

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 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) evolves 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.* 2021*b*, 2023*a*; 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 (Messenger & 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 , \mathbf{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

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

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

$$\theta_j = n v_{th}^\alpha V^{3-\alpha} \quad (2.5)$$

for some integer $\alpha \leq 3$, where $v_{th} = \sqrt{T} = \sqrt{p/n}$, defining $T = p/n$ 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 1D1V.

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
$v_{th}/(0.01c)$:	0.373	0.422	0.503	0.678	0.860	1.03	1.05
$\delta_e k_{char}$:	-60.8	-59.9	-60.2	-62.3	-64.3	-67.6	-67.4
Resulting $ k_{char} \lambda_{D,e}$:	0.227	0.253	0.303	0.422	0.553	0.696	0.708
$\omega_r/\omega_{pe} (\pm 0.03^*)$:	1.00	0.96	0.88	0.80	0.66	0.19	0
γ/ω_{pe} :	0.0515	0.0837	0.137	0.184	0.234	0.273	0.276
$\omega_{pe}\Delta t_L$:	31.3	23.7	20.9	19.0	13.3	12.3	11.9
$\omega_{pe}t_b$:	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 $|k_{char}|\lambda_{D,e}$), 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 $t_b = \sqrt{m_e/(e|k_{char}|E_{rms})}$, with both k_{char} and E_{rms} in this case being evaluated at the time of maximum average 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 $u_{th} = 3.16 \times 10^{-3}c$. We vary n_b as a fraction of the total electron density n_e , keeping the relative velocity $V_{rel} = V_c - V_b$ 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 $v_{th} = \sqrt{p/n} > u_{th}$. More specifically, this combined thermal speed is $v_{th} = [u_{th}^2 + (n_b/n_e)(1 - n_b/n_e)V_{rel}^2]^{1/2}$, with a maximum of $v_{th} \approx 1.05 \times 10^{-2}c$ 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 $1.9\omega_{pe}^{-1}$ 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 $19.0 < \omega_{pe}t < 38.0$, highlighted in red in [figure 1\(b\)](#), meaning $\omega_{pe}\Delta t_L = 19.0$. 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.

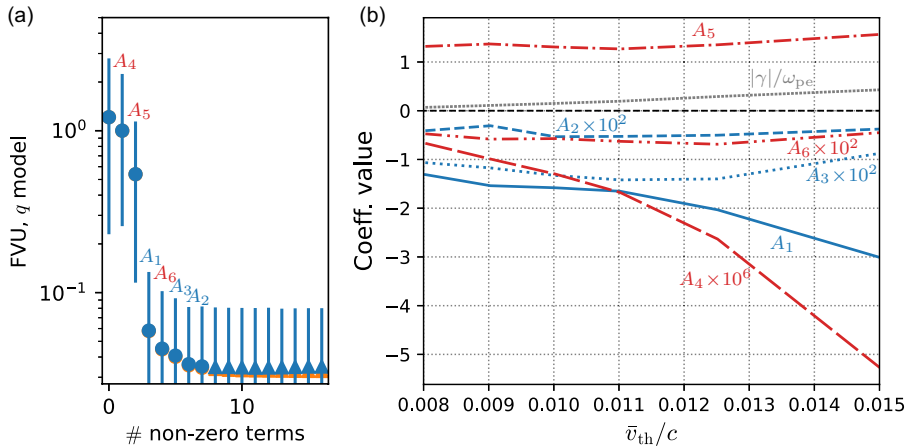


FIGURE 2. Results from the Landau damping simulations. (a) FVU of the successive closures found for $v_{th}/c = 0.01$ (where $\gamma = -0.150$), with each new term indicated by its coefficient in (3.1). Circles, unlike triangles, denote consistently found terms and blue (orange) marker colour corresponds 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 $FVU = \sum_i (y_i - \hat{y}_i)^2 / \sum_i (y_i - \bar{y})^2$, 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

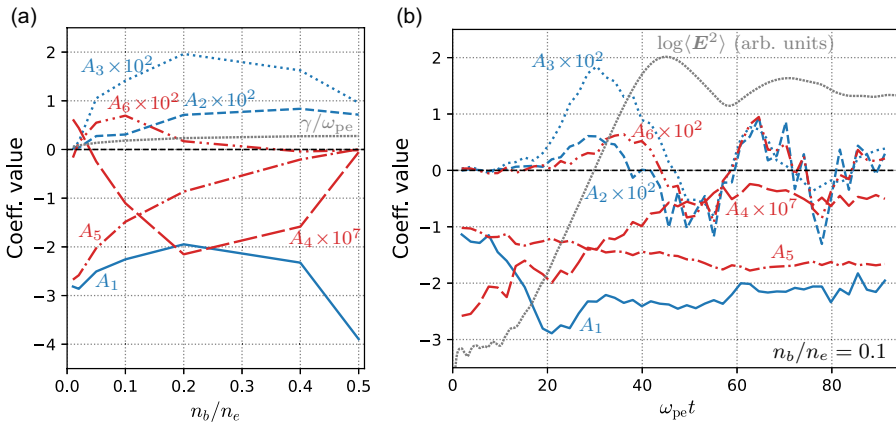


FIGURE 3. Results from the two-stream instability simulations. (a) Dependence of the growth phase closure coefficients on n_b/n_e . For comparison, γ/ω_{pe} is plotted as a grey tightly dotted line. (b) Evolution of the coefficients over time in the case where $n_b/n_e = 0.1$. Blue (red) lines indicate q_{even} (q_{odd}) coefficients. Here, the grey tightly dotted line is the logarithm of the spatially averaged \mathbf{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 $A_1 nv_{\text{th}}^2 V$ term into the model is significantly larger than the one coming from the Hammett–Perkins-like $A_3 nv_{\text{th}}^2 \partial_x v_{\text{th}}$ 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 $\propto v_{\text{th}} \partial_x p = v_{\text{th}}^3 \partial_x n + 2nv_{\text{th}}^2 \partial_x v_{\text{th}}$, 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 $\gamma(t) = \frac{1}{2} \partial_t \ln \langle E^2 \rangle$ 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.

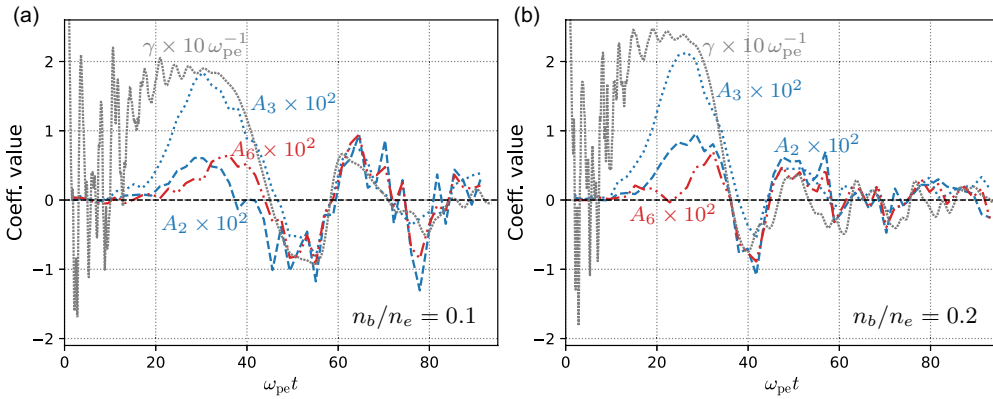


FIGURE 4. Three growth/damping-related coefficients A_2 , A_3 and A_6 , compared with the instantaneous growth rate (a) in the case where $n_b/n_e = 0.1$, also shown in figure 3(b), and in the case where $n_b/n_e = 0.2$. As we can see, A_3 is very nearly precisely proportional to γ , while the other 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 $\{\mathbf{V}\mathbf{p}\}$ 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} : $\mathbf{Q} = n\mathbf{V}^{(3)} + 3\{\mathbf{V}\mathbf{p}\} + \mathbf{q}$. Thus, having a term in our q model equal to specifically $-3\{\mathbf{V}\mathbf{p}\}$ (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

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

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} .

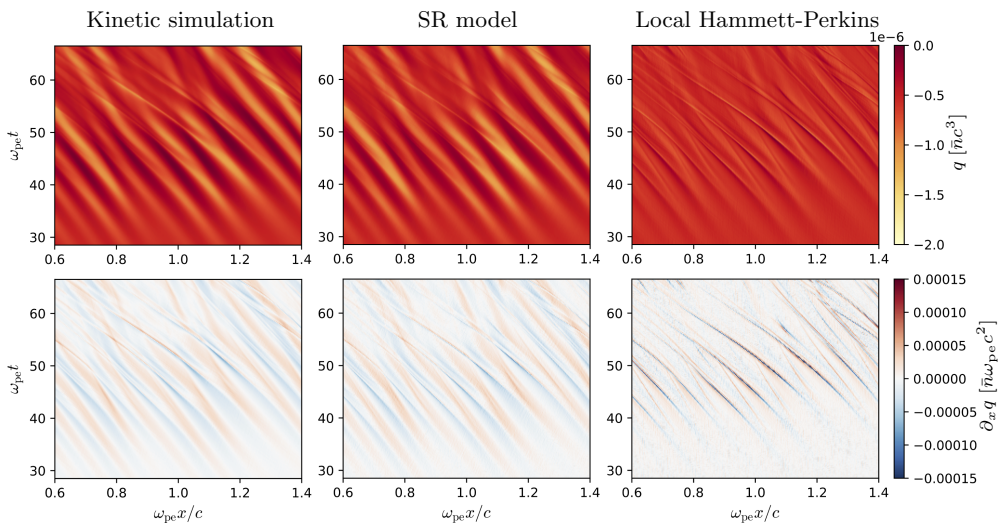


FIGURE 6. A comparison of the OSIRIS q data (top left) from the $n_b/n_e = 0.1$ simulation with our six-term SR model (top middle) and a local Hammett–Perkins model equivalent to keeping only A_3 and A_4 (top right), as well as the resulting $\partial_x q$ (bottom row). Both models are trained solely 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 $\text{FVU} \sim 2\% - 5\%$ 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

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

where we are using shorthand notation

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

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

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

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 $\omega \sim \omega_{\text{pe}} \sim |k|\bar{v}_{\text{th}}$, 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 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 $|\gamma(t)|/\omega_{\text{pe}} < 0.02$ 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

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

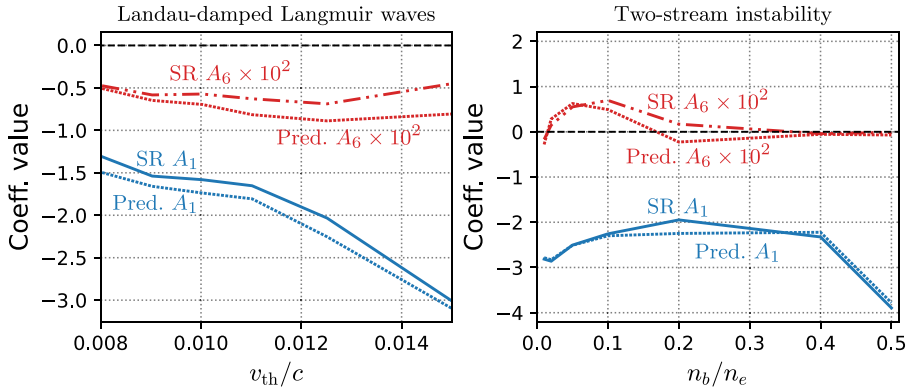


FIGURE 8. Values of A_1 and A_6 found by SR during the growth phase compared with the linear theory prediction given by (3.3), inserting the values of the other coefficients found by SR. The plot on the left contains the results from the Landau-damped Langmuir wave simulations, while the plot on the 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 $Cn v_{th} V^2$ 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 $n v_{th} V^2$ is second-order in r since $V \sim r v_{ph}$, 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 local 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, $\propto n v_{\text{th}} V^2$, 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.

$$\begin{cases} \partial_t f_\sigma + \mathbf{v} \cdot \nabla f_\sigma + \frac{q_\sigma}{m_\sigma} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_\sigma = 0, \\ \nabla \cdot \mathbf{E} = \frac{1}{\varepsilon_0} \sum_\sigma q_\sigma n_\sigma, \\ \nabla \cdot \mathbf{B} = 0, \\ \nabla \times \mathbf{E} = -\partial_t \mathbf{B}, \\ \nabla \times \mathbf{B} = c^{-2} \partial_t \mathbf{E} + \mu_0 \sum_\sigma q_\sigma n_\sigma \mathbf{V}_\sigma, \end{cases} \quad (\text{A.1})$$

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.

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

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

$$\begin{cases} n (\partial_t + V \partial_x) V + \partial_x p = -\frac{e}{m_e} n E, \\ (\partial_t + V \partial_x) p + 3p \partial_x V + \partial_x q = 0, \\ \partial_x E = \frac{e}{\varepsilon_0} (\bar{n} - n), \\ \partial_t E = \frac{e}{\varepsilon_0} n V. \end{cases} \quad (\text{A.4})$$

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:

$$\begin{cases} -i\omega\bar{n}\tilde{V} + ik(\bar{n}\bar{v}_{\text{th}}^2 + 2\bar{n}\bar{v}_{\text{th}}\tilde{v}_{\text{th}}) = -\frac{e}{m_e}\bar{n}\tilde{E}, \\ -i\omega(\bar{n}\bar{v}_{\text{th}}^2 + 2\bar{n}\bar{v}_{\text{th}}\tilde{v}_{\text{th}}) + 3ik\bar{n}\bar{v}_{\text{th}}^2\tilde{V} + ik\tilde{q} = 0, \\ ik\tilde{E} = -\frac{e}{\varepsilon_0}\tilde{n}, \\ -i\omega\tilde{E} = \frac{e}{\varepsilon_0}\bar{n}\tilde{V}, \end{cases} \quad (\text{A.6})$$

or equivalently

$$\begin{cases} \tilde{E} = i\frac{e}{k\varepsilon_0}r\bar{n} \\ \tilde{V} = rv_{\text{ph}}, \\ \tilde{v}_{\text{th}} = -\frac{1}{2}\left(1 + \frac{\omega_{\text{pe}}^2 - \omega^2}{k^2\bar{v}_{\text{th}}^2}\right)r\bar{v}_{\text{th}}, \\ \tilde{q} = -\left(3 + \frac{\omega_{\text{pe}}^2 - \omega^2}{k^2\bar{v}_{\text{th}}^2}\right)r\bar{n}\bar{v}_{\text{th}}^2v_{\text{ph}}. \end{cases} \quad (\text{A.7})$$

To make the notation neater, we have introduced the (complex) phase velocity $v_{\text{ph}} = \omega/k$ and the shorthand notation r for \tilde{n}/\bar{n} , quantifying the amplitude of the perturbation. We have also introduced the electron plasma frequency $\omega_{\text{pe}} = \sqrt{e^2\bar{n}/m_e\varepsilon_0}$. 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.

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

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

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

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

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

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

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

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. & SCHEKOCHIHIN, 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. & GERMASCHESKI, 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. & GERMASCHESKI, 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.

- KHOTYAINTESEV, Y.V., GRAHAM, D.B., NORGREN, 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.
- MESSENGER, 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. & GERMASCHESKI, 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. & SCHEKOCHIHIN, 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. & GERMASCHESKI, 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.