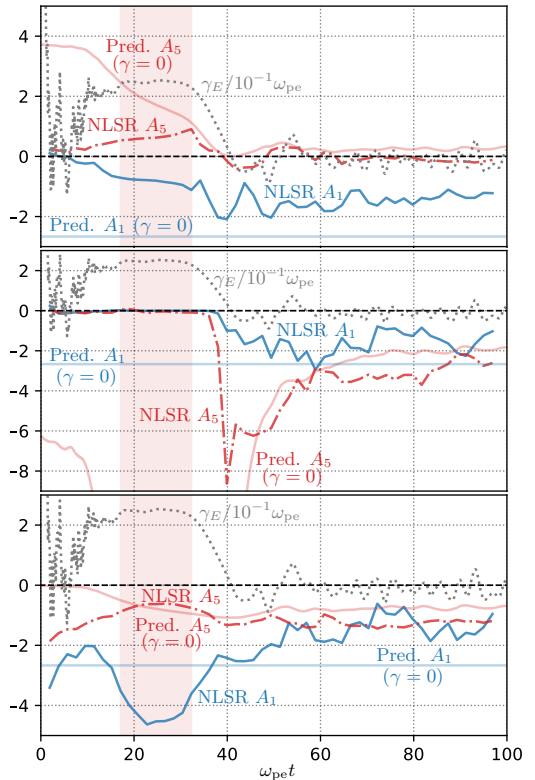


FIG. 5: The two 3-term closure coefficients A_1 (solid blue) and A_5 (dash-dotted red) compared to predictions from linear theory at $\gamma = 0$ (in faded blue and red, respectively) shown for the beam (top), core (middle) and combined population (bottom) in the simulation with initial condition $\{n_b, V_{rel}, v_{th,b}, v_{th,c}\} = \{0.17 \bar{n}, 3.4 \times 10^{-2} c, 5.0 \times 10^{-3} c, 6.4 \times 10^{-3} c\}$. Furthermore, the growth rate of the electric field perturbation is plotted in dotted gray and the linear growth phase is highlighted in red.

Note also that while these predicted values of A_1 and A_5 in theory entail a complete, usable closure (A_4 does not affect $\nabla \cdot \mathbf{q}$), they depend on wave parameters and on the validity of linear theory, meaning the resulting closure would be nonlocal and likely inaccurate for modeling nonlinear phenomena. Thus, we will instead utilize neural networks and nonlinear sparse regression, as described in Secs. II B and II A 2, to estimate A_1 , A_4 and A_5 from the four box-averaged fluid quantities \bar{n}_b/\bar{n} , \bar{V}_{rel}/c , $\bar{v}_{th,b}/c$ and $\bar{v}_{th,c}/c$.

B. Closure coefficient modeling

Using neural networks of the type described in Sec. II B, the 3-term closure coefficients can be predicted at decent accuracy—varying from a total FVU of $\sim 3\%$ for the combined-species A_5 coefficient to a total FVU of $\sim 23\%$ for the combined-species A_1 coefficient. The full range of FVU values on the training, test and total datasets from these neural network models, as well as the learning rate η used in each case can be seen

in Tab. I, and plots illustrating the performance for the core-species A_5 coefficient is shown in Fig. 6. As discussed in Sec. II B, these values can be regarded as estimates of the lower bound on the FVU reachable when attempting to model $A_{1,4,5}$ as a function of the four box-averaged quantities we have selected. The majority of the remaining error is likely to be “irreducible” without introducing a dependence on wave parameters or similar into the model, which would entail an increase in computational complexity.

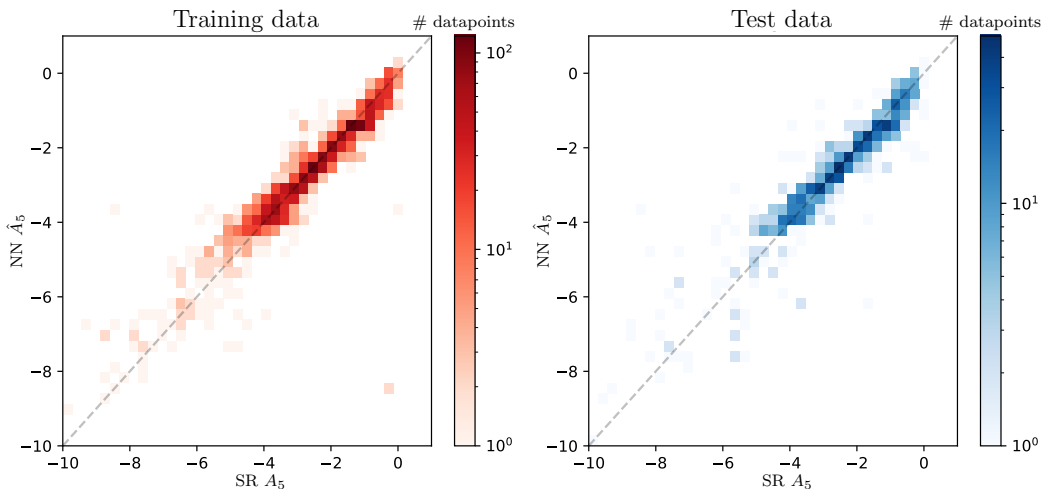


FIG. 6: Neural network performance when estimating $A_{5,c}$ from \bar{n}_b/\bar{n} , V_{rel}/c , $\bar{v}_{th,b}/c$ and $\bar{v}_{th,c}/c$: a binned scatterplot of predicted vs actual coefficient values (left: training data, right: test data). In this case, a learning rate of $\eta = 3.5 \times 10^{-3}$ was used, with FVU scores of 7% (training), 8% (test) and 7% (total).

Using NLSR with a monomial-basis rational ansatz as outlined in Sec. IIA 2, we find models for all species and closure coefficients with largely neural network-equivalent accuracy. In fact, the accuracy of the NLSR models in some cases, e.g. for the beam-species A_4 coefficient, even slightly surpasses that of the neural network models, as can be seen in Tab. II. As the dependence of $A_{1,4,5}$ upon \bar{n}_b/\bar{n} , V_{rel}/c , $\bar{v}_{th,b}/c$ and $\bar{v}_{th,c}/c$ is relatively complex, the minimum model complexity m required for

neural network-equivalent performance is also high, ranging from 25 to 73 terms for the monomial-basis models and from 31 to 82 terms for the Bernstein-basis models.

For example, the combined-species A_5 coefficient, being one of the more easily expressible, still requires 25 terms. Writing the model out explicitly—which is only feasible due to the relatively limited number of terms involved, we find that $\hat{A}_5 = N(\mathbf{x})/D(\mathbf{x})$ for $\mathbf{x} = [\bar{n}_b/\bar{n}, \bar{v}_{th,b}/c, \bar{v}_{th,c}/c, \bar{V}_{rel}/c]$ with

$$\left\{ \begin{array}{l} N(\mathbf{x}) = \xi_{34}\bar{n}_b\bar{v}_{th,b}^2\bar{V}_{rel} + \xi_{36}\bar{v}_{th,c}\bar{V}_{rel} + \xi_{37}\bar{n}_b\bar{v}_{th,c}\bar{V}_{rel} + \xi_{38}\bar{n}_b^2\bar{v}_{th,c}\bar{V}_{rel} + \xi_{54}\bar{V}_{rel}^2 + \xi_{73}\bar{n}_b\bar{v}_{th,c}^2\bar{V}_{rel}^2 + \xi_{74}\bar{n}_b^2\bar{v}_{th,c}^2\bar{V}_{rel}^2 \\ D(\mathbf{x}) = 1 + \zeta_8\bar{n}_b^2\bar{v}_{th,b}^2 + \zeta_9\bar{v}_{th,c} + \zeta_{11}\bar{n}_b^2\bar{v}_{th,c} + \zeta_{15}\bar{v}_{th,b}^2\bar{v}_{th,c} + \zeta_{18}\bar{v}_{th,c}^2 + \zeta_{19}\bar{n}_b\bar{v}_{th,c}^2 + \zeta_{20}\bar{n}_b^2\bar{v}_{th,c}^2 + \\ + \zeta_{37}\bar{n}_b\bar{v}_{th,c}\bar{V}_{rel} + \zeta_{38}\bar{n}_b^2\bar{v}_{th,c}\bar{V}_{rel} + \zeta_{45}\bar{v}_{th,c}^2\bar{V}_{rel} + \zeta_{56}\bar{n}_b^2\bar{V}_{rel}^2 + \zeta_{59}\bar{n}_b^2\bar{v}_{th,b}\bar{V}_{rel}^2 + \zeta_{62}\bar{n}_b^2\bar{v}_{th,b}\bar{V}_{rel}^2 + \\ + \zeta_{63}\bar{v}_{th,c}\bar{V}_{rel}^2 + \zeta_{64}\bar{n}_b\bar{v}_{th,c}\bar{V}_{rel}^2 + \zeta_{65}\bar{n}_b^2\bar{v}_{th,c}\bar{V}_{rel}^2 + \zeta_{72}\bar{v}_{th,c}^2\bar{V}_{rel}^2 + \zeta_{73}\bar{n}_b\bar{v}_{th,c}^2\bar{V}_{rel}^2, \end{array} \right. \quad (29)$$

TABLE I: FVU performance and learning rates η for the neural network models used to estimate $A_{1,4,5}$ from the four box-averaged quantities \bar{n}_b/\bar{n} , \bar{V}_{rel}/c , $\bar{v}_{\text{th},b}/c$ and $\bar{v}_{\text{th},c}/c$.

	Training FVU	Test FVU	Total FVU	η [10^{-3}]
$A_{1,b}$	0.10	0.25	0.14	3.5
$A_{1,c}$	0.07	0.16	0.10	3.5
A_1	0.21	0.26	0.23	3.5
$A_{4,b}$	0.16	0.30	0.21	1.4
$A_{4,c}$	0.04	0.08	0.05	1.4
A_4	0.07	0.11	0.08	3.5
$A_{5,b}$	0.07	0.22	0.11	1.0
$A_{5,c}$	0.07	0.08	0.07	3.5
A_5	0.03	0.04	0.03	3.5

where the normalization factors \bar{n} and c have been omitted for readability. Note that ξ is 0-indexed, while ζ is 1-indexed, in agreement with Sec. IIA 2.

As it would be cumbersome and offer little immediate insight, we do not explicitly write out the remaining NLSR coefficient models, nor the nonzero components of ξ and ζ . In principle, however, such expressions could be written out for each model, and they still provide more interpretable information than the optimized parameters of a deep neural network. In addition, NLSR models are easier to differentiate analytically, as they lack the recursive structure of neural networks. Although they are not as interpretable as traditional linear SR models, they thus remain significantly more interpretable than neural networks, while retaining comparable expressive power. It should also be noted that while the NLSR models do require a large number of coefficients to reach neural network-equivalent accuracy, the number of free parameters in the NLSR models is still lower than that of our MLP models by more than an order of magnitude.

Because of the nevertheless relatively high model complexity required for neural network-equivalent performance, and due to the stochastic nature of the algorithm, consistent model identification is very rare with NLSR. Instead, there are often many different viable models at each model complexity. Thus, it is more useful to show the FVU of all models found in the ten cross-validation folds at each complexity in a scatterplot, as in Fig. 7, rather than only showing the average FVU and the range of FVUs reached as is usually done for linear SR. In the figure, we specifically show the A_5 models found for the beam electron population, but the behavior of NLSR for other coefficients and species is similar—the main difference being what the optimal FVU is, and what complexity is required to reach it. We note that for some coefficients and species, NLSR has difficulty converging at high complexity, leading to a U-shaped apparent Pareto front. In all such cases, however, FVU scores on the same order as those of the corresponding neural network models were still reached, at a lower complexity.

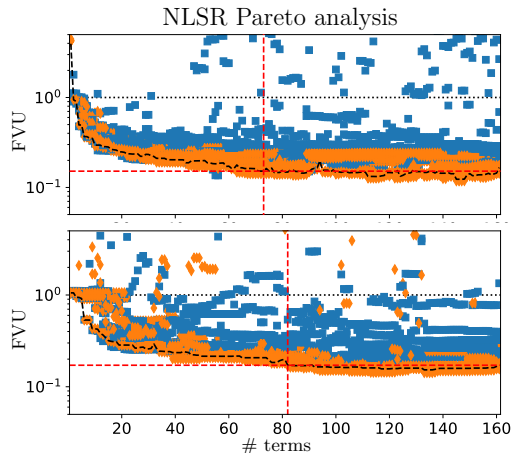


FIG. 7: Scatterplot of FVU for the rational monomial-basis (top) and Bernstein-basis (bottom) beam-species A_5 models found by NLSR at each complexity in the ten cross-validation folds, for training data (orange diamonds) and test data (blue squares). The total dataset FVU for the optimal models at each complexity, approximately tracing the Pareto front, is additionally shown in dashed black. Furthermore, the complexity and total FVU of the Pareto-optimal models included in Tab. II, i.e. $m = 73$, FVU 15% (monomial) and $m = 82$, FVU 17% (Bernstein), are indicated by the red dashed lines.

As discussed in Sec. IIA 3, however, these models sometimes exhibit problematic behavior when applied outside of the training dataset due to poles within the parameter regime of interest. If one instead expresses the denominator polynomials in the Bernstein basis and restricts to nonnegative coefficients as outlined in the section in question, this issue is eliminated at the cost of a slight decrease in accuracy for the training dataset. The FVU scores reached by NLSR—both with and without these restrictions—are shown in Tab. II for selected models, together with the corresponding model complexities. The models shown in the table were selected to have FVU scores comparable to the minimum FVU reached, using a minimal number of terms, to maximize ease of use and limit overfitting.

As we can see in the table, the more restricted Bernstein-basis models fulfilling these criteria tend to have comparable but slightly higher FVU than the monomial-basis ones. For the complexity, on the other hand, behavior varies depending on coefficient: A_1 is seemingly more easily expressible using the Bernstein basis, while A_4 and A_5 are sparser in the monomial basis. Note that while the latter two are captured fully by the rational NLSR models, the NLSR A_1 models have somewhat higher FVUs than the neural network models, sug-

TABLE II: FVU performance and model complexity m (number of non-zero terms in ξ') for selected monomial- and Bernstein-basis rational models of $A_{1,4,5}$ in terms of the four box-averaged quantities \bar{n}_b/\bar{n} , \bar{V}_{rel}/c , $\bar{v}_{\text{th},b}/c$ and $\bar{v}_{\text{th},c}/c$ found by NLSR.

	Training FVU		Test FVU		Total FVU		m	
	Mon.	Ber.	Mon.	Ber.	Mon.	Ber.	Mon.	Ber.
$A_{1,b}$	0.27	0.28	0.27	0.29	0.27	0.28	60	56
$A_{1,c}$	0.19	0.29	0.21	0.35	0.20	0.30	67	49
A_1	0.27	0.31	0.30	0.38	0.28	0.32	42	31
$A_{4,b}$	0.15	0.17	0.29	0.30	0.16	0.18	25	62
$A_{4,c}$	0.06	0.07	0.11	0.14	0.06	0.08	25	52
A_4	0.08	0.08	0.09	0.10	0.08	0.08	25	48
$A_{5,b}$	0.15	0.17	0.15	0.20	0.15	0.17	73	82
$A_{5,c}$	0.08	0.12	0.16	0.14	0.09	0.12	42	78
A_5	0.06	0.05	0.07	0.05	0.06	0.05	25	76

gesting performance may be improved further by altering the term library or functional form.

While modeling these coefficients with largely neural network-equivalent accuracy is promising, fully assessing the viability of the $A_{1,4,5}$ models necessitates examining how well a closure implementing them, via

$$\hat{q}_\sigma = \hat{A}_{1,\sigma} n_\sigma v_{\text{th},\sigma}^2 (V_\sigma - \bar{V}_\sigma) + \hat{A}_{4,\sigma} + \hat{A}_{5,\sigma} n_\sigma v_{\text{th},\sigma}^3, \quad (30)$$

predicts q_σ and $\partial_x q_\sigma$. As can be seen in Fig. 8, these models of q_σ are generally quite accurate, as expected based on the relatively low individual-coefficient FVU scores. For the simulation depicted in the figure, the FVU of this resulting model \hat{q}_σ over all spacetime points i in the simulation, meaning

$$\text{FVU}_{\text{sim}}[\hat{q}_\sigma] = \frac{\sum_i (q_{\sigma,i} - \hat{q}_{\sigma,i})^2}{\sum_i (q_{\sigma,i} - \bar{q}_\sigma)^2}, \quad (31)$$

is 16% for the beam species, 11% for the core species and 14% for the combined species. At early timesteps, we can see that the plots of PIC q_σ and \hat{q}_σ rational models differ slightly in shade, signifying that the \hat{A}_4 is somewhat inaccurate pre-instability. This is unsurprising, since this region was excluded from the optimization dataset, and the A_4 term does not naturally scale with the perturbations in other fluid quantities, unlike the terms associated with A_1 and A_5 . Regardless, an incorrect value of the A_4 constant term has no influence on the viability of the closure, since only $\partial_x \hat{q}_\sigma$ is actually inserted into the pressure equation when the closure is used. Measuring the accuracy with which this quantity is predicted is somewhat less straightforward than for undifferentiated q_σ , however, since numerical differentiation amplifies particle noise.

Calculating $\text{FVU}_{\text{sim}}[\partial_x \hat{q}_\sigma]$ naively using an $\mathcal{O}(\Delta x^2)$ accurate symmetric scheme with spatial step $\delta x = 2\Delta x$, via the same method as for \hat{q}_σ , we get an FVU of 23% for the beam species, 27% for the core species and

18% for the combined species. Differentiating with a longer spatial step—or, similarly, passing the differentiated data through a moving average filter, these FVU scores are reduced. For example, applying a moving average with a window size of 11 (equivalent to averaging simple symmetric differentiation schemes with spatial steps $\delta x = 10\Delta x$ and $\delta x = 12\Delta x$) yields FVU scores of 20%, 23% and 16% respectively for the three cases. With larger window sizes, FVU is decreased further—a window size of 21 e.g. yields FVUs of 19%, 20% and 13%. However, increasing the window size in this way also risks coarse-graining out physical small-scale perturbations, the importance of which is difficult to judge solely from PIC simulation data. Thus, fully evaluating the performance of our \hat{q}_σ closures still requires implementation within a fluid code.

Even without actually performing such an implementation, however, there is a quantity which likely correlates more directly with fluid code performance than even the $\partial_x \hat{q}_\sigma$ FVU: the FVU for our resulting model of $\partial_t p$. After all, this is the quantity we are ultimately aiming to predict accurately by inserting our closure into the pressure equation. Using this quantity, we take into account that predicting $\partial_x q_\sigma$ accurately is only important insofar as the term in question is significant compared to the other terms on the right-hand side of the pressure equation expressed as

$$\partial_t p_\sigma = -V_\sigma \partial_x p_\sigma - 3p_\sigma \partial_x V_\sigma - \partial_x q_\sigma, \quad (32)$$

i.e. $V_\sigma \partial_x p_\sigma$ and $3p_\sigma \partial_x V_\sigma$. For the simulation depicted in Fig. 8, the right-hand side of this equation is predicted with FVU scores of 9%, 14% and 7% for the beam, core and combined electron populations—a significant improvement from e.g. taking $\partial_x q_\sigma \approx 0$, which yields FVU scores of 39%, 51% and 38% for the same cases. In fact, our model predicts the right-hand side of the pressure equation approximately as accurately as $\partial_t p_\sigma$ itself does, calculated using a symmetric temporal $\mathcal{O}(\Delta t^2)$ -accurate finite-differences scheme applied to the PIC p_σ . The FVU scores in the latter case for optimally smoothed spatial derivatives (using a moving average window size of 7) are 16%, 10% and 7%, respectively. Note, however, that the agreement between the two sides of the pressure equation is significantly better than this locally—i.e., when considering the state of the simulation at neighboring *simulation-internal* timesteps; as noted in Sec. II C, these are a factor 100 smaller than the dataset effective timestep Δt .

For most of the other simulations, the accuracy is similar to that of the case shown in Fig. 8. Some simulations are predicted even more accurately—for example, the one with initial condition $\{n_b, V_{\text{rel}}, v_{\text{th},b}, v_{\text{th},c}\} = \{0.033 \bar{n}, 4.7 \times 10^{-2} c, 5.5 \times 10^{-5} c, 7.2 \times 10^{-3} c\}$ reaches FVU scores for the right-hand side of the pressure equation of 6%, 12% and 5% for the beam, core and combined populations respectively. Notably, these values are almost an order of magnitude lower than the 42%, 85% and 48% FVU scores reached with a naive $\hat{q}_\sigma = 0$ model.




FIG. 8: PIC q_σ (left) compared to three-term models \hat{q}_σ implementing the Bernstein-basis rational models of $A_{1,4,5}$ (middle), with the $\hat{q}_\sigma - q_\sigma$ error shown on the right—for the beam (top), core (middle) and combined (bottom) electron species. The depicted simulation uses an initial condition $\{n_b, V_{\text{rel}}, v_{\text{th},b}, v_{\text{th},c}\} = \{0.13 \bar{n}, 3.4 \times 10^{-2} c, 1.8 \times 10^{-3} c, 4.5 \times 10^{-3} c\}$.

In certain cases, however, accuracy is somewhat lower. For example, accuracy as measured by \hat{q}_c FVU is generally poor when the beam—and thus the perturbation of the core—is very weak. This is not a major problem, though, as the core heat flux itself is very weak in such cases (and thus not as important to model). A further example is that of localized heat flux spikes, which sometimes occur during saturation, the magnitude of which is poorly modeled for q_c in simulations with very cold and/or dense core populations. However, even if inaccurate prediction of the exact spike amplitude impacts FVU significantly, it is unclear how large the effect would be on a fluid code simulation implementing the closure.

Another relatively minor issue is that of small-scale perturbations propagating at relativistic speeds ($> 0.1c$), which commonly arise during saturation in simulations with very high initial V_{rel} ($\sim 0.1c$), which are not captured accurately by the closure. This is unsurprising, since our training dataset is overwhelmingly non-

relativistic. Additionally, the moment equations we use are only valid in the non-relativistic limit in the first place—for relativistic systems, closure construction is significantly more complicated^{113–116}.

A somewhat more serious issue at such high- V_{rel} initial conditions is that they can yield phase-space dynamics violent enough to effectively swap the two electron populations for some electron holes during the saturation process—especially in conjunction with high beam densities. This causes V_{rel} to be strongly negative ($\sim -0.05c$) in some parts of the simulation box and strongly positive ($\sim +0.05c$) in others, meaning \bar{V}_{rel} is a poor predictor of local conditions. More accurate simulation of such cases thus likely requires $A_{1,4,5}$ models sensitive to variations in V_{rel} on the scale of the electron holes created by the instability, or a modified implementation where particles can be exchanged between the two species. Regardless, the total simulation FVU in such cases for the right-hand side of the pressure equation remains below 50% even

in the simulations most affected by these phenomena—significantly lower than that of e.g. the $\hat{q}_\sigma = 0$ model, which is often approximately unity in such cases.

IV. DISCUSSION AND CONCLUSIONS

Achieving both high efficiency and high fidelity is a central goal of numerical plasma physics. For fluid models, this requires closures able to capture kinetic effects without fully kinetic simulation. Yet deriving such closures from first principles remains challenging for many collisionless processes. To overcome this, we employ data-driven approaches from machine learning, including neural networks and sparse regression. In a recent article⁸⁹, we used sparse regression to discover a heat flux closure accurately capturing the physics of both Landau-damped Langmuir waves and two-stream instability.

Two significant questions remained: Firstly, while the closure had been demonstrated to capture the heat flux of a single electron species well, it remained unproven for multi-electron species modeling, which is more suitable for fluid modeling of two-stream-unstable setups. Secondly, the coefficients in front of the discovered closure terms were found to vary depending on plasma conditions, meaning the closure needs to be supplemented with a method for estimating the coefficients from, e.g., lower-order fluid moments when implemented in a fluid code.

In this work, we propose and test possible strategies to address both of these issues, through a combination of theoretical and computational approaches:

- We derived a multi-species generalization of the previously found single-species six-term closure from linear collisionless theory and demonstrated how this modified closure is indeed identified consistently by sparse regression for relevant two-stream-unstable setups using a suitable term library.
- Furthermore, we developed a new framework for non-linear sparse regression, and showed how this can be used to estimate the most important free parameters in the original closure from box-averaged fluid quantities at largely neural network-equivalent accuracy.
- Additionally, we showed how the resulting rational coefficient models can be protected from divergence over a domain of interest, at the cost of only a limited decrease in accuracy, by expressing the denominator polynomials in the Bernstein basis.

Using low-order rational functions protected against poles may have an advantage compared to neural network models, since the latter may reduce the convergence order of the simulation due to their uncontrolled character, even with smooth activation functions⁸⁰.

- To ensure that the resulting closure is able to accurately capture the underlying kinetic physics over a large parameter domain, we used first-principles kinetic

simulations in the OSIRIS code with a range of initial conditions to generate our dataset. The resulting models regularly capture 80–90% of the variation in the heat fluxes q_σ for the various species, and predict $\partial_t p_\sigma$ with a typical accuracy of 85–95% over the course of a simulation, further demonstrating the utility of the closure. For comparison, a naive zero heat flux closure typically only yields a 50–60% accuracy for $\partial_t p_\sigma$.

Notably, this high accuracy was achieved despite our rational NLSR models having an FVU error rate of 30% for the A_1 coefficient, roughly twice that of optimal neural network models, which reach 10–20% FVU. A clear pathway towards further increasing the accuracy of the closure is thus to investigate extended term libraries for this coefficient. For the other coefficients of interest, A_4 and A_5 , neural network-equivalent accuracy is reached.

This is not the only potential pathway towards a further improved closure, however. Looking instead towards specific circumstances where closure performance is decreased, we see that strong inhomogeneity over the simulation box is one of the major culprits. For example, performance suffers when electron hole growth is uneven or the sign of V_σ varies over the simulation box. A promising avenue for future work is thus to use more local averages (potentially with e.g. a Gaussian kernel) to predict $A_{1,4,5}$, rather than the full simulation-box averages used in this work. Notably, this approach is also better aligned with ensuring parallelizability of the resulting model.

A further limitation is that the closures we discuss in this work are all non-relativistic. As outlined in Sec. III B, performance is decreased for setups where significant numbers of particles travel at speeds $\gtrsim 0.1c$ relative to one another. Since the framework we are working within is based on the non-relativistic moment equations, this is expected. Generalizing these closure models to handle relativistic dynamics is thus another potential future area of investigation.

Generalizing data-driven closures to three-dimensional electromagnetic plasma phenomena remains a major challenge. When extending sparse regression closures to higher dimensions, the requirement that the resulting heat flux transforms as a symmetric third-rank tensor must be enforced in order to keep the size of the resulting term library manageable. Owing to the strong anisotropy of magnetized plasma dynamics, extending beyond the electrostatic phenomena considered here, the inclusion of electromagnetic fields in the term library is likely unavoidable. This substantially enlarges an already large term library and increases the likelihood of strongly correlated terms, which may render the regression problem poorly conditioned. Ensuring rotational invariance (as well as Lorentz or Galilean invariance, as appropriate) is essential for physical consistency and is also practically beneficial in reducing the search space. Both requirements can be satisfied through an appropriate construction of the closure, while rotational invariance may alternatively be enforced by augmenting the training data⁸⁵ with randomly rotated samples.

The definitive test of a closure’s accuracy is its performance in a fluid code. While implementing and directly evaluating such a closure lies beyond the scope of this article, we have shown that the key three coefficients can be reliably estimated from bulk fluid quantities, suggesting no fundamental barrier to implementation.

Our NLSR framework (Sec. II A 2) offers both advantages and drawbacks relative to neural networks. In its current form, training equivalent neural networks is considerably less computationally demanding, though the cost of NLSR could likely be reduced with further optimization. The key strength of NLSR lies in its relative interpretability: it allows straightforward protection against divergences across any domain of interest and enables analytic computation of derivatives—features that are difficult to achieve with the more opaque structure of neural networks.

Apart from the type of closures discussed in this paper, NLSR could also be applied to discover data-driven sub-grid scale closures. More traditionally derived such models have previously seen use in both neutral fluid dynamics and magnetohydrodynamics to account for spatiotemporally under-resolved physics^{117–119}. There has also been some work directed towards the development of theory-based sub-grid fluid closures in collisionless plasma physics, where the effects of small-scale, pressure anisotropy or heat flux driven instabilities on momentum and heat transport are captured^{4,120}. Work on *data-driven* sub-grid scale modeling of collisionless plasmas, however, has thus far been limited to the acceleration or super-resolution of kinetic simulations^{121,122}.

Ultimately, significant further study is needed to develop closures for the three-dimensional and multi-scale collisionless processes which fundamentally motivate this work, being infeasible to model kinetically. However, the closures and techniques developed here serve as a useful stepping stone in such efforts. With a proper tensorial generalization it stands to reason that the six- and three-term closures developed here can even find direct application in large-scale two- or three-dimensional simulations if used in conjunction with e.g. a neural network to identify regions where electrostatic processes dominate.

ACKNOWLEDGMENTS

The authors are grateful to D. Graham and T. Fülöp for fruitful discussions. The computations used the OSIRIS particle-in-cell simulation code, and were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement No. 2022-06725. The work was supported by the Knut and Alice Wallenberg foundation (Dnr. 2022.0087 and 2023.0249) and the Swedish Research Council (Dnr. 2021-03943), as well as the National Science Foundation Grants No. PHY-2018087 and PHY-2018089.

AUTHOR DECLARATIONS

Conflicts of interest

The authors have no conflicts to disclose.

DATA AVAILABILITY STATEMENT

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

Appendix A: Constraints from multi-species linear theory

Applying linear collisionless theory to fluid models with multiple electron species yields results which are broadly similar to those using a single electron species. There are, however, some significant differences—mainly due to the fact that the species can interact through the electromagnetic field. Within 1D electrostatic linear theory this coupling can be boiled down to each species feeling its own distinct effective complex plasma frequency, $\omega_{p\sigma}$, which can be defined through

$$r_\sigma \omega_{p\sigma}^2 = r \omega_{pe}^2, \quad (\text{A1})$$

where $r_\sigma = \tilde{n}_\sigma / \tilde{n}$ and $r = \tilde{n} / \tilde{n}$, following the notational conventions of Ref. 89. We can also express r in terms of individual species quantities:

$$r = \frac{\sum_\sigma r_\sigma \tilde{n}_\sigma}{\sum_\sigma n_\sigma} = \{2 \text{ species}\} = \frac{r_1 \tilde{n}_1 + r_2 \tilde{n}_2}{\tilde{n}_1 + \tilde{n}_2}. \quad (\text{A2})$$

Apart from this, there is also Doppler shifting of frequencies, since even though we are in the combined-species CoM frame, each species may still individually have a nonzero average flow velocity.

The lowest three moments of the Vlasov equation for each species’ distribution function f_σ , together with Maxwell’s equations, are given by

$$\begin{cases} \partial_t n_\sigma + \partial_x (n_\sigma V_\sigma) = 0 \\ n_\sigma (\partial_t + V_\sigma \partial_x) V_\sigma + \partial_x p_\sigma = -\frac{e}{m_e} n_\sigma E \\ (\partial_t + V_\sigma \partial_x) p_\sigma + 3p_\sigma \partial_x V_\sigma + \partial_x q_\sigma = 0 \\ \partial_x E = \frac{e}{\epsilon_0} (\tilde{n} - n) \\ \partial_t E = \frac{e}{\epsilon_0} n V. \end{cases} \quad (\text{A3})$$

With S different electron species, this is a set of $3S + 2$ equations relating $4S + 1$ unknowns. Since the two remaining Maxwell’s equations imply the combined-species continuity equation, the equations actually only amount to $3S + 1$ constraints, however. Thus, closures for S different quantities (in our case q_σ) are needed to have a

solvable system of PDEs. Since we are only interested in processes occurring at electron timescales, we have taken

the ions to be immobile with number density equal to the average electron density \bar{n} to ensure quasi-neutrality.

Now, let us consider a small wave-like perturbation around equilibrium in the combined-species CoM frame, i.e.

$$\begin{cases} n_\sigma = \bar{n}_\sigma + \tilde{n}_\sigma e^{i(kx - \omega t)} \\ V_\sigma = \bar{V}_\sigma + \tilde{V}_\sigma e^{i(kx - \omega t)} \\ p_\sigma = n_\sigma v_{\text{th},\sigma}^2, \quad v_{\text{th},\sigma} = \bar{v}_{\text{th},\sigma} + \tilde{v}_{\text{th},\sigma} e^{i(kx - \omega t)} \\ q_\sigma = \bar{q}_\sigma + \tilde{q}_\sigma e^{i(kx - \omega t)} \\ E = \tilde{E} e^{i(kx - \omega t)}, \end{cases} \quad (\text{A4})$$

with $\sum_\sigma \bar{n}_\sigma \bar{V}_\sigma = 0$. We assume the wavenumber k to be real, but the frequency $\omega = \omega_r + i\gamma$ is allowed to be complex, γ being the growth rate. Inserting ansatz (A4) into Eq. (A3) and keeping only terms up to first order in the perturbations, we obtain the relations making up multi-species 1D linear collisionless theory:

$$\begin{cases} -i\omega \tilde{n}_\sigma + ik(\tilde{n}_\sigma \bar{V}_\sigma + \bar{n}_\sigma \tilde{V}_\sigma) = 0 \\ \bar{n}_\sigma(-i\omega + ik\bar{V}_\sigma)\tilde{V}_\sigma + ik(\tilde{n}_\sigma \bar{v}_{\text{th},\sigma}^2 + 2\bar{n}_\sigma \bar{v}_{\text{th},\sigma} \tilde{v}_{\text{th},\sigma}) = -\frac{e}{m_e} \bar{n}_\sigma \tilde{E} \\ (-i\omega + ik\bar{V}_\sigma)(\tilde{n}_\sigma \bar{v}_{\text{th},\sigma}^2 + 2\bar{n}_\sigma \bar{v}_{\text{th},\sigma} \tilde{v}_{\text{th},\sigma}) + 3ik\bar{n}_\sigma \bar{v}_{\text{th},\sigma}^2 \tilde{V}_\sigma + ik\tilde{q}_\sigma = 0 \\ ik\tilde{E} = -\frac{e}{\varepsilon_0} \tilde{n}, \end{cases} \quad (\text{A5})$$

where $\tilde{n} = \sum_\sigma \tilde{n}_\sigma$.

Reorganizing these equations to solve for the variations in electric field strength, density, thermal speed and heat flux, we find

$$\begin{cases} \tilde{E} = i\frac{e}{k\varepsilon_0} r\tilde{n} \\ \tilde{V}_\sigma = r_\sigma v'_{\text{ph},\sigma} \\ \tilde{v}_{\text{th},\sigma} = -\frac{1}{2} \left[1 + \frac{\omega_{\text{p}\sigma}^2 - \omega_\sigma'^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \right] r_\sigma \bar{v}_{\text{th},\sigma} \\ \tilde{q}_\sigma = - \left[3 + \frac{\omega_{\text{p}\sigma}^2 - \omega_\sigma'^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \right] r_\sigma \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^2 v'_{\text{ph},\sigma}. \end{cases} \quad (\text{A6})$$

Here, we have introduced the notation $v'_{\text{ph},\sigma} = v_{\text{ph}} - \bar{V}_\sigma$ and $\omega'_\sigma = \omega - k\bar{V}_\sigma$ for the (species-specific) Doppler-shifted complex phase velocity and frequency, respectively. Note that just like with a single electron species, having negligible ion dynamics implies that a heat flux closure needs to agree with the expression for \tilde{q}_σ to first order in r_σ in order to be viable for modeling weak wave-like perturbations. If we make an ansatz

$$\begin{cases} q_{\text{even},\sigma} = A_1 n_\sigma v_{\text{th},\sigma}^2 (V_\sigma - \bar{V}_\sigma) + A_2 \partial_x n_\sigma v_{\text{th},\sigma}^3 + A_3 n_\sigma v_{\text{th},\sigma}^2 \partial_x v_{\text{th},\sigma} \\ q_{\text{odd},\sigma} = A_4 + A_5 n_\sigma v_{\text{th},\sigma}^3 + A_6 n_\sigma v_{\text{th},\sigma}^2 \partial_x V_\sigma, \end{cases} \quad (\text{A7})$$

motivated by the single-species 6-term closure model, we get

$$\begin{aligned} \tilde{q}_{\text{even},\sigma} &= A_1 \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^2 \tilde{V}_\sigma + ikA_2 \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^3 + ikA_3 \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^2 \tilde{v}_{\text{th},\sigma} = \\ &= A_1 r_\sigma \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^2 (v_{\text{ph}} - k\bar{V}_\sigma) + ikA_2 r_\sigma \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^3 - \frac{1}{2} ikA_3 r_\sigma \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^3 \left[1 + \frac{\omega_{\text{p}\sigma}^2 - (\omega - k\bar{V}_\sigma)^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \right] = \\ &= - \left[-A_1 + ik \left(\frac{1}{2} A_3 - A_2 \right) \frac{\bar{v}_{\text{th},\sigma}}{v'_{\text{ph},\sigma}} + \frac{1}{2} ikA_3 \frac{\bar{v}_{\text{th},\sigma}}{v'_{\text{ph},\sigma}} \frac{\omega_{\text{p}\sigma}^2 - \omega_\sigma'^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \right] r_\sigma \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^2 v'_{\text{ph},\sigma} \end{aligned} \quad (\text{A8})$$

and

$$\begin{aligned} \tilde{q}_{\text{odd},\sigma} &= A_5 (\bar{n}_\sigma \bar{v}_{\text{th},\sigma}^3 + 3\bar{n}_\sigma \bar{v}_{\text{th},\sigma}^2 \tilde{v}_{\text{th},\sigma}) + ikA_6 \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^2 \tilde{V}_\sigma = \\ &= A_5 \left(1 - \frac{3}{2} \left[1 + \frac{\omega_{\text{p}\sigma}^2 - \omega_\sigma'^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \right] \right) r_\sigma \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^3 + ikA_6 r_\sigma \bar{n}_\sigma \bar{v}_{\text{th},\sigma}^2 v'_{\text{ph},\sigma}, \end{aligned} \quad (\text{A9})$$

recovering expressions similar to those found in the single-species case, except with $\omega \rightarrow \omega'_\sigma$ (implying $v_{\text{ph}} \rightarrow v'_{\text{ph},\sigma}$) and $\omega_{\text{pe}} \rightarrow \omega_{\text{p}\sigma}$.

Note that the A_1 term needs to be set proportional to $V_\sigma - \bar{V}_\sigma$ rather than proportional to V_σ to recover these expressions—otherwise, additional terms $\sim \tilde{n}_\sigma \bar{v}_{\text{th},\sigma}^2 \bar{V}_\sigma$ and $\sim \tilde{n}_\sigma \tilde{v}_{\text{th},\sigma} \bar{v}_{\text{th},\sigma} \bar{V}_\sigma$ appear. This can be viewed as using the same definition as before, if one interprets the combined-species quantity V present in the original closure as $V - \bar{V}$, which is reasonable, as we are explicitly working in the combined-species CoM frame.

If we for each species further define

$$\Phi_\sigma(\omega, k) = \frac{\omega_{\text{p}\sigma}^2 - \omega'^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \quad (\text{A10})$$

as well as

$$\begin{cases} \alpha_\sigma = \text{Re } \Phi_\sigma = \frac{\xi_\sigma \omega_{\text{pe}}^2 + \gamma^2 - \omega_r'^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \\ \beta_\sigma = -\text{Im } \Phi_\sigma = \frac{2\omega_r' \gamma - \zeta_\sigma \omega_{\text{pe}}^2}{k^2 \bar{v}_{\text{th},\sigma}^2}, \end{cases} \quad (\text{A11})$$

where we have introduced the additional notation $\omega_r' = \text{Re } \omega'_\sigma = \omega_r - k\bar{V}_\sigma$ and $\frac{\gamma}{r_\sigma} = \rho_\sigma = \xi_\sigma + i\zeta_\sigma$ for real ξ_σ, ζ_σ , inserting Eqs. (A8) and (A9) into Eq. (A6) yields a single complex-valued constraint on the closure coefficients for each species:

$$(3 + \Phi_\sigma)\omega'_\sigma = -(A_1 + ikA_6)\omega'_\sigma + \left[-ikA_2 + \frac{1}{2}ikA_3(1 + \Phi_\sigma) + \frac{1}{2}A_5(1 + 3\Phi_\sigma) \right] k\bar{v}_{\text{th},\sigma}. \quad (\text{A12})$$

Just like in the single-species case, we can split this equation into real and (negative) imaginary parts to get two real-valued constraints

$$\begin{cases} (3 + \alpha_\sigma)\omega_r' + \beta_\sigma \gamma = -A_1 \omega_r' + kA_6 \gamma + \frac{1}{2}[kA_3 \beta_\sigma + A_5(1 + 3\alpha_\sigma)]k\bar{v}_{\text{th},\sigma} \\ -(3 + \alpha_\sigma)\gamma + \beta_\sigma \omega_r' = A_1 \gamma + kA_6 \omega_r' + \left[kA_2 - \frac{1}{2}kA_3(1 + \alpha_\sigma) + \frac{3}{2}A_5 \beta_\sigma \right] k\bar{v}_{\text{th},\sigma}. \end{cases} \quad (\text{A13})$$

Solving for A_1 and kA_6 gives

$$\begin{cases} kA_6 = \beta_\sigma + \left[-kA_2 \omega_r' + \frac{1}{2}kA_3(1 + \Phi_{\sigma-})\omega_r' - \frac{1}{2}A_5(1 + 3\Phi_{\sigma+})\gamma \right] \frac{k\bar{v}_{\text{th},\sigma}}{|\omega'_\sigma|^2} \\ A_1 = -3 - \alpha_\sigma + \left[-kA_2 \gamma + \frac{1}{2}kA_3(1 + \Phi_{\sigma+})\gamma + \frac{1}{2}A_5(1 + 3\Phi_{\sigma-})\omega_r' \right] \frac{k\bar{v}_{\text{th},\sigma}}{|\omega'_\sigma|^2}, \end{cases} \quad (\text{A14})$$

where we, similarly to what was done in the single-species analysis in Ref. 89, have introduced the shorthand

$$\Phi_{\sigma\pm} = \frac{\omega_{\text{p}\sigma\pm}^2 \pm |\omega'_\sigma|^2}{k^2 \bar{v}_{\text{th},\sigma}^2}, \quad \omega_{\text{p}\sigma\pm}^2 = (\xi_\sigma + \eta_\pm \zeta_\sigma) \omega_{\text{pe}}^2, \quad \eta_\pm = \mp \left(\frac{\gamma}{\omega_r'} \right)^{\mp 1} \quad (\text{A15})$$

Taking $\gamma \rightarrow 0$, we have no issues with $\Phi_{\sigma-}$, since

$$\Phi_{\sigma-} \omega_r' = \frac{(\omega_r' \xi_\sigma + \gamma \zeta_\sigma) \omega_{\text{pe}}^2 - \omega_r' |\omega'_\sigma|^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \xrightarrow{\gamma \rightarrow 0} \frac{\xi_\sigma \omega_{\text{pe}}^2 - \omega_r'^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \omega_r', \quad (\text{A16})$$

but for $\Phi_{\sigma+}$ we need to be careful. What we have in our expressions is

$$\Phi_{\sigma+} \gamma = \frac{(\gamma \xi_\sigma - \omega_r' \zeta_\sigma) \omega_{\text{pe}}^2 + \gamma |\omega'_\sigma|^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \xrightarrow{\gamma \rightarrow 0} -\frac{\zeta_\sigma \omega_{\text{pe}}^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \omega_r'. \quad (\text{A17})$$

Thus, our constraints become

$$\begin{cases} kA_6 = -\frac{\zeta_\sigma \omega_{\text{pe}}^2}{k^2 \bar{v}_{\text{th},\sigma}^2} + \left[-kA_2 + \frac{1}{2}kA_3 \left(1 + \frac{\xi_\sigma \omega_{\text{pe}}^2 - \omega_r'^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \right) + \frac{3}{2}A_5 \frac{\zeta_\sigma \omega_{\text{pe}}^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \right] \frac{k\bar{v}_{\text{th},\sigma}}{\omega_r'} \\ A_1 = -3 - \frac{\xi_\sigma \omega_{\text{pe}}^2 - \omega_r'^2}{k^2 \bar{v}_{\text{th},\sigma}^2} + \left[-\frac{1}{2}kA_3 \frac{\zeta_\sigma \omega_{\text{pe}}^2}{k^2 \bar{v}_{\text{th},\sigma}^2} + \frac{1}{2}A_5 \left(1 + 3 \frac{\xi_\sigma \omega_{\text{pe}}^2 - \omega_r'^2}{k^2 \bar{v}_{\text{th},\sigma}^2} \right) \right] \frac{k\bar{v}_{\text{th},\sigma}}{\omega_r'}, \end{cases} \quad (\text{A18})$$

which is significantly more complicated than for the single species, due to the presence of ζ_σ . In particular, $A_2 = A_3 = A_6 = 0$ is no longer always a solution of the first equation. Indeed, this can only be true if

$$A_5 = \frac{2\omega'_r}{3k\bar{v}_{\text{th},\sigma}}, \quad (\text{A19})$$

which is a priori unlikely to be fulfilled for an arbitrary 6-term model found by SR. In other words, one should not necessarily expect $A_{2,3,6}$ to correlate clearly with the growth rate when modeling several electron species separately, like they do in the single-species case. If we nevertheless do demand this, inserting this expression for A_5 into the second constraint (together with $A_3 = 0$), several cancellations occur, leaving us with

$$A_1 = -3 + \frac{1}{3} = -\frac{8}{3}, \quad (\text{A20})$$

implying that

$$\begin{cases} A_1 = -\frac{8}{3} \\ A_5 = \frac{2\omega'_r}{3k\bar{v}_{\text{th},\sigma}} \end{cases} \quad (\text{A21})$$

are the linear theory predictions of the A_1 and A_5 coefficients for the 3-term $A_{1,4,5}$ model in the $\gamma \rightarrow 0$ limit.

If we instead set $A_{2,3,6} = 0$ from the start, to get expressions for $A_{1,5}$ in the 3-term $A_{1,4,5}$ model without taking $\gamma \rightarrow 0$, we get the significantly more complicated relations

$$\begin{cases} \frac{2\omega'_r\gamma - \zeta_\sigma\omega_{\text{pe}}^2}{k^2\bar{v}_{\text{th},\sigma}^2} = \frac{3}{2}A_5 \frac{\frac{1}{3}\gamma k^2\bar{v}_{\text{th},\sigma}^2 + (\gamma\xi_\sigma - \omega'_r\zeta_\sigma)\omega_{\text{pe}}^2 + \gamma|\omega'_\sigma|^2}{|\omega'_\sigma|^2 k\bar{v}_{\text{th},\sigma}} \\ A_1 = -3 - \frac{\xi_\sigma\omega_{\text{pe}}^2 + \gamma^2 - \omega_r'^2}{k^2\bar{v}_{\text{th},\sigma}^2} + \frac{3}{2}A_5 \frac{\frac{1}{3}\omega'_r k^2\bar{v}_{\text{th},\sigma}^2 + (\omega'_r\xi_\sigma + \gamma\zeta_\sigma)\omega_{\text{pe}}^2 - \omega'_r|\omega'_\sigma|^2}{|\omega'_\sigma|^2 k\bar{v}_{\text{th},\sigma}}, \end{cases} \quad (\text{A22})$$

which can be rearranged into

$$\begin{cases} A_1 = -3 - \frac{\xi_\sigma\omega_{\text{pe}}^2 + \gamma^2 - \omega_r'^2}{k^2\bar{v}_{\text{th},\sigma}^2} + \frac{2\omega'_r\gamma - \zeta_\sigma\omega_{\text{pe}}^2}{k^2\bar{v}_{\text{th},\sigma}^2} \frac{\frac{1}{3}\omega'_r k^2\bar{v}_{\text{th},\sigma}^2 + (\omega'_r\xi_\sigma + \gamma\zeta_\sigma)\omega_{\text{pe}}^2 - \omega'_r|\omega'_\sigma|^2}{\frac{1}{3}\gamma k^2\bar{v}_{\text{th},\sigma}^2 + (\gamma\xi_\sigma - \omega'_r\zeta_\sigma)\omega_{\text{pe}}^2 + \gamma|\omega'_\sigma|^2} \\ A_5 = \frac{2}{3} \frac{2\omega'_r\gamma - \zeta_\sigma\omega_{\text{pe}}^2}{k^2\bar{v}_{\text{th},\sigma}^2} \frac{|\omega'_\sigma|^2 k\bar{v}_{\text{th},\sigma}}{\frac{1}{3}\gamma k^2\bar{v}_{\text{th},\sigma}^2 + (\gamma\xi_\sigma - \omega'_r\zeta_\sigma)\omega_{\text{pe}}^2 + \gamma|\omega'_\sigma|^2}. \end{cases} \quad (\text{A23})$$

Appendix B: Separate- and combined-species fluid quantities

By definition, the k^{th} order raw moment for a species σ is given by

$$\mathbf{Y}_\sigma = \int d^3\mathbf{v} \mathbf{v}^{(k)} f_\sigma. \quad (\text{B1})$$

These add linearly when species are combined, since the corresponding combined quantity is given by

$$\mathbf{Y} = \int d^3\mathbf{v} \mathbf{v}^{(k)} f = \int d^3\mathbf{v} \mathbf{v}^{(k)} \sum_\sigma f_\sigma = \sum_\sigma \mathbf{Y}_\sigma, \quad (\text{B2})$$

by linearity of integrals. In other words, we have

$$n = \sum_\sigma n_\sigma, \quad (\text{B3})$$

and

$$n\mathbf{V} = \sum_\sigma n_\sigma \mathbf{V}_\sigma \Rightarrow \mathbf{V} = \frac{1}{n} \sum_\sigma n_\sigma \mathbf{V}_\sigma. \quad (\text{B4})$$

For central moments, relating sub-species quantities and combined quantities is a bit more complicated. Taking pressure as an example, the relation is

$$\begin{aligned} \mathbf{p} &= \mathbf{P} - n\mathbf{V}^{(2)} = \sum_\sigma (\mathbf{p}_\sigma + n_\sigma \mathbf{V}_\sigma^{(2)}) - n\mathbf{V}^{(2)} = \\ &= \sum_\sigma [\mathbf{p}_\sigma + n_\sigma \mathbf{V}_\sigma (\mathbf{V}_\sigma - \mathbf{V})]. \end{aligned} \quad (\text{B5})$$

When there are precisely two subspecies, we have

$$\begin{cases} \mathbf{V}_1 - \mathbf{V} = \frac{1}{n} [(n - n_1)\mathbf{V}_1 - n_2\mathbf{V}_2] = \frac{n_2}{n} (\mathbf{V}_1 - \mathbf{V}_2) \\ \mathbf{V}_2 - \mathbf{V} = \dots = -\frac{n_1}{n} (\mathbf{V}_1 - \mathbf{V}_2), \end{cases} \quad (\text{B6})$$

implying

$$n_1(\mathbf{V}_1 - \mathbf{V}) = -n_2(\mathbf{V}_2 - \mathbf{V}) = \frac{n_1 n_2}{n} (\mathbf{V}_1 - \mathbf{V}_2), \quad (\text{B7})$$

which means that our expression for \mathbf{p} reduces to

$$\mathbf{p} = \sum_\sigma \mathbf{p}_\sigma + \frac{n_1 n_2}{n} (\mathbf{V}_1 - \mathbf{V}_2)^{(2)}, \quad (\text{B8})$$

or in 1D

$$p = \sum_{\sigma} p_{\sigma} + \frac{n_1 n_2}{n} (V_1 - V_2)^2. \quad (\text{B9})$$

Analogously, for the heat flux we have

$$\begin{aligned} \mathbf{q} &= \mathbf{Q} - 3\{\mathbf{P}\mathbf{V}\} + 2n\mathbf{V}^{(3)} = \\ &= \sum_{\sigma} \left(\mathbf{q}_{\sigma} + 3\{\mathbf{p}_{\sigma}\mathbf{V}_{\sigma}\} + n_{\sigma}\mathbf{V}_{\sigma}^{(3)} \right) - 3 \left\{ \sum_{\sigma} \left(\mathbf{p}_{\sigma} + n_{\sigma}\mathbf{V}_{\sigma}^{(2)} \right) \mathbf{V} \right\} + 2n\mathbf{V}^{(3)} = \\ &= \sum_{\sigma} \mathbf{q}_{\sigma} + 3 \sum_{\sigma} \{ \mathbf{p}_{\sigma}(\mathbf{V}_{\sigma} - \mathbf{V}) \} + \sum_{\sigma} n_{\sigma} \{ \mathbf{V}_{\sigma}^{(2)}(\mathbf{V}_{\sigma} - \mathbf{V}) \} - 2 \sum_{\sigma} n_{\sigma} \{ \mathbf{V}_{\sigma}(\mathbf{V}_{\sigma} - \mathbf{V}) \mathbf{V} \} = \\ &= \sum_{\sigma} \left[\mathbf{q}_{\sigma} + \left\{ \left(3\mathbf{p}_{\sigma} + n_{\sigma}\mathbf{V}_{\sigma}^{(2)} - 2n_{\sigma}\mathbf{V}_{\sigma}\mathbf{V} \right) (\mathbf{V}_{\sigma} - \mathbf{V}) \right\} \right], \end{aligned} \quad (\text{B10})$$

which for specifically two-species combination reduces according to

$$\begin{cases} 3 \sum_{\sigma} \mathbf{p}_{\sigma}(\mathbf{V}_{\sigma} - \mathbf{V}) = 3(\mathbf{T}_1 - \mathbf{T}_2) \frac{n_1 n_2}{n} (\mathbf{V}_1 - \mathbf{V}_2), \\ \sum_{\sigma} n_{\sigma} \mathbf{V}_{\sigma}^{(2)}(\mathbf{V}_{\sigma} - \mathbf{V}) = (\mathbf{V}_1^{(2)} - \mathbf{V}_2^{(2)}) \frac{n_1 n_2}{n} (\mathbf{V}_1 - \mathbf{V}_2), \\ 2\mathbf{V} \sum_{\sigma} n_{\sigma} \mathbf{V}_{\sigma}(\mathbf{V}_{\sigma} - \mathbf{V}) = 2\mathbf{V} \frac{n_1 n_2}{n} (\mathbf{V}_1 - \mathbf{V}_2)^{(2)}, \end{cases} \quad (\text{B11})$$

so that

$$\begin{aligned} \mathbf{q} &= \sum_{\sigma} \mathbf{q}_{\sigma} + \frac{n_1 n_2}{n} \{ [3(\mathbf{T}_1 - \mathbf{T}_2) + (\mathbf{V}_1 + \mathbf{V}_2 - 2\mathbf{V})(\mathbf{V}_1 - \mathbf{V}_2)] (\mathbf{V}_1 - \mathbf{V}_2) \} = \\ &= \sum_{\sigma} \mathbf{q}_{\sigma} + \frac{n_1 n_2}{n} \left\{ \left[3(\mathbf{T}_1 - \mathbf{T}_2) - \frac{n_1 - n_2}{n} (\mathbf{V}_1 - \mathbf{V}_2)^{(2)} \right] (\mathbf{V}_1 - \mathbf{V}_2) \right\}, \end{aligned} \quad (\text{B12})$$

or in 1D

$$\begin{aligned} q &= \sum_{\sigma} q_{\sigma} + \frac{n_1 n_2}{n} [3(T_1 - T_2) + (V_1 + V_2 - 2V)(V_1 - V_2)](V_1 - V_2) = \\ &= \sum_{\sigma} q_{\sigma} + \frac{n_1 n_2}{n} \left[3(T_1 - T_2) - \frac{n_1 - n_2}{n} (V_1 - V_2)^2 \right] (V_1 - V_2). \end{aligned} \quad (\text{B13})$$

As we can see, the extra term contributing to the combined heat flux is significantly more involved than the corresponding term for the pressure.

Appendix C: The single-species 6-term model

In Ref. 89, the single-species analogue of the term library used in this work was utilized to discover the six-term closure $q = q_{\text{even}} + q_{\text{odd}}$, with

$$\begin{cases} q_{\text{even}} = A_1 n v_{\text{th}}^2 V + A_2 v_{\text{th}}^3 \partial_x n + A_3 n v_{\text{th}}^2 \partial_x v_{\text{th}} \\ q_{\text{odd}} = A_4 + A_5 n v_{\text{th}}^3 + A_6 n v_{\text{th}}^2 \partial_x V, \end{cases} \quad (\text{C1})$$

where V is in the center-of-mass frame, i.e. equivalent to $V - \bar{V}$. This closure was consistently identified by SR for

In other words, p depends not only on p_{σ} , but also on an extra term involving the harmonic mean of the two species' densities and their difference in flow velocity.

both Landau-damped Langmuir waves and two-stream instability. Here, A_1 through A_6 are free parameters, which depend on plasma conditions and must be estimated from other known quantities, as outlined above. The coefficients in q_{even} are “ k -even”, meaning they are invariant under mirroring of the domain (equivalent to exchanging k with $-k$ in Fourier space), while those in q_{odd} are “ k -odd”, and pick up a minus sign under this operation.

Coefficients A_2 , A_3 and A_6 were found to correlate strongly with the instantaneous perturbation growth rate γ in the single-species case; this also carries over into the multi-species case, with some caveats. On the other hand, the A_1 , A_4 and A_5 terms are necessary for accurate q modeling also when the growth rate is zero, such as in the saturated phase of the two-stream instabil-

ity. These latter three coefficients were also consistently found to be more important in general—i.e., even when γ is nonzero—as measured by their respective contributions to lowering the model error rate.

While A_1 and A_5 are dimensionless and A_4 has units of (mass-normalized) heat flux, the three growth-related terms all contain a spatial derivative, meaning the $A_{2,3,6}$ coefficients have units of length. Consequently, they are approximately proportional to the characteristic length scale of variations in the density, temperature and flow velocity, respectively, while A_1 and A_5 remain order unity regardless of these length scales.

That these specific six terms should be identified consistently is not entirely unexpected, as they are the six terms which are nonzero at zeroth or first order in perturbation theory. Specifically, terms A_4 and A_5 are nonzero at zeroth order, and every term except for A_4 contributes at first order. The A_3 term corresponds to the local approximation of the Hammett-Perkins closure originally derived in Ref. 36, corresponding to a heat flux driven by temperature gradients. The A_2 and A_6 terms may similarly be understood as heat fluxes driven by gradients in density (or pressure) and flow velocity, respectively. The three terms found to be most important for accurately modeling q are less transparent when it comes to interpretation, however.

The A_1 term corresponds to an advective heat flux contribution, chiefly carried by fluctuations in V . As discussed in Ref. 89, it may be interpreted as partially canceling the terms proportional to $V\partial_x p$ and $p\partial_x V$ in the 1D pressure equation. Notably, a linearised version of this term also shows up in a heat flux closure which can be derived by using a different three-pole approximation of the plasma response function than the one used to derive the Hammett-Perkins closure^{123,124}. In real space, this closure, which concerns the first-order perturbation δq to the heat flux, looks like

$$\delta q = A_1 \bar{n} \bar{v}_{\text{th}}^2 \delta V - \chi \bar{n} \bar{v}_{\text{th}} \mathcal{H} \delta T, \quad (C2)$$

as outlined in Ref. 124. Here, \mathcal{H} denotes the negative *Hilbert transform*, which also occurs in the exact Hammett-Perkins closure when formulated in real space. In fact, the second term in this alternative closure is precisely the full, nonlocal Hammett-Perkins closure. The only difference is that the value of the coefficient χ which satisfies the linearized moment equations differs depending on whether the first term is included or not. Interestingly, the derivation given in Ref. 124, based in linear theory, gives $A_1 = (3\pi - 8)/(4 - \pi) \approx 1.66$, which has the opposite sign compared to the values we find (but is similar in magnitude). This discrepancy can be ascribed to the presence of the A_5 term in our closure, since a closure with only the $A_{1,4,5}$ terms gives $A_1 = -8/3$ in linear theory (see Appendix A).

The reason for the A_5 term's importance is less clear. It is possible that it helps compensate for the omission of the nonlocal aspects of the Hammett-Perkins heat flux in the temperature gradient approximation. Since the

term in question is also identified as important throughout the nonlinear saturation process in two-stream instability simulations, it may also be connected to nonlinear effects. Finally, the A_4 term is unique in that it does not actually affect closure performance, since the heat flux only appears in the pressure equation through its divergence. It is instead only needed to correct the zeroth-order heat flux induced by the A_5 term.

REFERENCES

- ¹J. M. TenBarge, J. Ng, J. Juno, L. Wang, A. H. Hakim, and A. Bhattacharjee, “An extended MHD study of the 16 October 2015 MMS diffusion region crossing,” *Journal of Geophysical Research: Space Physics* **124**, 8474–8487 (2019).
- ²C. Dong, L. Wang, A. Hakim, A. Bhattacharjee, J. A. Slavin, G. A. DiBraccio, and K. Germaschewski, “Global ten-moment multifluid simulations of the solar wind interaction with Mercury: From the planetary conducting core to the dynamic magnetosphere,” *Geophysical Research Letters* **46**, 11584–11596 (2019).
- ³J. Ng, A. Hakim, L. Wang, and A. Bhattacharjee, “An improved ten-moment closure for reconnection and instabilities,” *Physics of Plasmas* **27**, 082106 (2020).
- ⁴D. A. St-Onge, M. W. Kunz, J. Squire, and A. A. Schekochihin, “Fluctuation dynamo in a weakly collisional plasma,” *Journal of Plasma Physics* **86**, 905860503 (2020).
- ⁵F. Shi, Y. Lin, X. Wang, B. Wang, and Y. Nishimura, “3-D global hybrid simulations of magnetospheric response to foreshock processes,” *Earth, Planets and Space* **73**, 138 (2021).
- ⁶L. Arzamasskiy, M. W. Kunz, J. Squire, E. Quataert, and A. A. Schekochihin, “Kinetic turbulence in collisionless high- β plasmas,” *Phys. Rev. X* **13**, 021014 (2023).
- ⁷R. Achikanath Chirakkara, A. Seta, C. Federrath, and M. W. Kunz, “Critical magnetic Reynolds number of the turbulent dynamo in collisionless plasmas,” *Monthly Notices of the Royal Astronomical Society* **528**, 937–953 (2023).
- ⁸S. Chapman, “The kinetic theory of simple and composite monatomic gases; viscosity, thermal conduction, and diffusion,” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **93**, 1–20 (1916).
- ⁹D. Enskog, *Kinetische Theorie der Vorgänge in mässig verdünnten Gasen.*, Ph.D. thesis, Uppsala universitet (1917).
- ¹⁰S. I. Braginskii, “Transport phenomena in a completely ionized two-temperature plasma,” *Sov. Phys. JETP* **6**, 358–369 (1958).
- ¹¹S. I. Braginskii, “Transport processes in a plasma,” *Reviews of Plasma Physics* **1**, 205–309 (1965).
- ¹²S. Chapman and T. G. Cowling, *The mathematical theory of non-uniform gases: An account of the kinetic theory of viscosity, thermal conduction and diffusion in gases* (Cambridge University Press, Cambridge, England, 1991).
- ¹³W. H. Matthaeus, “Turbulence in space plasmas: Who needs it?” *Physics of Plasmas* **28**, 032306 (2021).
- ¹⁴I. Svenningsson, E. Yordanova, G. Cozzani, Y. V. Khotyaintsev, and M. André, “Kinetic generation of whistler waves in the turbulent magnetosheath,” *Geophysical Research Letters* **49**, e2022GL099065 (2022).
- ¹⁵Y. V. Khotyaintsev, D. B. Graham, K. Steinvall, L. Alm, A. Vaivads, A. Johlander, C. Norgren, W. Li, A. Divin, H. S. Fu, K.-J. Hwang, J. L. Burch, N. Ahmadi, O. Le Contel, D. J. Gershman, C. T. Russell, and R. B. Torbert, “Electron heating by debye-scale turbulence in guide-field reconnection,” *Physical Review Letters* **124**, 045101 (2020).
- ¹⁶L. Richard, L. Sorriso-Valvo, E. Yordanova, D. B. Graham, and Y. V. Khotyaintsev, “Turbulence in magnetic reconnection jets

- from injection to sub-ion scales,” *Physical Review Letters* **132**, 105201 (2024).
- ¹⁷J. E. Stawarz, P. A. Muñoz, N. Bessho, R. Bandyopadhyay, T. K. M. Nakamura, S. Eriksson, D. B. Graham, J. Büchner, A. Chasapis, J. F. Drake, M. A. Shay, R. E. Ergun, H. Hasegawa, Y. V. Khotyaintsev, M. Swisdak, and F. D. Wilder, “The interplay between collisionless magnetic reconnection and turbulence,” *Space Science Reviews* **220**, 90 (2024).
 - ¹⁸J. C. Holmes, R. E. Ergun, D. L. Newman, N. Ahmadi, L. Andersson, O. Le Contel, R. B. Torbert, B. L. Giles, R. J. Strangeway, and J. L. Burch, “Electron phase-space holes in three dimensions: Multispacecraft observations by Magnetospheric Multiscale,” *Journal of Geophysical Research: Space Physics* **123**, 9963–9978 (2018).
 - ¹⁹C. Norgren, M. André, A. Vaivads, and Y. V. Khotyaintsev, “Slow electron phase space holes: Magnetotail observations,” *Geophysical Research Letters* **42**, 1654–1661 (2015).
 - ²⁰K. Steinvall, Y. V. Khotyaintsev, D. B. Graham, A. Vaivads, P.-A. Lindqvist, C. T. Russell, and J. L. Burch, “Multispacecraft analysis of electron holes,” *Geophysical Research Letters* **46**, 55–63 (2018).
 - ²¹K. Steinvall, Y. V. Khotyaintsev, D. B. Graham, A. Vaivads, O. Le Contel, and C. T. Russell, “Observations of electromagnetic electron holes and evidence of Cherenkov whistler emission,” *Physical Review Letters* **123**, 255101 (2019).
 - ²²Y. Dong, Z. Yuan, S. Huang, Z. Xue, X. Yu, C. J. Pollock, R. B. Torbert, and J. L. Burch, “Observational evidence of accelerating electron holes and their effects on passing ions,” *Nature Communications* **14**, 7276 (2023).
 - ²³I. Svenningsson, E. Yordanova, Y. V. Khotyaintsev, M. André, G. Cozzani, and K. Steinvall, “Whistler waves in the quasi-parallel and quasi-perpendicular magnetosheath,” *Journal of Geophysical Research: Space Physics* **129**, e2024JA032661 (2024).
 - ²⁴J.-H. Li, X.-Z. Zhou, Z.-Y. Liu, S. Wang, Y. Omura, L. Li, C. Yue, Q.-G. Zong, G. Le, C. T. Russell, and J. L. Burch, “Direct observations of cross-scale wave-particle energy transfer in space plasmas,” *Science Advances* **11**, eadr8227 (2025).
 - ²⁵M. F. Ivarsen, Y. Miyashita, J.-P. St-Maurice, G. C. Hussey, B. Pitzel, D. Galeschuk, S. Marei, R. B. Horne, Y. Kasahara, S. Matsuda, S. Kasahara, K. Keika, Y. Miyoshi, K. Yamamoto, A. Shinbori, D. R. Huyghebaert, A. Matsuoka, S. Yokota, and F. Tsuchiya, “Characteristic e-region plasma signature of magnetospheric wave-particle interactions,” *Physical Review Letters* **134**, 145201 (2025).
 - ²⁶S. F. Tigik, D. B. Graham, and Y. V. Khotyaintsev, “Electron-scale energy transfer due to lower hybrid waves during asymmetric reconnection,” *Journal of Geophysical Research: Space Physics* **130**, e2024JA033503 (2025).
 - ²⁷J. L. Burch, R. B. Torbert, T. D. Phan, L.-J. Chen, T. E. Moore, R. E. Ergun, J. P. Eastwood, D. J. Gershman, P. A. Cassak, M. R. Argall, S. Wang, M. Hesse, C. J. Pollock, B. L. Giles, R. Nakamura, B. H. Mauk, S. A. Fuselier, C. T. Russell, R. J. Strangeway, J. F. Drake, M. A. Shay, Y. V. Khotyaintsev, P.-A. Lindqvist, G. Marklund, F. D. Wilder, D. T. Young, K. Torkar, J. Goldstein, J. C. Dorelli, L. A. Avanov, M. Oka, D. N. Baker, A. N. Jaynes, K. A. Goodrich, I. J. Cohen, D. L. Turner, J. F. Fennell, J. B. Blake, J. Clemmons, M. Goldman, D. Newman, S. M. Petrinc, K. J. Trattner, B. Lavraud, P. H. Reiff, W. Baumjohann, W. Magnes, M. Steller, W. Lewis, Y. Saito, V. Coffey, and M. Chandler, “Electron-scale measurements of magnetic reconnection in space,” *Science* **352**, aaf2939 (2016).
 - ²⁸E. Yordanova, Z. Vörös, A. Varsani, D. B. Graham, C. Norgren, Y. V. Khotyaintsev, A. Vaivads, E. Eriksson, R. Nakamura, P.-A. Lindqvist, G. Marklund, R. E. Ergun, W. Magnes, W. Baumjohann, D. Fischer, F. Plaschke, Y. Narita, C. T. Russell, R. J. Strangeway, O. Le Contel, C. Pollock, R. B. Torbert, B. J. Giles, J. L. Burch, L. A. Avanov, J. C. Dorelli, D. J. Gershman, W. R. Paterson, B. Lavraud, and Y. Saito, “Electron scale structures and magnetic reconnection signatures in the turbulent magnetosheath,” *Geophysical Research Letters* **43**, 5969–5978 (2016).
 - ²⁹M. Oka, J. Birn, J. Egedal, F. Guo, R. E. Ergun, D. L. Turner, Y. Khotyaintsev, K.-J. Hwang, I. J. Cohen, and J. F. Drake, “Particle acceleration by magnetic reconnection in geospace,” *Space Science Reviews* **219**, 75 (2023).
 - ³⁰L. Richard, Y. V. Khotyaintsev, C. Norgren, K. Steinvall, D. B. Graham, J. Egedal, A. Vaivads, and R. Nakamura, “Electron heating by parallel electric fields in magnetotail reconnection,” *Physical Review Letters* **134**, 215201 (2025).
 - ³¹Q. M. Lu, B. Lembege, J. B. Tao, and S. Wang, “Perpendicular electric field in two-dimensional electron phase-holes: A parameter study,” *Journal of Geophysical Research: Space Physics* **113** (2008), <https://doi.org/10.1029/2008JA013693>.
 - ³²Q. Lu, K. Huang, Y. Guan, S. Lu, and R. Wang, “Energy dissipation in magnetic islands formed during magnetic reconnection,” *The Astrophysical Journal* **954**, 146 (2023).
 - ³³S. Wang, R. Wang, Q. Lu, H. Fu, and S. Wang, “Direct evidence of secondary reconnection inside filamentary currents of magnetic flux ropes during magnetic reconnection,” *Nature Communications* **11**, 3964 (2020).
 - ³⁴Y. V. Khotyaintsev, D. B. Graham, C. Norgren, and A. Vaivads, “Collisionless magnetic reconnection and waves: Progress review,” *Frontiers in Astronomy and Space Sciences* **6** (2019), 10.3389/fspas.2019.00070.
 - ³⁵D. B. Graham, G. Cozzani, Y. V. Khotyaintsev, V. D. Wilder, J. C. Holmes, T. K. M. Nakamura, J. Büchner, K. Dokko, L. Richard, K. Steinvall, C. Norgren, L.-J. Chen, H. Ji, J. F. Drake, J. E. Stawarz, and S. Eriksson, “The role of kinetic instabilities and waves in collisionless magnetic reconnection,” *Space Science Reviews* **221**, 20 (2025).
 - ³⁶G. W. Hammett and F. W. Perkins, “Fluid moment models for Landau damping with application to the ion-temperature-gradient instability,” *Physical Review Letters* **64**, 3019–3022 (1990).
 - ³⁷R. D. Hazeltine, “Transport theory in the collisionless limit,” *Physics of Plasmas* **5**, 3282–3286 (1998).
 - ³⁸G. F. Chew, M. L. Goldberger, and F. E. Low, “The Boltzmann equation and the one-fluid hydromagnetic equations in the absence of particle collisions,” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **236**, 112–118 (1956).
 - ³⁹A. Le, J. Egedal, W. Daughton, W. Fox, and N. Katz, “Equations of state for collisionless guide-field reconnection,” *Physical Review Letters* **102**, 085001 (2009).
 - ⁴⁰O. Ohia, J. Egedal, V. S. Lukin, W. Daughton, and A. Le, “Demonstration of anisotropic fluid closure capturing the kinetic structure of magnetic reconnection,” *Physical Review Letters* **109**, 115004 (2012).
 - ⁴¹L. Wang, A. H. Hakim, A. Bhattacharjee, and K. Germaschewski, “Comparison of multi-fluid moment models with particle-in-cell simulations of collisionless magnetic reconnection,” *Physics of Plasmas* **22**, 012108 (2015).
 - ⁴²F. Rosenblatt, “The perception: A probabilistic model for information storage and organization in the brain,” *Psychological Review* **65**, 386–408 (1958).
 - ⁴³F. Rosenblatt, *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms* (Spartan Books, 1962).
 - ⁴⁴K. Fukushima, “Neocognitron: A Self-organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position,” *Biological Cybernetics* **36**, 193–202 (1980).
 - ⁴⁵J. R. Koza, *Genetic Programming: On the programming of computers by means of natural selection* (MIT Press, 1992).
 - ⁴⁶J. Bongard and H. Lipson, “Automated reverse engineering in nonlinear dynamical systems,” *Proceedings of the National Academy of Sciences* **104**, 9943–9948 (2007).
 - ⁴⁷M. Schmidt and H. Lipson, “Distilling free-form natural laws from experimental data,” *Science* **324**, 81–85 (2009).
 - ⁴⁸N. Makke and S. Chawla, “Interpretable scientific discovery with

- symbolic regression: a review,” *Artificial Intelligence Review* **57**, 2 (2024).
- ⁴⁹R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society: Series B (Methodological)* **58**, 267–288 (1996).
- ⁵⁰T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (Springer, New York, NY, USA, 2009).
- ⁵¹G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning* (Springer, New York, NY, USA, 2013).
- ⁵²S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Discovering governing equations from data by sparse identification of nonlinear dynamical systems,” *Proceedings of the National Academy of Sciences of the United States of America* **113**, 3932–3937 (2016).
- ⁵³S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Data-driven discovery of partial differential equations,” *Science Advances* **3** (2017), 10.1126/sciadv.1602614.
- ⁵⁴T. Mikolov, K. Chen, G. S. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” in *International Conference on Learning Representations* (2013).
- ⁵⁵V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. A. Riedmiller, “Playing atari with deep reinforcement learning,” *ArXiv abs/1312.5602* (2013).
- ⁵⁶V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, “Human-level control through deep reinforcement learning,” *Nature* **518**, 529–533 (2015).
- ⁵⁷J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *Proceedings of the 32nd International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 37, edited by F. Bach and D. Blei (PMLR, Lille, France, 2015) pp. 2256–2265.
- ⁵⁸A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, Vol. 30, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., 2017).
- ⁵⁹J. Hermann, J. Spencer, K. Choo, A. Mezzacapo, W. M. C. Foulkes, D. Pfau, G. Carleo, and F. Noé, “Ab initio quantum chemistry with neural-network wavefunctions,” *Nature Reviews Chemistry* **7**, 692–709 (2023).
- ⁶⁰S.-M. Udrescu and M. Tegmark, “AI Feynman: A physics-inspired method for symbolic regression,” *Science Advances* **6**, eaay2631 (2020).
- ⁶¹R. Dubčáková, “Eureqa: software review,” *Genetic Programming and Evolvable Machines* **12**, 173–178 (2011).
- ⁶²A. Hernandez, A. Balasubramanian, F. Yuan, S. A. M. Mason, and T. Mueller, “Fast, accurate, and transferable many-body interatomic potentials by symbolic regression,” *npj Computational Materials* **5**, 112 (2019).
- ⁶³B. Weng, Z. Song, R. Zhu, Q. Yan, Q. Sun, C. G. Grice, Y. Yan, and W.-J. Yin, “Simple descriptor derived from symbolic regression accelerating the discovery of new perovskite catalysts,” *Nature Communications* **11**, 3513 (2020).
- ⁶⁴J. Martínez-Gil and J. M. Chaves-Gonzalez, “A novel method based on symbolic regression for interpretable semantic similarity measurement,” *Expert Systems with Applications* **160**, 113663 (2020).
- ⁶⁵M. Virgolin, Z. Wang, T. Alderliesten, and P. A. N. Bosman, “Machine learning for the prediction of pseudorealistic pediatric abdominal phantoms for radiation dose reconstruction,” *Journal of Medical Imaging* **7**, 046501 (2020).
- ⁶⁶I. A. Abdellouai and S. Mehrkanoo, “Symbolic regression for scientific discovery: an application to wind speed forecasting,” in *2021 IEEE Symposium Series on Computational Intelligence (SSCI)* (2021) pp. 01–08.
- ⁶⁷P. Lemos, N. Jeffrey, M. Cranmer, S. Ho, and P. Battaglia, “Rediscovering orbital mechanics with machine learning,” *Machine Learning: Science and Technology* **4**, 045002 (2023).
- ⁶⁸E. Candes, J. Romberg, and T. Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory* **52**, 489–509 (2006).
- ⁶⁹J. Yang, K. Yu, Y. Gong, and T. Huang, “Linear spatial pyramid matching using sparse coding for image classification,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition* (2009) pp. 1794–1801.
- ⁷⁰M. Sorokina, S. Sygletos, and S. Turitsyn, “Sparse identification for nonlinear optical communication systems: Sino method,” *Optics Express* **24**, 30433–30443 (2016).
- ⁷¹L. Boninsegni, F. Nüske, and C. Clementi, “Sparse learning of stochastic dynamical equations,” *The Journal of Chemical Physics* **148**, 241723 (2018).
- ⁷²J.-C. Loiseau, B. R. Noack, and S. L. Brunton, “Sparse reduced-order modelling: sensor-based dynamics to full-state estimation,” *Journal of Fluid Mechanics* **844**, 459–490 (2018).
- ⁷³N. T. T. Nguyen, G. T. Kenyon, and B. Yoon, “A regression algorithm for accelerated lattice qcd that exploits sparse inference on the D-Wave quantum annealer,” *Scientific Reports* **10**, 10915 (2020).
- ⁷⁴L. Zanna and T. Bolton, “Data-driven equation discovery of ocean mesoscale closures,” *Geophysical Research Letters* **47**, e2020GL088376 (2020).
- ⁷⁵S. Beetham and J. Capecehatro, “Formulating turbulence closures using sparse regression with embedded form invariance,” *Phys. Rev. Fluids* **5**, 084611 (2020).
- ⁷⁶N. Barbour, W. Dorland, I. G. Abel, and M. Landreman, “Machine-learning closure for vlasov-poisson dynamics in fourier-hermite space,” *arXiv* (2025), accepted for publication in the *Journal of Plasma Physics*, 2504.13313 [physics.plasm-ph].
- ⁷⁷Z. Huang, C. Dong, and L. Wang, “Machine-learning heat flux closure for multi-moment fluid modeling of nonlinear landau damping,” *Proceedings of the National Academy of Sciences* **122**, e2419073122 (2025).
- ⁷⁸M. Luo, C. Heaton, Y. Wang, D. Plummer, M. Fitzgerald, F. Miniati, S. M. Vinko, and G. Gregori, “Time-embedded convolutional neural networks for modeling plasma heat transport,” (2025), *arXiv:2509.06088* [physics.plasm-ph].
- ⁷⁹C. Lamy, B. Dubroca, P. Nicolai, V. Tikhonchuk, and J.-L. Feugeas, “Modeling of electron nonlocal transport in plasmas using artificial neural networks,” *Phys. Rev. E* **105**, 055201 (2022).
- ⁸⁰F. Miniati and G. Gregori, “Learning transport processes with machine intelligence,” *Scientific Reports* **12**, 11709 (2022).
- ⁸¹M. Dam, M. Brons, J. Juul Rasmussen, V. Naulin, and J. S. Hesthaven, “Sparse identification of a predator-prey system from simulation data of a convection model,” *Physics of Plasmas* **24**, 022310 (2017).
- ⁸²A. A. Kaptanoglu, K. D. Morgan, C. J. Hansen, and S. L. Brunton, “Physics-constrained, low-dimensional models for magneto-hydrodynamics: First-principles and data-driven approaches,” *Phys. Rev. E* **104**, 015206 (2021).
- ⁸³E. P. Alves and F. Fiuza, “Data-driven discovery of reduced plasma physics models from fully kinetic simulations,” *Phys. Rev. Res.* **4**, 033192 (2022).
- ⁸⁴A. A. Kaptanoglu, C. Hansen, J. D. Lore, M. Landreman, and S. L. Brunton, “Sparse regression for plasma physics,” *Physics of Plasmas* **30**, 033906 (2023).
- ⁸⁵M. C. McGrae-Menge, J. R. Pierce, F. Fiuza, and E. P. Alves, “Embedding physical symmetries into machine-learned reduced plasma physics models via data augmentation,” (2025), *arXiv:2506.14048* [physics.plasm-ph].
- ⁸⁶J. Donaghy and K. Germaschewski, “In search of a data-driven symbolic multi-fluid ten-moment model closure,” *Journal of Plasma Physics* **89**, 895890105 (2023).

- ⁸⁷W. Cheng, H. Fu, L. Wang, C. Dong, Y. Jin, M. Jiang, J. Ma, Y. Qin, and K. Liu, “Data-driven, multi-moment fluid modeling of Landau damping,” *Computer Physics Communications* **282**, 108538 (2023).
- ⁸⁸P. Sharma, G. W. Hammett, E. Quataert, and J. M. Stone, “Shearing box simulations of the MRI in a collisionless plasma,” *The Astrophysical Journal* **637**, 952 (2006).
- ⁸⁹E. R. Ingelsten, M. C. McGrae-Menge, E. P. Alves, and I. Pusztai, “Data-driven discovery of a heat flux closure for electrostatic plasma phenomena,” *Journal of Plasma Physics* **91**, E64 (2025).
- ⁹⁰J. Ng, *Fluid closures for the modelling of reconnection and instabilities in magnetotail current sheets*, Ph.D. thesis, Princeton University (2019).
- ⁹¹To construct a model that remains accurate across a parameter space with large variations in stability and heat flux, we use the fraction of variance unexplained (FVU), defined in Eq. 4, whose scale invariance facilitates comparison across datasets. Other metrics are also informative: the (root-)mean-square error directly measures the deviation magnitude but also correlates with the variance (standard deviation) of the dataset, while the maximum error is useful for enforcing stability and constraints and for identifying rare but potentially catastrophic errors, although it is statistically unstable, as it is determined by a single sample. As the FVU is effectively the mean square error normalized to the dataset variance, it is dominated by contributions from high-amplitude regions of the dataset. In our case, this is a desirable quality, since we expect our low-amplitude data to be highly noisy, meaning even a good closure may appear to have large relative errors in this regime.
- ⁹²R. A. Fonseca, L. O. Silva, F. S. Tsung, V. K. Decyk, W. Lu, C. Ren, W. B. Mori, S. Deng, S. Lee, T. Katsouleas, and J. C. Adam, “OSIRIS: A three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators,” in *Computational Science — ICCS 2002*, edited by P. M. A. Sloot, A. G. Hoekstra, C. J. K. Tan, and J. J. Dongarra (Springer Berlin Heidelberg, Berlin, Heidelberg, 2002) pp. 342–351.
- ⁹³R. A. Fonseca, S. F. Martins, L. O. Silva, J. W. Tonge, F. S. Tsung, and W. B. Mori, “One-to-one direct modeling of experiments and astrophysical scenarios: pushing the envelope on kinetic plasma simulations,” *Plasma Physics and Controlled Fusion* **50**, 124034 (2008).
- ⁹⁴J. P. Boris and R. A. Shanny, *Proceedings [of the] 4th Conference on Numerical Simulation of Plasmas* (Naval Research Laboratory, 1972).
- ⁹⁵R. W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles* (CRC Press, 2021).
- ⁹⁶H. Schaeffer, “Learning partial differential equations via data discovery and sparse optimization,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **473**, 20160446 (2017).
- ⁹⁷H. Schaeffer and S. G. McCalla, “Sparse model selection via integral terms,” *Phys. Rev. E* **96**, 023302 (2017).
- ⁹⁸D. A. Messinger and D. M. Bortz, “Weak sindy: Galerkin-based data-driven model selection,” *Multiscale Modeling & Simulation* **19**, 1474–1497 (2021).
- ⁹⁹D. R. Gurevich, M. R. Golden, P. A. Reinbold, and R. O. Grigoriev, “Learning fluid physics from highly turbulent data using sparse physics-informed discovery of empirical relations (spider),” *Journal of Fluid Mechanics* **996**, A25 (2024).
- ¹⁰⁰N. M. Mangan, S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Inferring biological networks by sparse identification of nonlinear dynamics,” *IEEE Transactions on Molecular, Biological, and Multi-Scale Communications* **2**, 52–63 (2016).
- ¹⁰¹K. Kaheman, J. N. Kutz, and S. L. Brunton, “SINDy-PI: a robust algorithm for parallel implicit sparse identification of nonlinear dynamics,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **476**, 20200279 (2020).
- ¹⁰²P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods* **17**, 261–272 (2020).
- ¹⁰³S. K. Lam, A. Pitrou, and S. Seibert, “Numba: a LLVM-based Python JIT compiler,” in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM ’15* (Association for Computing Machinery, New York, NY, USA, 2015).
- ¹⁰⁴J. Chok and G. Vasil, “Rational function approximation with normalized positive denominators,” *SIAM Journal on Scientific Computing* **47**, A2699–A2721 (2025).
- ¹⁰⁵I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (The MIT Press, 2016).
- ¹⁰⁶K. Hornik, M. Stinchcombe, and H. White, “Multilayer feedforward networks are universal approximators,” *Neural Networks* **2**, 359–366 (1989).
- ¹⁰⁷The number of free parameters are $25 \cdot 4 + 25 = 125$ in the input layer, $25 \cdot 25 + 25 = 650$ in each hidden layer and $1 \cdot 25 + 1 = 26$ in the output layer, for a total of 2101.
- ¹⁰⁸A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: an imperative style, high-performance deep learning library,” in *Proceedings of the 33rd International Conference on Neural Information Processing Systems* (Curran Associates Inc., Red Hook, NY, USA, 2019).
- ¹⁰⁹J. Ansel, E. Yang, H. He, N. Gimelshein, A. Jain, M. Voznesensky, B. Bao, P. Bell, D. Berard, E. Burovski, G. Chauhan, A. Chourdia, W. Constable, A. Desmaison, Z. DeVito, E. Ellison, W. Feng, J. Gong, M. Gschwind, B. Hirsh, S. Huang, K. Kalambarkar, L. Kirsch, M. Lazos, M. Lezcano, Y. Liang, J. Liang, Y. Lu, C. K. Luk, B. Maher, Y. Pan, C. Puhirsch, M. Reso, M. Saroufim, M. Y. Siraichi, H. Suk, S. Zhang, M. Suo, P. Tillet, X. Zhao, E. Wang, K. Zhou, R. Zou, X. Wang, A. Mathews, W. Wen, G. Chanan, P. Wu, and S. Chintala, “Pytorch 2: Faster machine learning through dynamic Python bytecode transformation and graph compilation,” in *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2, ASPLOS ’24* (Association for Computing Machinery, New York, NY, USA, 2024) pp. 929–947.
- ¹¹⁰S. P. Gary, *Theory of Space Plasma Microinstabilities* (Cambridge University Press, Cambridge, UK, 1993).
- ¹¹¹A. Sauvée, “Scaleogram,” GitHub repository (2021).
- ¹¹²G. R. Lee, R. Gommers, F. Waselewski, K. Wohlfahrt, and A. O’Leary, “PyWavelets: A Python package for wavelet analysis,” *Journal of Open Source Software* **4**, 1237 (2019).
- ¹¹³R. D. Hazeltine and S. M. Mahajan, “Fluid description of relativistic, magnetized plasma,” *The Astrophysical Journal* **567**, 1262 (2002).
- ¹¹⁴E. A. Johnson, “The relativistic Vlasov equation,” (2011).
- ¹¹⁵L. Tinti, G. Vujanovic, J. Noronha, and U. Heinz, “Resummed hydrodynamic expansion for a plasma of particles interacting with fields,” *Phys. Rev. D* **99**, 016009 (2019).
- ¹¹⁶E. R. Most, J. Noronha, and A. A. Philippov, “Modelling general-relativistic plasmas with collisionless moments and dissipative two-fluid magnetohydrodynamics,” *Monthly Notices of the Royal Astronomical Society* **514**, 4989–5003 (2022).
- ¹¹⁷P. Grete, D. G. Vlaykov, W. Schmidt, and D. R. G. Schleicher, “A nonlinear structural subgrid-scale closure for compressible MHD. II. a priori comparison on turbulence simulation data,” *Physics of Plasmas* **23**, 062317 (2016).
- ¹¹⁸M. Nabavi and J. Kim, “Optimising subgrid-scale closures for spectral energy transfer in turbulent flows,” *Journal of Fluid Mechanics* **982**, A18 (2024).

- ¹¹⁹K. Jakhar, Y. Guan, R. Mojjani, A. Chattopadhyay, and P. Hassanzadeh, "Learning closed-form equations for subgrid-scale closures from high-fidelity data: Promises and challenges," *Journal of Advances in Modeling Earth Systems* **16**, e2023MS003874 (2024).
- ¹²⁰J. F. Drake, C. Pfrommer, C. S. Reynolds, M. Ruzskowski, M. Swisdak, A. Einarsson, T. Thomas, A. B. Hassam, , and G. T. Roberg-Clark, "Whistler-regulated magnetohydrodynamics: Transport equations for electron thermal conduction in the high- β intracluster medium of galaxy clusters," *The Astrophysical Journal* **923**, 245 (2021).
- ¹²¹M. Reza, F. Faraji, and J. N. Kutz, "Data-driven inference of high-dimensional spatiotemporal state of plasma systems," *Journal of Applied Physics* **136**, 183301 (2024).
- ¹²²F. Faraji and M. Reza, "Machine learning applications to computational plasma physics and reduced-order plasma modeling: a perspective," *Journal of Physics D: Applied Physics* **58**, 102002 (2025).
- ¹²³P. Hunana, G. P. Zank, M. Laurenza, A. Tenerani, G. M. Webb, M. L. Goldstein, M. Velli, and L. Adhikari, "New closures for more precise modeling of landau damping in the fluid framework," *Phys. Rev. Lett.* **121**, 135101 (2018).
- ¹²⁴P. Hunana, A. Tenerani, G. P. Zank, M. L. Goldstein, G. M. Webb, E. Khomenko, M. Collados, P. S. Cally, L. Adhikari, and M. Velli, "An introductory guide to fluid models with anisotropic temperatures. part 2. kinetic theory, padé approximants and landau fluid closures," *Journal of Plasma Physics* **85**, 205850603 (2019).