A Learned Closure Method Applied to Phase Mixing in a Turbulent Gradient-Driven Gyrokinetic System in Simple Geometry

A. Shukla[1], D. R. Hatch[1], W. Dorland[2], and C. Michoski[1]

[1]Institute for Fusion Studies, University of Texas at Austin
[2]Department of Physics, University of Maryland

Abstract

We present a new method for formulating closures that learn from kinetic simulation data. We apply this method to phase mixing in a simple gyrokinetic turbulent system—temperature gradient driven turbulence in an unsheared slab. The closure, called the learned multi-mode (LMM) closure, is constructed by, first, extracting an optimal basis from a nonlinear kinetic simulation using singular value decomposition (SVD). Subsequent nonlinear fluid simulations are projected onto this basis and the results are used to formulate the closure. We compare the closure with other closures schemes over a broad range of the relevant 2D parameter space (collisionality and gradient drive). We find that the turbulent kinetic system produces phase mixing rates much lower than the linear expectations, which the LMM closure is capable of capturing. We also compare radial heat fluxes. A Hammett-Perkins closure, generalized to include collisional effects, is quite successful throughout the parameter space, producing $\sim$ 14% Root-Mean-Square (RMS) error. The LMM closure is also very effective: when trained at three (two) points (in a 35 point parameter grid), the LMM closure produces 9% (12%) RMS errors. The LMM procedure can be readily generalized to other closure problems.

1 Introduction

The gyrokinetic model [1, 2, 3], in which the fast gyration of particles around the magnetic field is averaged out, has proven to be a useful description of strongly magnetized plasmas. The kinetic system is reduced from 6D to 5D (3 spatial dimensions and 2 velocity dimensions) and the fast cyclotron timescale is eliminated. The kinetic system is thus greatly simplified for both analysis and simulation.

Gyrokinetics has become the standard tool for describing turbulent transport in magnetic fusion devices, and more broadly, has found fruitful applications ranging from basic plasma physics to space / astro systems [4, 5, 6, 7]. In fusion applications, in particular, gyrokinetic simulations have demonstrated increasing explanatory power with respect to experimental observations [8, 9, 10]. Despite these developments, nonlinear gyrokinetics remains too expensive to be routinely used to predict confinement (i.e., to evolve profiles) or broadly explore parameter space for optimal confinement configurations. Consequently, further reductions in complexity remain highly desirable.

One such approach to further reducing the gyrokinetic system, the gyrofluid framework, was introduced in Ref. [11, 12]. A critical component of gyrofluid models is closures that model important kinetic effects within a fluid treatment. In this paper we study closures in a reduced gyrokinetic system with a Hermite polynomial basis (in velocity space) for a relatively simple turbulent system—gradient-driven turbulence in an unsheared slab. The Hermite basis facilitates a direct comparison of closed fluid simulations (truncated at low Hermite number) with kinetic simulations (truncated at high Hermite number).

A prototypical example of gyrofluid closures is the Hammett-Perkins (HP)[11] closure. The HP procedure closes a fluid system using the linear ki-

netic response. This was a major breakthrough, providing a much more rigorous treatment of collisionless plasmas than conventional fluid theory. It effectively models phase mixing / Landau damping rates resulting in linear growth rates and frequencies in quite good agreement with the true (kinetic) values. Its utility is evidenced by continued vigorous development and application to broad-ranging plasma systems such as turbulent transport in tokamaks [13, 14], tokamak edge turbulence [15, 16, 17], and space plasma turbulence [18, 19]. In this paper, we examine the HP closure in a turbulent system.

We also introduce a new method for learning closures from kinetic simulation data. This method, which we call the learned multi-mode (LMM) closure, is motivated by the notion that a closure for a turbulent system may benefit from the versatility to capture aspects of the nonlinear state. Our closure procedure first extracts, from a single nonlinear kinetic simulation, an optimal basis using singular value decomposition (SVD). Subsequent fluid simulations are projected onto the 'fluid' components of this basis and the projection is used to formulate the closure.

The closure schemes are examined in nonlinear simulations over a broad range of parameter space through the lens of two metrics: (1) the phase mixing rate, and (2) the radial heat flux. The HP closures substantially over-predict the phase mixing rates, which are greatly reduced in comparison with the linear predictions. This is consistent with several recent papers that have noted that Landau damping rates can be greatly modified from linear expectations in turbulent systems (see, e.g. Refs. [20, 21, 22, 23, 24, 25]). In contrast, the LMM closure reproduces phase mixing rates quite accurately.

Despite limitations in reproducing turbulent phase mixing rates, the HP closure is much more accurate in reproducing the kinetic values of the radial heat flux. To be quantitative, an HP closure generalized to include collisional effects results in an RMS error of $\sim 14\%$ over the parameter space.

The LMM closure produces accurate heat fluxes in regions near the training parameter point, with performance deteriorating with distance. Training at multiple, sparsely separated, points results in a highly effective closure. When trained at three (two) points (in a 35 point parameter grid), the LMM closure produces 9% (12%) errors. We envision the utility of the closure to be maximized within a rigorous statistical framework like Bayesian optimization to guide selection of training points.

This paper is outlined as follows: in Sec. 2 we describe the simplified gyrokinetic model and DNA code, which we use to test the performance of various closures. In Sec. 3, we briefly describe HP-style closures and introduce the new learned multi-mode closure. In Sec. 4, we analyze the linear and nonlinear validity of closures by comparing growth rates and phase mixing rates produced by each closure to those of the kinetic system. Sec. 5 evaluates the performance of each closure by comparing the radial heat fluxes throughout a broad parameter space. Advantages, limitations, and possible future avenues of research are described in the concluding section 6.

2 Reduced Gyrokinetic Equations in a Hermite Representation

In order to explore various closure ideas, we study a relatively simple kinetic turbulent system—gradient driven electrostatic instabilities and turbulence in an unsheared slab. The underlying model is a reduction of gyrokinetics to one dimension (parallel to the magnetic field) in velocity space and retaining rudimentary finite Larmor radius (FLR) effects. As a starting point, we consider the gyrokinetic equations for a single kinetic species s in a triply periodic Fourier representation of an unsheared slab:

$$\begin{aligned}\frac{\partial f_s}{\partial t} = &-\left[\omega_n + \omega_T\left(v_{||}^2 + \mu - \frac{3}{2}\right)\right] F_{0s} i k_y \bar{\phi}_s \\ &- \sqrt{2} v_{||} i k_z \left(f_s + F_{0s}\bar{\phi}_s\right) + C(f_s) \\ &+ \sum_{\mathbf{k}'_\perp} \left(k'_x k_y - k_x k'_y\right) \bar{\phi}_{s,\mathbf{k}'_\perp} f_{s,\mathbf{k}_\perp - \mathbf{k}'_\perp},\end{aligned}$$

where the overbar denotes a gyroaverage (i.e. multiplication by the zeroth order Bessel function $J_0(\sqrt{2\mu}k_\perp)$), $k_x(\rho_s)$ is the radial wavenumber (in the direction of the background gradients) and $\rho_s = \left(\frac{T_s}{m_s}\right)^{1/2} \frac{m_s}{q_s B_0}$ is the gyroradius, $k_y(\rho_s)$ is

the binormal (to the magnetic field and the x direction) wavenumber, $k_\perp = \sqrt{k_x^2 + k_y^2}$, $k_z(L_{ref})$ is the parallel (to the magnetic field) wavenumber, $v_{||}(1/v_{th,s})$ is the parallel (to the magnetic field) velocity, the thermal velocity is defined as $v_{th,s} = \left(\frac{2T_s}{m_s}\right)^{1/2}$, $\mu = \frac{m_s v_\perp^2}{2B_0}(B_0/T_{0s})$ is the magnetic moment and acts as the perpendicular velocity coordinate, $t(\frac{\sqrt{T_s/m_s}}{L_{ref}})$ is time, $C \to C(\frac{L_{ref}}{v_{th,s}})$ is a collision operator (defined below), $f_s(\frac{L_{ref}}{\rho_s}\frac{v_{th,s}^3}{n_{0s}})$ is the perturbed distribution function, $F_{0s}(\frac{v_{th,s}^3}{n_{0s}})$ is the background Maxwellian distribution function, $\phi(\frac{L_{ref}}{\rho_s}\frac{T_{0s}}{e})$ is the electrostatic potential, $\omega_n = L_{ref}/L_n$ is the inverse normalized density gradient scale length, $\omega_T = L_{ref}/L_T$ is the inverse normalized temperature gradient scale length, and L_{ref} is a reference macroscopic scale length (normalization is shown in parentheses). The gyrocenter distribution function, $f_{k_x,k_y,k_z}(v_{||},\mu)$, is a function of the three spatial wavenumbers and two velocity coordinates.

The field equation for the electrostatic potential is,

$$\phi_{k_x,k_y} = \frac{\int \bar{f}_s dv_{||} d\mu + \tau\langle\phi\rangle_{FS}\delta_{k_y,0}}{\tau + [1 - \Gamma_0(b)]}, \tag{1}$$

where τ is the ratio of ion to electron temperature, $\Gamma_0(x) = I_0(x)e^{-x}$ with $I_0(x)$ the zeroth order modified Bessel function, $b = k_\perp^2$, and the flux-surface averaged potential is,

$$\langle\phi\rangle_{FS} = \frac{\pi\langle\int \bar{f}_s dv_{||} d\mu\rangle_{FS}}{[1 - \Gamma_0(b)]}. \tag{2}$$

The inclusion of the flux-surface averaged potential in Eq. 2 is appropriate for an ion species driven by the ion-temperature gradient (ITG). Such a system favors strong zonal flow production. In this unsheared slab system, this results in strongly suppressed turbulence [26]. Consequently, we neglect this term in our simulations, which is appropriate for an electron species with adiabatic ions. The gradient drive is now due to the electron-temperature gradient (ETG). In the following, the species labels are dropped with the understanding that all quantities should be considered to be electrons.

The reduced system studied in this work is derived by first integrating over the μ coordinate and retaining rudimentary FLR effects of the form $J_0(\sqrt{2\mu}k_\perp) \to e^{-k_\perp^2/2}$. This approximation is exact only when integrating over a Maxwellian distribution function, as is done for all gyroaverages of the electrostatic potential. More sophisticated treatments are noted in the literature [12], but this rudimentary treatment is sufficient for our purposes, namely to stabilize the instabilities at $k_\perp\rho_s \gtrapprox 1$.

The parallel velocity dimension is then decomposed on a basis of Hermite polynomials $f(v) = \sum_{n=0}^{\infty} f_n H_n(v_{||})e^{-v_{||}^2}$, where n denotes now the order of the Hermite polynomial. The Hermite representation facilitates analysis of the system in both fluid (truncation at low n) and kinetic (truncation at high n) limits and is thus well suited for studying closures. There is a simple connection between these Hermite moments, f_n, and the conventional fluid moments (see Appendix C for details). The Hermite-based equations are as follows [26, 27]:

$$\begin{aligned}\frac{\partial f_{\mathbf{k},n}}{\partial t} = &\frac{\omega_T i k_y}{\pi^{1/4}}\frac{k_\perp^2}{2}\bar{\phi}_{\mathbf{k}}\delta_{n,0} - \frac{\omega_n i k_y}{\pi^{1/4}}\bar{\phi}_{\mathbf{k}}\delta_{n,0} \\ &- \frac{\omega_T i k_y}{\sqrt{2}\pi^{1/4}}\bar{\phi}_{\mathbf{k}}\delta_{n,2} - \frac{i k_z}{\pi^{1/4}}\bar{\phi}_{\mathbf{k}}\delta_{n,1} \\ &- i k_z[\sqrt{n} f_{\mathbf{k},n-1} + \sqrt{n+1} f_{\mathbf{k},n+1}] - \nu n f_{\mathbf{k},n} \\ &+ \sum_{\mathbf{k}'}(k_x' k_y - k_x k_y')\bar{\phi}_{\mathbf{k}'} f_{\mathbf{k}-\mathbf{k}'}\end{aligned} \tag{3}$$

The electrostatic potential is directly proportional to the zeroth-order Hermite polynomial:

$$\bar{\phi}_{\mathbf{k}} = \frac{\pi^{1/4} e^{-k_\perp^2/2} f_{\mathbf{k},0}}{1 + \tau - \Gamma_0(b)}, \tag{4}$$

The first three terms on the right hand side of 3 correspond to the gradient drive, the 4^{th} to landau damping, the 5^{th} to phase mixing, the 6^{th} to collisions, and the last is the nonlinearity.

This system of equations is numerically solved using the DNA code [26, 27].

The phase mixing term, $ik_z[\sqrt{n} f_{\mathbf{k},n-1} + \sqrt{n+1} f_{\mathbf{k},n+1}]$, depends on $f_{\mathbf{k},n\pm1}$ and results in the transfer of energy between scales in phase space (see the following section for a detailed discussion). The dependence of the equation for $f_{\mathbf{k},n}$ on $f_{\mathbf{k},n+1}$ is responsible for the closure problem; the evolution of a given moment depends directly on the next higher order moment, so the set of equations is not closed. Some truncation strategy is required. The simplest closure scheme is naive truncation:

explicitly evolve $n_{\max}$ moment equations, and set $f_{\mathbf{k},n_{\max}+1} = 0$. If the system is sufficiently collisional, low-moment truncation is viable [28]. In a weakly collisional system, if a sufficiently high number of moments is retained, the simulation can be considered to be kinetic and closure by truncation, or via a simple high-n closure [29], generally does not disturb the low order moments [26, 27]. If, however, one wishes to evolve a fluid system (i.e. evolve only a few moments), simple truncation will generally produce deviations from the kinetic system, particularly at low collisionality where Landau damping / phase mixing is an important effect.

2.1 Free Energy Equations

In order to understand the effects and limitations of various closures, it is useful to conceptualize the turbulent dynamics in the context of an energy equation. The free energy [27] is given by:

$$\varepsilon_{\mathbf{k},n} = \varepsilon^{(\phi)}_{\mathbf{k}} \delta_{n,0} + \varepsilon^{(f)}_{\mathbf{k},n} \tag{5}$$

with field component

$$\varepsilon^{(\phi)}_{\mathbf{k}} = \frac{1}{2}\left(\tau + 1 - \Gamma_0\left(k_\perp^2\right)\right)^{-1} |\phi_{\mathbf{k}}|^2 \tag{6}$$

and entropy component

$$\varepsilon^{(f)}_{\mathbf{k},n} = \frac{1}{2}\pi^{1/2} |f_{\mathbf{k},n}|^2 . \tag{7}$$

The free energy evolution equation can be obtained from 3 and 4:

$$\frac{\partial \varepsilon^{(\phi)}_{\mathbf{k},n}}{\partial t} = J^{(\phi)}_{\mathbf{k}} \delta_{n,0} + N^{(\phi)}_{\mathbf{k},n} \tag{8}$$

and

$$\frac{\partial \varepsilon^{(f)}_{\mathbf{k},n}}{\partial t} = \omega_T Q_{\mathbf{k}} \delta_{n,2} - C_{\mathbf{k},n} - J^{(\phi)}_{\mathbf{k}} \delta_{n,1} + J_{\mathbf{k},n-1/2} - J_{\mathbf{k},n+1/2} + N^{(f)}_{\mathbf{k},n}. \tag{9}$$

The terms on the RHS of 8 and 9 represent various energy injection, dissipation, and transfer channels. The only energy sink—collisional dissipation $C_{\mathrm{k},n} = 2\nu n \varepsilon_{\mathrm{k},n}$—is directly proportional to the Hermite number n multiplied by the free energy. The energy source $\omega_T Q_{\mathrm{k}} = \omega_T \operatorname{Re}\left[-\frac{\pi^{1/4}}{2^{1/2}} i k_y f_2^* \bar{\phi}\right]$ is proportional to the perpendicular heat flux $Q_{\mathbf{k}}$.

There are also two conservative energy transfer channels. The nonlinear energy transfer $N^{(f)}_{\mathbf{k},n}$ redistributes energy in k space but does not transfer energy between different n and is not a net source or sink (it vanishes under summation in k-space).

$J^{(\phi)}_{\mathbf{k}} = \operatorname{Re}\left[-i k_z \phi^{1/4} \bar{\phi}^* f_{\mathbf{k},1}\right]$ is the energy transferred between the field component at $n = 0$ and the entropy component (i.e., Landau damping).

For our purposes of studying closures, the most important terms are the linear phase mixing terms $J_{\mathrm{k},n-1/2} = \operatorname{Re}\left[-\pi^{1/2} i k_z \sqrt{n} f^*_{\mathrm{k},n} f_{\mathrm{k},n-1}\right]$ and $J_{\mathbf{k},n+1/2} = \operatorname{Re}\left[\pi^{1/2} i k_z \sqrt{n+1} f^*_{\mathbf{k},n} f_{\mathbf{k},n+1}\right]$. These terms also represent a conservative energy transfer channel, albeit in velocity space. They conservatively transfer energy between n and $n-1$, $n+1$ respectively but do not transfer energy in k-space. One way to characterize the closure problem is determining the proper value of f_{n+1} so that $J_{\mathbf{k},n+1/2}$ sends the proper amount of energy to higher order moments—or, as the case may be, *receives* the proper amount of energy *from* higher order moments. Below, in Sec. 4, we will analyze several closures in terms of their capacity to recover the proper (turbulent, kinetic) rates of energy transfer in phase space.

3 Closures

In this section, we describe several closure schemes as applied to our reduced gyrokinetic system. All closure schemes are of the same class: $f_{\mathbf{k},4} = \sum_{i=0}^{3} A_{\mathbf{k},i} f_{\mathbf{k},i}$, i.e., closures that express the last moment in terms of a linear combination of the lower moments. Some closures will have coefficients $A_{\mathbf{k},i}$ that are specific to the wave-vector, $\mathbf{k}$ but others will not, instead having $A_{\mathbf{k},i} = A_i$ for all $\mathbf{k}$.

In a kinetic model where a large number of moments is retained, truncation, which entails setting $f_{\mathbf{k},n_{max}+1} = 0$, can be used. Alternatively, a simple high-n closure as described in Ref. [29] can be applied. However, our goal in this work is to formulate a fluid model that captures the relevant kinetic physics while retaining only the most thermodynamically-relevant quantities, namely the first four moments. To achieve this, we require a closure for $f_{\mathbf{k},4}$ that is more intelligent than simple truncation. The following subsections describe HP-

style closures and the new LMM closure scheme.

3.1 HP-style closures

Here we provide a brief description of HP-style closures. Derivations and verification of these closures can be found in Appendix A. The HP closure [11, 30] is designed so that the dispersion relation, also referred to as the kinetic response function, arising from the hierarchy of closed moment equations matches the linear kinetic dispersion relation arising from the Vlasov-Poisson kinetic system. The exact kinetic response function, which involves the plasma dispersion function, $Z(\omega)$, is

$$R_{00}(\omega) = -iZ(\omega). \tag{10}$$

The HP closure for the Nth moment takes the form $f_N = \sum_{i=0}^{N-1} A_i f_i(\omega)$. Combining this closure ansatz with the hierarchy of moment equations results in an approximate response function $R^a_{00}(\omega)$, a polynomial in ω involving the closure coefficients, A_i.

The HP closure enforces a match in the low frequency limit, $\omega \to 0$, so the Taylor expansion for plasma dispersion function can be used, which turns Eq. 10 into a polynomial in ω. The closure coefficients A_i can then be chosen so that $R^a_{00}(\omega) = R_{00}(\omega)$. A detailed derivation of the HP closure can be found in Appendix A.

The HP closure for the 4^{th} moment in our system is

$$f_{\mathbf{k},4} = sgn(k_z) A_3 f_{\mathbf{k},3} + A_2 f_{\mathbf{k},2} \tag{11}$$

where the coefficients are $A_3 = -1.759i$ and $A_2 = 0.755$. These coefficients are the same for all $\mathbf{k}$, so the only $\mathbf{k}$-dependence for this closure comes from the $sgn(k_z)$.

In order to test HP-style closures in collisional regimes, we consider a generalization of the HP closure, developed by Snyder in [31], which also includes the effects of collisionality. We have developed a collisional extension of the HP closure, the Hammet-Perkins-Collisional (HPC) closure, which is inspired by Snyder's method but modified to match the low frequency limit through second order.

The procedure for arriving at the HPC closure for the Nth moment, $f_{\mathbf{k},N}$, is as follows. First, use the HP method to determine the closure for the N+1th moment, then substitute this expression for $f_{\mathbf{k},N+1}$ into the linearized time evolution equation for $f_{\mathbf{k},N}$ and take the low frequency limit of this equation ($\partial f_{\mathbf{k},N}/\partial t = 0$). Differentiating this equation with respect to time and then using the low frequency limit of the time evolution equation for the $N-1^{th}$ moment yields a collisional closure for the N^{th} moment.

The HPC closure for the 4^{th} moment in our system is

$$f_{\mathbf{k},4} = \frac{-3.051ik_z\nu - 1.759ik_z^2 sgn(k_z)}{1.838\nu^2 + 3.709k_z\nu sgn(k_z) + k_z^2} f_{\mathbf{k},3} + \frac{0.755k_z^2}{1.838\nu^2 + 3.709k_z\nu sgn(k_z) + k_z^2} f_{\mathbf{k},2} \tag{12}$$

This is a 2^{nd} order accurate (for small ω) closure for $f_{\mathbf{k},4}$ in terms of $f_{\mathbf{k},3}$ and $f_{\mathbf{k},2}$ including collisional effects. A detailed derivation of this closure and its coefficients can be found in Appendix B.

Note that if one takes the collisionless limit, $\nu \to 0$, of Eq. 12, the collisionless closure given in Eq. 11 is recovered.

Both the HPC and HP closures were initially designed for models based on the conventional fluid moments in which the n^{th} fluid moment is calculated by integrating the kinetic distribution times velocity to the n^{th} power. Subsequent work generalized the procedure for Hermite-based systems [30]. The relationship between the Hermite moments and the fluid moments is very simple and is shown in Appendix C.

3.2 The Learned Multi-Mode Closure

We now ask the question of how a closure may be generalized for a nonlinear system in which the turbulent dynamics continually perturb the relationships between the low order moments retained in the system.

This is motivated, in part, by several recent results showing discrepancies between linear and nonlinear phase mixing dynamics. Refs. [20, 21] investigate the effect of a stochastic forcing term on Landau damping rates, demonstrating large deviations from the linear expectations for some parameters. Refs. [22, 24, 25] demonstrate a 'fluidization' of collisionless plasma turbulence—i.e., a large reduction of Landau damping rates due to the cancellation of

the forward velocity space cascade due to turbulence. Likewise, Ref. [23] observes Landau damping rates far smaller than the linear predictions in a turbulent system (see Fig. 10 of that paper). We thus posit that in order to capture the phase mixing rates appropriate for a turbulent, kinetic system, a closure should be endowed with the versatility to adapt to the nonlinear state. To this end, we propose a closure scheme that learns directly from the turbulent kinetic system.

To illustrate the closure strategy, consider the Hermite-based system described in Sec. 2 at two different truncation levels: (1) a four-moment fluid system, and (2) a kinetic system of N Hermite moments, where N is large enough that the system is effectively kinetic (in our simulations we opt for $N = 48$). For a given wavevector, $\mathbf{k}$, an eigenvector of the linear operator is simply a vector with the complex values of each moment—i.e., a 4D vector in the fluid system and an ND vector in the kinetic system.

The following closure approach is conceptually similar to the HP approach. For a given set of physical parameters (gradient drive, collisionality), solve for the linear eigenvector of the kinetic system. Then use the relationship between g_4 and g_3 from this kinetic eigenvector to close the fluid system. If the linear eigenvector persists unmodified in the nonlinear state, this approach would be sufficient. However, as described above, important nonlinear modifications are observed in turbulent systems. Consequently, our strategy is to 'learn' an appropriate closure directly from the turbulent kinetic system.

We do so by extracting from a nonlinear kinetic simulation an 'optimal' basis for the nonlinear turbulent state at each wave vector. For a four-field fluid model, we extract this optimal basis from the first five moments of the kinetic system in order to retain the information necessary to close the system. The turbulent *fluid* state is then projected onto these basis vectors (with the fifth moment of each removed). Since these basis vectors are attached also to the kinetic information (i.e. the fifth moment), this projection can be used to close the system. Mathematical details are described in the next subsection.

Since an optimal basis was extracted from a kinetic simulation, we would expect this procedure to be effective at the parameter point of the kinetic 'training' simulation. The utility of this method, however, will depend on the closure retaining efficacy in some non-negligible parameter domain surrounding the training point. We demonstrate below that this is the case.

We call this closure strategy the 'Learned multi-mode' (LMM) closure because (1) it 'learns' the closure coefficients from the full turbulent kinetic system, and (2) it employs multiple modes (basis vectors) in order to better capture the dynamical variations in the turbulent state.

We end this section by noting some connections with other lines of research. First, this closure approach is related to various strategies for projection-based model reduction [32, 33, 34, 35, 36, 37], wherein basis vectors are extracted (often via SVD) from data describing a complex system to reduce the complexity of the underlying models.

We also note some connections with the closure proposed in Refs. [38, 39]. This closure scheme employs two modes (the ITG mode and its complex conjugate) in order to enforce a 'non-dissipative' closure—i.e., it eliminates any energy transfer between the fluid moments and higher order moments. Consequently, it produces damping rates that are far below (i.e., zero) the linear values, qualitatively similar to the nonlinear results cited above. However, the true turbulent system allows energy to shift dynamically between lower and higher order moments. Consequently, we view this closure as a compelling idea, but one that is perhaps too restrictive.

We also note the connection between the LMM closure and the line of research exploring the role of damped eigenmodes in plasma microturbulence [40, 41, 42, 23, 43], which shows that multiple modes coexisting at a single wavevector play a crucial role in turbulent energetics. Our 'multi-mode' closure also acknowledges the activity of multiple eigenmodes per wavevector and defines the closure coefficients in terms of the relative amplitude of these modes in the nonlinear state.

3.3 Implementation of the LMM Closure

Here we describe the mathematical details of the approach outlined in the previous section. The closure requires a nonlinear kinetic simulation to formulate a set of basis vectors. In our case, we use

48 Hermite moments for the full kinetic simulation. Any number of subsequent fast fluid simulations can then be run requiring explicit computation of only f_0, f_1, f_2, and f_3. In this section we will use bold uppercase letters to denote matrices and bold lowercase letters to refer to vectors.

The full kinetic simulation is used as follows. Let $\mathbf{F}_{N\times M}$ (M is the number of time points and N is number of moments retained in the fluid model plus one) be the matrix created from the simulated distribution function at a single wave vector. The distribution function at a single wavevector is written $f_i(t)$, where $i = 0, 1, ..., N-1$ denotes the Hermite number and t takes on discrete values t_j with $j = 0, 1, ..., M-1$ (the wave vector is suppressed for clarity), so that element ij of $\mathbf{F}$ is $\mathbf{F}_{ij} = f_i(t_j)$:

$$\mathbf{F} = \begin{bmatrix} f_0(t_0) & f_0(t_1) & \cdots & f_0(t_{M-1}) \\ f_1(t_0) & f_1(t_1) & \cdots & f_1(t_{M-1}) \\ \vdots & \vdots & \ddots & \vdots \\ f_{N-1}(t_0) & f_{N-1}(t_1) & \cdots & f_{N-1}(t_{M-1}) \end{bmatrix} \tag{13}$$

The SVD of $\mathbf{F}$ is given by

$$\mathbf{F}_{N\times M} = \mathbf{U}_{N\times N}\mathbf{\Sigma}_{N\times N}\mathbf{V}^H_{N\times M} \tag{14}$$

where $\mathbf{U}$ and $\mathbf{V}$ are unitary and $\mathbf{\Sigma}$ is diagonal with real entries. General background information about this extremely useful matrix decomposition and be found in Ref. [44] and a review on its application to turbulence as proper orthogonal decomposition (POD) can be found in Ref. [45]. The columns of the matrix $\mathbf{U}$ are called the left singular vectors. In our application, they define N basis vectors for the distribution function. The rows of $\mathbf{V}^H$ are the time traces of the amplitude of each of these vectors. The diagonal entries in $\mathbf{\Sigma}$ define the singular values, which encompass all the amplitude information. The utility of the SVD lies in its property that the outer product between the first basis vector and the first time trace (weighted by the corresponding singular value) reproduces more of the fluctuation data (as measured by the Frobenius norm) than any other possible decomposition of this form. Likewise the superposition of the first two (n) outer products captures more of the fluctuation data than any other rank two (n) decomposition and so forth. For convenience, we define a matrix $\mathbf{B}$, which weights the basis vectors by their corresponding singular values so that they include the amplitude information: $\mathbf{B} = \mathbf{U}\mathbf{\Sigma}$.

For the purposes of our desired four moment model, we select $N = 5$ (i.e. only a small subset of the 48 total Hermite moments). Since different Hermite moments are only connected to their direct neighbors, this is sufficient to fully exploit the information in the simulation defining the natural (kinetic, turbulent) relations between f_3 and f_4.

Let $\mathbf{f}$ represent the column vector of the first four moments at a single time step: $\mathbf{f} = [f_0\ f_1\ f_2\ f_3]^T$. In each time step of a subsequent fluid simulation, we numerically advance $\mathbf{f}$ explicitly via Eq. 3. The truncated moment, f_4, is calculated as follows. First, we project the state vector $\mathbf{f}$ onto the basis formed by the columns of $\mathbf{B}$. This entails finding the projection coefficients that define the amount of each SVD mode in the turbulent state at a given point in time. We will call the column vector containing these projection coefficients $\mathbf{c}$. We can do this by removing the row corresponding to the unknown N^{th} moment (the 5^{th} row) from $\mathbf{B}$ and extracting $\mathbf{c}$ from the following equation:

$$\mathbf{f} = \mathbf{M}\mathbf{c} \tag{15}$$

where $\mathbf{M}$ denotes the submatrix of $\mathbf{B}$ consisting of the first 4 rows and all 5 columns of $\mathbf{B}$, i.e., $\mathbf{M}$ is the submatrix produced by removing the last (5^{th}) row of $\mathbf{B}$. This gives

$$\mathbf{c} = \mathbf{M}^\dagger\mathbf{f} \tag{16}$$

where † denotes the pseudo-inverse.

Now that we have $\mathbf{c}$, a length N vector of the inferred mode amplitudes, we can predict f_4 by applying these mode strengths to the previously removed row of $\mathbf{B}$, $\mathbf{b}_5$. This gives

$$f_4 = \mathbf{b}_5\mathbf{c} = \mathbf{b}_5\mathbf{M}^\dagger\mathbf{f} = \mathbf{c}_{LMM}\mathbf{f} \tag{17}$$

where $\mathbf{b}_5$ is the 5^{th} row of $\mathbf{B}$ and $\mathbf{c}_{LMM} = \mathbf{b}_5 M^\dagger$ is the vector containing the 4 LMM closure coefficients. This procedure is repeated at each wavevector $\mathbf{k}$ to obtain a full set of coefficients that can be used to conduct an LMM-closed simulation.

This procedure results in a closure that reflects the natural relations between moments in the turbulent kinetic system and adapts to the relative amplitude of each basis vector in the nonlinear state.

Regarding computational cost, the LMM closure comes down to the dot product between two length

4 vectors: the closure coefficients, c_{LMM}, and the lower order moments, $\mathbf{f}$. The closure coefficients are computed ahead of time and saved to a file, which is loaded at the beginning of the simulation. During the simulation, the computational expense of the LMM closure is very similar to that of the HP closure; the HP closure requires two complex multiplications per wavevector per time step (one for each of the two HP closure coefficients), and the LMM closure requires four complex multiplications per wavevector per time step. This is much less demanding than the pseudo-spectral computation of the nonlinearity, so the increased expense is negligible. The main additional expense is in running nonlinear kinetic simulations for training. If this can be done sparsely, then the LMM closure is viable.

4 Preliminary Closure Tests

In this section we probe the properties of several closures in comparison with the kinetic system in both linear and nonlinear scenarios.

The HP closure has been shown to faithfully reproduce kinetic Landau damping rates and linear growth rates. We reproduce this result for our system: simulations exhibit good agreement between kinetic linear growth rates and fluid growth rates using the HP closure. A representative example is shown in Fig. 1 (top panel), where it is seen that the HP closure, the HPC closure, and the LMM closure all reproduce the growth rates of the linear kinetic system. Growth rates are produced by solving the linearized eigenvalue problem given by Eq. 3 for the 48-moment (kinetic system) and the 4-moment fluid system with each of the closures.

Fig. 1 (lower panels) also shows a simple test of the eigenmode structures by plotting the 4^{th} moment, f_4, normalized to the zeroth moment. These plots are highly relevant since ratios of moments are closely connected to the closure problem. In the 2^{nd} and 3^{rd} panels, we plot the real and imaginary parts of this quantity for the 48-moment (kinetic) linear system, the 48-moment nonlinear system, as well as all of the linear 4-moment systems closed by the HP, HPC, and LMM closures. In the 4^{th} and 5^{th} panels, we plot the real and imaginary parts of this quantity for the the 48-moment nonlinear system, as well as all of the nonlinear 4-moment systems closed by the HP, HPC, and LMM closures. For this example, the LMM closure was trained at parameter point $\omega_T = 9, \nu = 0.1$.

While the growth rates produced by all the closed systems match the linear kinetic growth rates very closely, the agreement is not as good in the plots of these moment ratios in the 2^{nd} and 3^{rd} panels. f_4 produced by the HP and HPC closures, which are both based on the linear system, exhibit a similar shape to the linear kinetic f_4. However, f_4 of the nonlinear kinetic system exhibits a significantly different shape in k_y: the ratio is much smaller. This is closely mirrored by f_4 produced by the linear system closed by the LMM closure. The capacity of the LMM closure to reproduce the nonlinear result is perhaps unsurprising, as it is based on the nonlinear system.

The 4^{th} and 5^{th} panels show that the ratios of f_4 to f_0 in nonlinear HP and HPC simulations matches the nonlinear kinetic simulation much more closely than the ratios from the linear HP and HPC systems. Apparently, the HP approach retains the capacity to adapt to the nonlinear state, which will be discussed further below. The imaginary part of the ratio from the nonlinear LMM-closed system matches the kinetic ratio the best of the three closures shown, but the real part of this ratio is consistently smaller than the kinetic ratio for all k_y shown.

This is an initial indication that dynamics of the linear and nonlinear systems are quite different, consistent with the literature discussed above [20, 21, 23, 25].

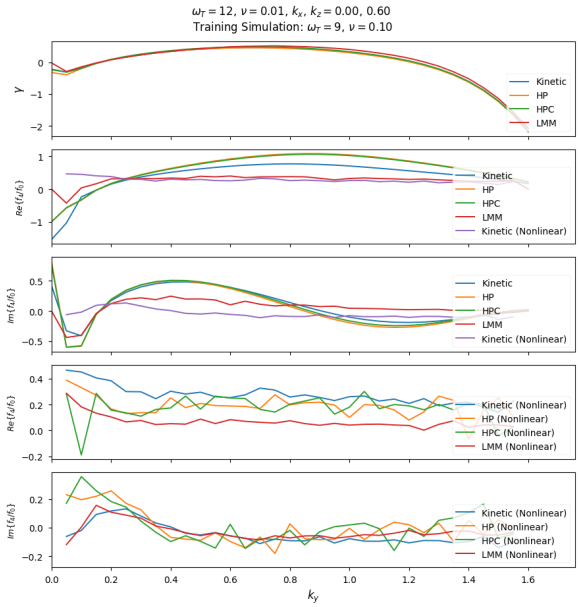

Figure 1: Linear growth rates (top panel) and real (2^{nd} panel) and imaginary (3^{rd} panel) parts of f_4 normalized to f_0 produced by solving the eigenvalue problem given by the linearized version of Eq. 3 plotted against k_y for temperature gradient drive (ω_T) = 12, collision frequency (ν) = 0.01 and k_x, k_z = 0, 0.6. The eigenvalue problem is solved using the linear 48-moment (kinetic) system and also using the 4 moment system closed with the HP closure, HPC closure, and LMM closure. The LMM closure coefficients used to produce this figure were extracted from the kinetic simulation at parameter point ($\omega_T, \nu = 9, 0.1$). Panels 2 and 3 also show the time averaged value of f_4/f_0 from the nonlinear kinetic simulation. Panels 4 and 5 show the real and imaginary parts of the time averaged value of f_4/f_0 from nonlinear kinetic simulations as well as nonlinear LMM,HP, and HPC simulations.

5 Nonlinear Closure Tests

In order to more thoroughly examine closure performance, simulations covering a wide range of temperature gradients, ω_T, and collision frequencies, ν, were conducted with a fully (reduced gyro-) kinetic model (48 moments : $n_{max} = 48$), simply truncated model (4 moments, 5^{th} is set to 0), the standard Hammett-Perkins (HP) closure retaining 4 moments, the HPC closure, and three different LMM closures.

The scan covers $\omega_T = 5, 6, 7, 8, 9, 12, 15$, and $\nu = 0.01, 0.05, 0.1, 0.2, 0.5$. All closures perform very poorly in the completely collisionless regime, $\nu = 0$. This is currently under investigation and may require a more careful treatment of the dissipation in the kinetic system, which is left for future work (collisionless results are shown in Appendix D). Other simulation parameters used are $\omega_n = 1$ and $\tau = 1$.

The grids used in $k-$space are $k_{x,min} = 0.05$, $k_{x,max} = 1.5$, $k_{y,min} = 0.05$, $k_{y,max} = 1.5$, and $k_{z,min} = 0.1$, $k_{z,max} = 3.6$. Hyper-collisions of the form $\nu_h(n/n_{(max)})^8 f_n$ are included in the full kinetic simulations in order to enforce decaying moments at high n. Hyper-diffusion of the form $\nu_\perp(k_{x,y}/k_{x,y}^{(max)})^8$ is included as a small scale dissipation mechanism (intended to roughly account for, e.g., nonlinear perpendicular phase mixing). We use $\nu_h = 0.1$ and $\nu_\perp = 1$ in our simulations. In addition, a Krook term is applied to the zero and minimum finite k_z modes in order to avoid slow growth of these modes that fails to saturate. Refs. [26, 27] describe such numerical considerations in more detail.

For reference, the exact HP closure used was $f_{\mathbf{k},4} = 0.755 f_{\mathbf{k},2} - i(1.759 sgn(k_z)) f_{\mathbf{k},3}$. The HPC closure used was

$$f_{\mathbf{k},4} = \frac{-3.051 i k_z \nu - 1.759 i k_z^2 sgn(k_z)}{1.838\nu^2 + 3.709 k_z \nu sgn(k_z) + k_z^2} f_{\mathbf{k},3} + \frac{0.755 k_z^2}{1.838\nu^2 + 3.709 k_z \nu sgn(k_z) + k_z^2} f_{\mathbf{k},2} \quad (18)$$

The derivation of these coefficients is described in Appendices A and B.

A key question for the LMM closure is the parameter domain over which the closure remains viable. We would expect the applicability of a set of LMM closure coefficients to deteriorate as the distance in parameter space from the training simulation increases. Of course, the computational expense of kinetic simulations requires that the number of training simulations be kept minimal in order for the closure to be useful. In order to probe the question of how far the closure applies throughout parameter space, we selected 3 kinetic training simulations spread throughout the parameter grid.

The three different LMM closures are obtained

by applying the method described in Sec. 3.3 to extract coefficients from kinetic simulations at $\omega_T, \nu = 6, 0.01$, $\omega_T, \nu = 9, 0.1$, at $\omega_T, \nu = 12, 0.5$. We refer to these three LMM closures as LMM-Left, LMM-Middle, and LMM-Right respectively, indicating the region of the scanned parameter space within which their training simulation lies.

5.1 Tests of Energy Dissipation

In order to gain insight into the nonlinear dynamics and their effect on the closure problem, we investigate the energy evolution equation, Eq. 9. In Eq. 9, the contribution from phase mixing defines the energy flux to higher order moments [27]. More specifically, $\chi_{n+1/2} \equiv J_{n+1/2}/(|k_z||f_n|^2)$, the normalized rate at which energy is transferred to/from higher order moments (the phase mixing rate), is defined by a correlation between two neighboring moments. The linear physics defines a fixed, dissipative, relationship between f_n and f_{n+1}. In the presence of turbulence, however, the various moments are continually perturbed by the nonlinearity, resulting in correlations that can differ substantially from the linear expectation.

These considerations are illustrated in Fig. 2, which shows the distribution (accumulated over time) of the energy transfer rate between the 3^{rd} and 4^{th} moments, $\chi_{3+1/2}$, for kinetic, LMM-closed, HP-closed, and HPC-closed simulations. The average dissipation, $\bar{\chi}_{3+1/2}$ resulting from the HP and HPC closures is much larger than the dissipation present in the kinetic system. We note that the HP and HPC closures would likely perform better by this metric with the inclusion of more moments, which may be explored in future work. The LMM-closure, being based on the nonlinear system, produces dissipation that matches the kinetic level quite closely. We note that the LMM closure coefficients used are extracted from a training simulation with different parameters than the simulation being examined (both parameter points are noted in the title of the figure), which indicates the effectiveness of the LMM closure even in simulations with parameters different from the training parameters.

Fig. 3 shows the ratio of the average value of $\chi^{Closed}_{3+1/2}$ to the average value of $\chi^{Kinetic}_{3+1/2}$ at the most unstable wavevector for the HP, HPC, and LMM closures for every point in the extended parameter scan. Three different sets of LMM closure coefficients are used to produce this plot: one set is extracted from the kinetic simulation at $\omega_T, \nu = 6, 0.01$, one from $\omega_T, \nu = 9, 0.1$, and one from $\omega_T, \nu = 12, 0.5$. These three different LMM closures are detailed in sec. 5. At each point in our parameter grid, we use the LMM closure that was trained closest to that grid point to produce the values of $\bar{\chi}^{LMM}_{3+1/2}$ shown in this plot.

The HP and HPC closures overestimate dissipation levels throughout most of the parameter grid. They perform best at low gradient drive (ω_T) with deteriorating performance as gradient drive is increased. The LMM closure matches kinetic dissipation levels significantly better than both HP closures throughout the parameter space.

It is important to note that the substantial disagreement between kinetic and HP/HPC dissipation levels is somewhat misleading. The discrepancy between the heat flux saturation levels examined in the next section is much less than is suggested by the large difference in phase mixing rates.

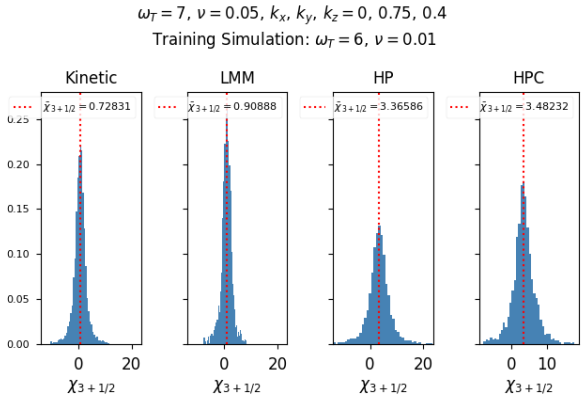

(a)

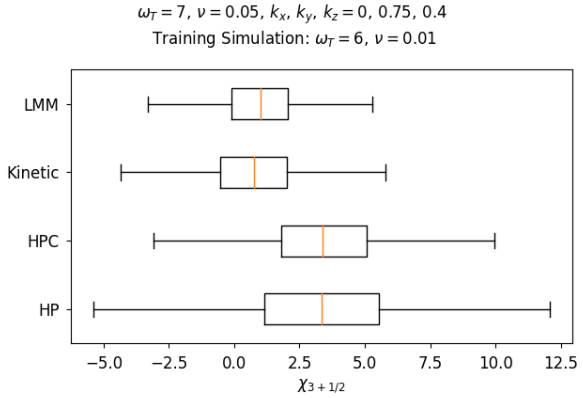

(b)

Figure 2: Probability distribution functions (a) and box and whisker plots (b) showing the distribution of $\chi_{3+1/2}$, the energy transferred between the 3^{rd} and 4^{th} moments, in the Kinetic, LMM-closed, HPC-closed, and HP-closed systems for $\omega_T = 7$ and $\nu = 0.05$ at the most unstable wavevector, $k_x = 0$, $k_y = 0.75$, $k_z = 0.4$. The LMM coefficients used to produce this plot were extracted from the kinetic simulation at $\omega_T, \nu = 6, 0.01$. Red dashed lines show the average value of each distribution, $\bar{\chi}_{3+1/2}$.

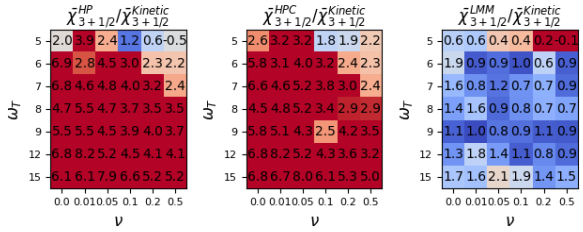

Figure 3: Ratios of the average value of $\chi^{closed}_{3+1/2}/\chi^{Kinetic}_{3+1/2}$ at the most unstable wavevector throughout parameter space for the HP, HPC, and LMM closures. Values below 1 indicate not enough dissipation and values above 1 indicate too much dissipation. This heatmap is set up such that fractions far from 1 in either direction are penalized the same way. For example, ratios of 0.5 and 2 will be the same color. As shown here, the LMM closure matches kinetic dissipation levels much better than the HP and HPC closures throughout most of our parameter grid.

5.2 Comparison of Heat Fluxes

Ultimately, we would like closed simulations to reproduce the most important macroscopic behavior of gyrokinetic simulations, notably the saturated value of the radial heat flux, $Q = \sum_{k_x,k_y,k_z} Q_{\mathbf{k}}$, where $Q_{\mathbf{k}}$ is defined in the discussion surrounding Eq. 9. We view this metric—the proximity of the saturated heat flux for a given closure scheme to that of the kinetic simulation—to be the most relevant metric for the performance of the closure.

Time traces of the heat fluxes produced in the kinetic simulation and seven closed simulations are shown for each combination of input parameters, ω_T and ν, in Fig. 4. The final saturation levels of each simulation type at each set of input parameters were calculated by averaging over the last 30% of the time trace. Each plot in Fig. 5 shows the percent error in saturated heat flux, $(Q^{Closed} - Q^{Kinetic})/Q^{Kinetic}$, for all parameter combinations for each closure scheme. Versions of figures 4 and 5 including $\nu = 0$ can be found in Appendix D. Comparisons are complicated somewhat by occasional shifts in transport levels that occur unpredictably in time, which introduces a level of uncertainty that cannot be eliminated within the

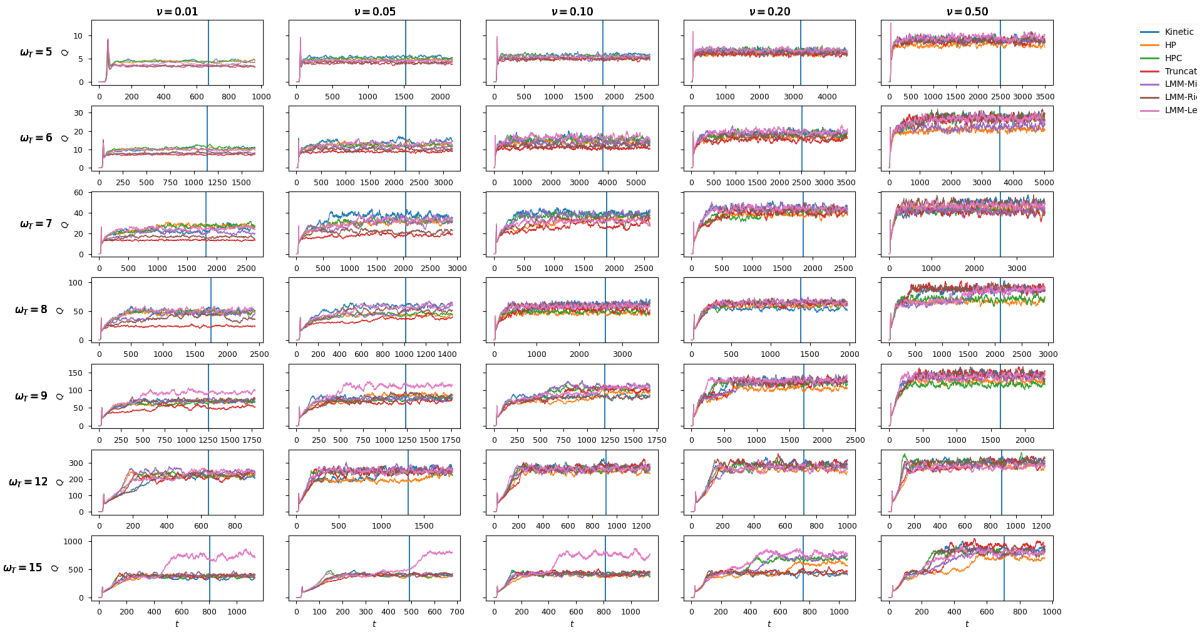

Figure 4: Time traces of the total radial heat flux (Q) for Kinetic (blue), HP (orange), HPC (green), truncated (red), LMM-Middle (purple), LMM-Right (brown), and LMM-Left (pink) simulations for temperature gradient drives (ω_T) ranging from 5 to 15 (increasing downward by panel) and collision frequencies (ν) ranging from 0.01 to 0.5 (increasing to the right by panel). The metric of performance is the final saturation level. The vertical blue lines show the cutoff point - 70% of the simulation time - after which each heat flux curve is averaged to get the final saturation level.

scope of this paper. This is a manifestation of metastable states, recently elucidated in Ref. [46].

As expected, truncation performs poorly at low values of collisionality, but improves at high collisionality where the simulations become more fluid-like. Truncation still performs relatively poorly at high collisionality for low gradient drive, but performs well when both collisionality and gradient drive are large—i.e., for parameters at which other effects (gradient drive or collisions) dominate phase mixing.

The HP closure works well in the low collisionality regime for which it was designed (note the small errors at $\nu = 0.01$). However, its performance deteriorates as collisionality is increased.

The HPC closure is designed to simultaneously include the effects of collisions and phase mixing in the appropriate limits. As expected, it exhibits notable improvement over the HP closure in the high collisionality regime while also retaining the strong performance of the collisionless HP closure at low collisionality. This closure performs poorly only at intermediate levels of gradient drive ($\omega_T = 7 - 9$) and in one simulation at $\omega_T, \nu = 15, 0.2$. This closure appears to be highly effective and its performance is only surpassed by the LMM closure trained at multiple parameter points, described below.

The LMM-Middle closure based on the kinetic simulation at $\omega_T, \nu = 9, 0.1$ surprisingly performs poorly in the simulation at its training parameter point, likely due to the propensity of this system toward metastable states (note the sudden jump in the LMM time trace toward the end of the simulation). The model does, however, perform well in nearby regions in the middle of our scanned parameter space. In fact, this closure extends throughout parameter space quite well, displaying low errors everywhere except in the top left corner (low collisionality and gradient drive).

The LMM-Right closure based on the kinetic simulation at $\omega_T, \nu = 12, 0.5$ produces very low er-

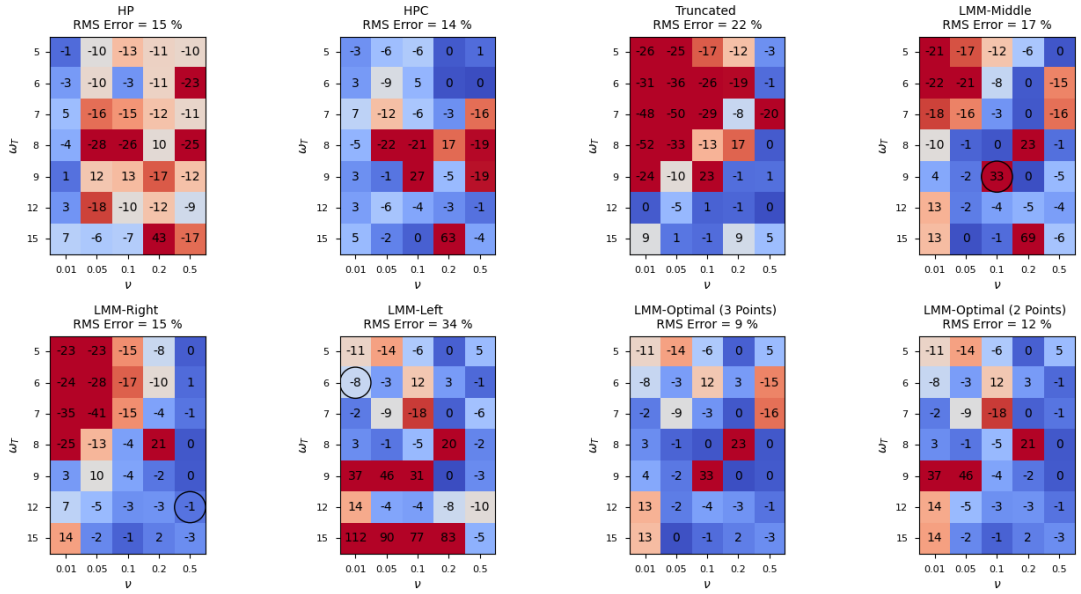

Figure 5: Percent error in saturated heat flux for each closure (HP, HPC, Truncation, LMM-Middle, LMM-Right, LMM-Left, LMM-Optimal) as compared against the kinetic simulation throughout parameter space. The circles on the 4^{th}, 5^{th}, and 6^{th} figures indicate which simulations from which the LMM coefficients were extracted. Percent errors are calculated as $(Q^{Closed} - Q^{Kinetic})/Q^{Kinetic} \times 100$ where Q^{closed} is calculated by averaging the last 30% of the time trace of the heat flux from the closed simulation and $Q^{Kinetic}$ is calculated by averaging the last 30% of the time trace of the heat flux from the kinetic simulation. The RMS error for each closure is also shown above each plot. The 7^{th} plot, LMM-Optimal (3 Points), displays error of the LMM closure trained at the nearest parameter point using three training points. The 8^{th} plot, LMM-Optimal (2 Points), displays the error of the LMM closure trained at the nearest parameter point using only the left and right training points.

rors at high collisionality and gradient drive, in the middle of the parameter region, and even at low gradient drive if collisionality is high. However, it also starts to deteriorate as simulations venture farther from the training point—at low collisionality and gradient drive.

The LMM-Left closure based on the kinetic simulation at $\omega_T, \nu = 6, 0.01$ works well at low gradient drive regardless of collisionality. At high values of gradient drive however, its performance is poor. Specifically at the highest value of gradient drive, $\omega_T = 15$, this closure has extremely large errors. Inspecting Fig. 4, one can see the reason for this. The LMM-Left heat flux initially appears to saturate at an accurate level, but part way through the simulation, it jumps to a higher level and resaturates. This behavior can also be seen in kinetic simulations at $\omega_T, \nu = 15, 0.2$ and $\omega_T, \nu = 15, 0.5$. When this jump occurs in the closed simulation but not the kinetic, it results in very large errors. These simulations illustrate the potential hazard of applying an LMM closure in a regime too far removed from its training point.

We note several sets of simulations with sudden jumps between saturation levels, likely indicative of metastable states. This is particularly clear in the lower right hand corner of the parameter grid, but can be observed throughout the parameter space. It is likely that some fraction of the errors can be attributable to this phenomenon. One example is the LMM-Middle training point, described above. The large HPC error at $\omega_T = 15, \nu = 0.2$ is another candidate. Unfortunately, this error can not be eliminated within the scope of this work. We do, however, probe this phenomenon in some more detail.

First, we ran a second set of kinetic simulations with identical physical parameters to the first set. The initial condition is also identical save for a random component, which, in some cases can result in deviations in the long-time simulations. Results are shown in Fig. 4. We compared the heat flux saturation levels between the two sets of simulations to quantify the uncertainty in kinetic saturation levels. The percent difference between the saturation level in each initial simulation was compared to the saturation level of its counterpart with slightly different initial conditions. There was an 5% RMS difference between the saturation levels in the two sets of simulations. There was one simulation at $\omega_T, \nu = 7, 0.5$ in which the simulation from the first set experienced a jump from one saturated state to another but the simulation from the second did not. So for this particular parameter point, there was a 20% difference in the saturation levels. This illustrates the metastable nature of our system, which introduces some uncertainty into the determination of saturated heat flux levels.

As another test, we probed the kinetic simulation at $\omega_T = 15, \nu = 0.5$. We changed the collisionality of this simulation from $\nu = 0.5$ to $\nu = 0.2$ and restarted the simulation with the end-point of the previous simulation as the initial condition. The results of this test are shown in Fig. 6. The system finds a saturated state considerably higher than its counterpart with identical physical parameters. The simulation which starts at $\nu = 0.5$ initially saturates at $Q \approxeq 840$ but resaturates at $Q \approxeq 760$ once ν is switched to $\nu = 0.2$. The simulation where $\nu = 0.2$ the entire time is saturated at $Q \approxeq 460$. It appears that $Q = 760$ and $Q = 460$ are both stable states for the system with parameters $\omega_T, \nu = 15, 0.2$. Which of these metastable states the system falls into depends on the previous state of the system. This test clearly demonstrates that our system has multiple metastable states and exhibits hysteresis. We refer the reader again to a recent, more detailed, study of this phenomenon [46, 47].

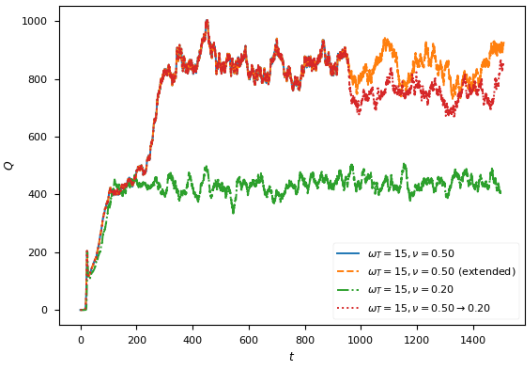

Figure 6

We note that the LMM-Right and LMM-Middle closures are similarly accurate to the HP and HPC closures throughout the parameter space, even taking into account the simulations far removed from the training point. The LMM-left closure exhibits the worst performance of all the closures due to the large late-time jumps in saturation levels at high gradient drive.

Ultimately, we envision the LMM closure applied in scenarios where kinetic training simulations can be supplied sparsely throughout parameter space. Consequently, we show in Fig. 5 the closure called *LMM-Optimal*, which is defined by selecting the LMM-closure that is trained at the nearest point in parameter space. For example, at $\omega_T, \nu = 5, 0.05$, the LMM-Left closure with coefficients extracted from the $\omega_T, \nu = 6, 0.01$ kinetic simulation is used and at $\omega_T, \nu = 15, 0.5$, the LMM-Right closure with coefficients extracted from the $\omega_T, \nu = 12, 0.5$ kinetic simulation is used. In Fig. 5, we include an LMM-Optimal closure that makes use of all three training points and another that uses only the left and right training points. Using two training points instead of three only slightly reduces performance; using three training points results in an RMS error of 9% and using two yields an RMS error of 12%. Tables 1 and 2 in Appendix D show which training simulation was used to run the LMM-closed simulation at each grid point for the LMM-Optimal (3 Point) and LMM-Optimal (2 Point) plots in Figure 5. Choosing the coefficients this way—using the LMM coefficients extracted from the nearest kinetic simulation to each grid point—results in excellent performance. Only two parameter points exhibit

errors above 20% and the RMS error is only 9% when three training points are used.

The utility of the LMM closure would depend on efficient exploration of parameter space. A database of closure information could be maintained as a parameter space is explored. Expensive kinetic simulations could be performed sparsely, guided by a rigorous statistical framework like Bayesian optimization. Fast fluid simulations, could then select the closure information at the nearest available parameter point.

6 Summary and Conclusions

We have introduced the LMM closure—a new method for formulating closures based on kinetic simulation data. Here it is applied to the problem of phase mixing in a relatively simple turbulent system—gradient driven turbulence in an unsheared slab. The LMM closure utilizes an optimal set of basis vectors from a nonlinear kinetic simulation to formulate closure coefficients in a fluid system. We have compared this method to several closure methods throughout the relevant 2D parameter space (collisionality and gradient drive).

First, we demonstrate that the LMM closure addresses the observation in several recent papers that phase mixing rates are substantially decreased in a turbulent system from the linear expectations. The HP closure is observed to be too dissipative, overestimating nonlinear values. We demonstrate that the LMM closure reproduces phase mixing rates quite accurately.

Comparisons between heat fluxes for four-moment fluid systems and the kinetic system demonstrate that simple truncation performs rather poorly. The HP and HPC closures perform much better than simple truncation. In particular, the HPC closure appears to be highly effective, resulting in relatively low errors and exhibiting no systematic breakdown in parameters space. The LMM closure is also highly effective when using multiple, sparse, training points, producing the lowest errors of all the closures.

Our results suggest that the LMM closure has various advantages as well as drawbacks, which we summarize here. Advantages:

- Capacity to capture the reduced phase mixing rates observed in turbulent systems.
- Accuracy throughout parameter space when trained sparsely.
- Generalizable, in principle, to more comprehensive systems (e.g., that described in Ref. [48]) and other closure terms (e.g. curvature terms).
- Applicable to scenarios where analytic closures have not been formulated or are inaccurate.

Drawbacks:

- Requirement for kinetic training simulations.
- Uncertain extrapolation throughout parameter space.

We expect these drawbacks can be mitigated by statistical frameworks (e.g., Bayesian optimization) designed to track and optimally minimize uncertainties throughout a parameter space.

Various generalizations can be imagined. For example, suitable basis vectors could perhaps be formulated without the need for a nonlinear kinetic simulation by, e.g., taking inspiration from (multiple) linear eigenmodes or conducting a deeper study of nonlinear modifications to the linear eigenmodes.

This work also represents one of the most thorough examinations of HP closures in comparison with nonlinear kinetic simulations. The HP closures were observed to produce phase mixing rates far above the nonlinear kinetic values in this gradient-driven system. Various improvements could be envisioned to mitigate this,for example, retaining some additional number of moments in the system. Notably, despite the discrepancy in phase mixing rates, the heat flux predictions were much more accurate and competitive with the LMM closure. The connection (or lack thereof) between phase mixing rates and heat fluxes is a compelling open question.

The application of closures like the one proposed here to simulations of more comprehensive tokamak or stellarator systems may enable efficient exploration of fusion configurations with the ultimate goal of optimizing fusion performance.

Acknowledgements.– This research used resources of the Texas Advanced Computing Center (TACC) at The University of Texas at Austin. The authors would like to acknowledge Michael Nastac for useful discussions. This work was supported by U.S. DOE Contract No. DE-FG02-04ER54742 and U.S. DOE Office of Fusion Energy Sciences Scientific Discovery through Advanced Computing (SciDAC) program under Award Number DE-SC0018429.

A HP Closure

As introduced in [11] and derived formally in [30], the HP/Smith-style closure coefficients are chosen by matching the approximate response $R^a_{00}(\omega)$ to the exact response $R_{00}(\omega)$, which is given by

$$R_{00}(\omega) = -iZ(\omega) \tag{19}$$

where Z is the plasma dispersion function.

The HP closure for the N^{th} moment, $f_N(\omega)$ takes the form $f_N = \sum_{i=0}^{N-1} A_i f_i(\omega)$. When the first p closure coefficients, A_i, are set to zero, the approximate response in terms of the orthogonal polynomials, $P_i(\omega)$, and conjugate polynomials, $Q_i(\omega)$, becomes

$$\mathrm{R}^{\mathrm{a}}_{00}(\omega) = i\frac{Q_n(\omega) - A_{n-1}Q_{n-1}(\omega) - \cdots - A_pQ_p(\omega)}{P_n(\omega) - A_{n-1}P_{n-1}(\omega) - \cdots - A_pP_p(\omega)} \tag{20}$$

and matches the exact response to order $(O(\omega^{n+1+p}))$. The orthogonal and conjugate polynomials are expressed as

$$\begin{aligned} P_j(\omega) &= P_{j,0} + P_{j,1}\omega + \cdots + P_{j,j}\omega^j \\ &= P^{\omega_0}_{j,0} + P^{\omega_0}_{j,1}(\omega - \omega_0) + \cdots + P_{j,j}(\omega - \omega_0)^j \end{aligned} \tag{21}$$

and

$$\begin{aligned} Q_j(\omega) &= Q_{j,0} + Q_{j,1}\omega + \cdots + Q_{j,j}\omega^j \\ &= Q^{\omega_0}_{j,0} + Q^{\omega_0}_{j,1}(\omega - \omega_0) + \cdots + Q_{j,j}(\omega - \omega_0)^j \end{aligned} \tag{22}$$

This becomes the $[n, q, \omega_0 = 0]$ Pade approximant to the exact response for $q = n - p$ if the q remaining coefficients are chosen to match the first q terms of the Taylor series

$$R_{00}(\omega) = r_0 + r_1\omega + r_2\omega^2 + \cdots \tag{23}$$

Since we are matching in the $\omega \to 0$ limit, we can substitute

$$Z(\zeta) = i\sqrt{\pi}e^{-\zeta^2} - 2\zeta\left(1 - \frac{2\zeta^2}{3} + \frac{4\zeta^4}{15} - \frac{8\zeta^6}{105} + \ldots\right), \tag{24}$$

the Taylor expansion for the plasma dispersion function in the low frequency limit, into 19 to define r_i in 23.

Setting $R^a_{00}(\omega) = R_{00}(\omega)$ results in the following equation

$$\begin{aligned} &\left(r_0 + \ldots + r_{q-1}\omega^{q-1}\right)\left(A_{n-1}P_{n-1}(\omega)\right. \\ &\left. + \ldots + A_pP_p(\omega)\right) \\ &- i\left(A_{n-1}Q_{n-1}(\omega) + \ldots + A_pQ_p(\omega)\right) \\ &= \left(r_0 + \ldots + r_{q-1}\omega^{q-1}\right)P_n(\omega) - iQ_n(\omega) \end{aligned} \tag{25}$$

Choosing the coefficients $A_{n-1}, \ldots, A_p$ so that the coefficients of the different powers of ω on the RHS and LHS match results in the closure.

This matching can be done by solving the following matrix equation

$$\begin{aligned} &\left(\begin{bmatrix} r_0 & 0 & \ldots & 0 \\ r_1 & r_0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ r_{q-1} & r_{q-2} & \ldots & r_0 \end{bmatrix}\right. \\ &\begin{bmatrix} P_{n-1,0} & P_{n-2,0} & \ldots & P_{p,0} \\ P_{n-1,1} & P_{n-2,1} & \ldots & P_{p,1} \\ \vdots & \vdots & \ddots & \vdots \\ P_{n-1,q-1} & P_{n-2,q-1} & \ldots & P_{p,q-1} \end{bmatrix} \\ &\left. -i\begin{bmatrix} Q_{n-1,0} & Q_{n-2,0} & \ldots & Q_{p,0} \\ Q_{n-1,1} & Q_{n-2,1} & \ldots & Q_{p,1} \\ \vdots & \vdots & \ddots & \vdots \\ Q_{n-1,q-1} & Q_{n-2,q-1} & \ldots & Q_{p,q-1} \end{bmatrix}\right) \times \begin{bmatrix} A_{n-1} \\ A_{n-2} \\ \vdots \\ A_p \end{bmatrix} \\ &= \begin{bmatrix} r_0 & 0 & \ldots & 0 \\ r_1 & r_0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ r_{q-1} & r_{q-2} & \ldots & r_0 \end{bmatrix}\begin{bmatrix} P_{n,0} \\ P_{n,1} \\ \vdots \\ P_{n,q-1} \end{bmatrix} - i\begin{bmatrix} Q_{n,0} \\ Q_{n,1} \\ \vdots \\ Q_{n,q-1} \end{bmatrix} \end{aligned} \tag{26}$$

for the closure coefficients $A_n, \ldots, A_p$.

We use $P_n(x) = H_n(x)$ defined by

$$H_n(x) = \frac{(-1)^n e^{x^2}}{\sqrt{2^n n!\sqrt{\pi}}}\frac{d^n}{dx^n}e^{-x^2} \tag{27}$$

satisfying the recurrence relation

$$H_n(x) = \sqrt{\frac{2}{n}}xH_{n-1}(x) - \sqrt{\frac{n-1}{n}}H_{n-2}(x) \tag{28}$$

and conjugate polynomials $Q_n(x)$ constructed by requiring that they satisfy the same recurrence relation as our Hermite polynomials and the conditions

$$Q_0(x) = 0,\ Q_1(x) = \pi^{-1/4}\sqrt{2}. \tag{29}$$

The orthogonal polynomial moments, our Hermite moments, are defined with respect to the Hermite polynomials as

$$f_n = \int_{-\infty}^{\infty} f(x)H_n(x)dx. \tag{30}$$

These are the Hermite polynomials and Hermite decomposition used in our model as described in Sec. 2.

A.1 Verification of HP Coefficients

The HP closure that we focus on closes at the 4^{th} moment with the N,q = 4,2 HP closure. For this closure we have $A_3 = -1.759i$, $A_2 = 0.755$. This gives the following HP closure for our system in DNA

$$f_{\mathbf{k},4} = -1.759i * sgn(k_z) f_{\mathbf{k},3} + 0.755 f_{\mathbf{k},2} \tag{31}$$

The $sgn(kz)$ is present in the $f_{\mathbf{k},3}$ term to ensure damping for both positive and negative k_z. As shown in Fig. 7, this closure produces an excellent match to the response function.

To construct the collisional HP closure we instead need the coefficients for a closure at the 5^{th} moment. We use the N,q = 5,2 closure which takes the form $f_{\mathbf{k},5} = A_4 sgn(k_z) f_{\mathbf{k},4} + A_3 f_{\mathbf{k},3}$ and has coefficients $A_4 = -1.805i$ and $A_3 = 0.801$.

To verify the correctness of our choice of coefficients for the $N, q = 4, 2$ and $N, q = 5, 2$ HP closures, we plot the approximate response functions produced by these closures along with the exact response function in Fig. 7. Both closures show good agreement.

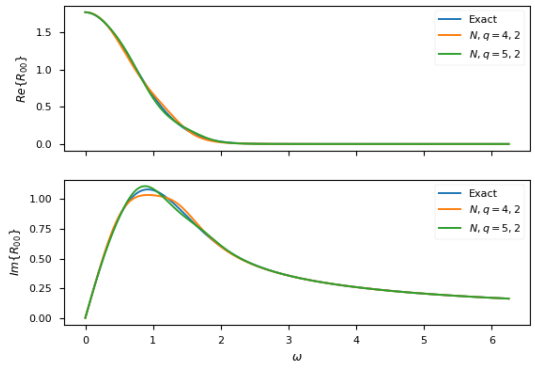

Figure 7: Real and imaginary parts of the exact response function and the approximate response functions for the N,q=4,2 and N,q=5,2 HP closures.

B The HPC Closure

The HP closure is designed for a collisionless regime. However, there is a method for modifying the HP closure to include the effects of collisions developed in [31].

In order to derive the collisional closure for our system, we start by using Smith's method to determine the optimal collisionless closure for $f_{\mathbf{k},5}$ with N,q=5,2 as described above. This yields the collisionless Smith/HP style closure for the 5^{th} moment:

$$f_{\mathbf{k},5} = sgn(k_z) A_4 f_{\mathbf{k},4} + A_3 f_{\mathbf{k},3} \tag{32}$$

with coefficients $A_4 = -1.805i$ and $A_3 = 0.801$. Notice the $sgn(kz)$ in the $f_{\mathbf{k},4}$ term. This is included to ensure damping at both positive and negative k_z.

We then plug this expression for $f_{\mathbf{k},5}$ into the linearized time evolution equation for $f_{\mathbf{k},4}$ including collisions which can be obtained by taking $n = 4$ in Eq. 3 and excluding the nonlinear term. This gives our first expression for $\partial f_4/\partial t$.

$$\begin{aligned} \frac{\partial f_4}{\partial t} &= -4\nu f_{\mathbf{k},4} - \sqrt{5} i k_z f_{\mathbf{k},5} - 2 i k_z f_3 \\ \frac{\partial f_4}{\partial t} &= -4\nu f_{\mathbf{k},4} - \sqrt{5} i k_z (sgn(k_z) A_4 f_{\mathbf{k},4} + A_3 f_{\mathbf{k},3}) - 2 i k_z f_3 \end{aligned} \tag{33}$$

We then take the low frequency limit, $\omega \to 0$, of Eq. 33($\partial f_{\mathbf{k},4}/\partial t = 0$),

$$0 = -4\nu f_{\mathbf{k},4} - \sqrt{5} i k_z (sgn(k_z) A_4 f_{\mathbf{k},4} + A_3 f_{\mathbf{k},3}) - 2 i k_z f_3 \tag{34}$$

and solve the resulting equation for $f_{\mathbf{k},4}$ in terms of $f_{\mathbf{k},3}$.

$$f_{\mathbf{k},4} = \frac{-(2 + \sqrt{5} A_3) i k_z f_{\mathbf{k},3}}{4\nu + \sqrt{5} A_4 sgn(k_z) i k_z} \tag{35}$$

This expression for $f_{\mathbf{k},4}$ is first order accurate in ω for small ω.

We seek an expression for $f_{\mathbf{k},4}$ that is second order accurate in ω, so we take the time derivative of both sides of Eq. 35.

$$\frac{\partial f_{\mathbf{k},4}}{\partial t} = \frac{-(2 + \sqrt{5} A_3) i k_z \partial f_{\mathbf{k},3}/\partial t}{4\nu + \sqrt{5} A_4 sgn(k_z) i k_z} \tag{36}$$

Substituting in $\partial f_{\mathbf{k},3}/\partial t$, which is obtained from Eq. 3, we get our second expression for $\partial f_{\mathbf{k},4}/\partial t$.

$$\frac{\partial f_{\mathbf{k},4}}{\partial t} = \frac{\begin{array}{r}-(2+\sqrt{5}A_3)ik_z(-ik_z\sqrt{3}f_{\mathbf{k},2} \\ -2ik_z f_{\mathbf{k},4} - 3\nu f_{\mathbf{k},3})\end{array}}{4\nu + \sqrt{5}A_4 sgn(k_z)ik_z} \quad (37)$$

Taking the low frequency limit again, we have

$$0 = \frac{\begin{array}{r}-(2+\sqrt{5}A_3)ik_z(-ik_z\sqrt{3}f_{\mathbf{k},2} \\ -2ik_z f_{\mathbf{k},4} - 3\nu f_{\mathbf{k},3})\end{array}}{4\nu + \sqrt{5}A_4 sgn(k_z)ik_z} \quad (38)$$

We then equate the right hand sides of Eq. 34 and Eq. 38 and solve for $f_{\mathbf{k},4}$ to obtain the collisional N,q=4,2 closure.

$$f_{\mathbf{k},4} = \frac{\begin{array}{l}(-(7\sqrt{5}iA_3 + 14i)k_z\nu + (5A_3 + 2\sqrt{5}) \\ A_4 sgn(k_z)k_z^2)f_{\mathbf{k},3} + (k_z^2(15A_3 + 2\sqrt{3}))f_{\mathbf{k},2}\end{array}}{\begin{array}{r}16\nu^2 + 8\sqrt{5}ik_z\nu A_4 sgn(k_z) \\ -k_z^2(5A_4^2 + 2\sqrt{5}A_3 + 4)\end{array}} \quad (39)$$

Plugging in $A_4 = -1.805i$ and $A_3 = 0.801$ which we know from using Smith's method to determine the collisionless 5,2 closure, we get the final numerical expression for our 4,2 collisional closure:

$$f_{\mathbf{k},4} = \frac{-3.051ik_z\nu - 1.759ik_z^2 sgn(k_z)}{1.838\nu^2 + 3.709k_z\nu sgn(k_z) + k_z^2} f_{\mathbf{k},3} + \frac{0.755k_z^2}{1.838\nu^2 + 3.709k_z\nu sgn(k_z) + k_z^2} f_{\mathbf{k},2} \quad (40)$$

This is the HPC closure, a 2^{nd} order accurate (for small ω) closure for $f_{\mathbf{k},4}$ in terms of $f_{\mathbf{k},3}$ and $f_{\mathbf{k},2}$ including collisional effects.

Note that if one takes the collisionless limit, $\nu \to 0$, of Eq. 40, the collisionless closure given in Eq. 31 is recovered.

C Hermite and Fluid Moments

In a fluid or gyrofluid model, the n^{th} moment is calculated as $\int v^n f dv$ where f is the kinetic or gyrokinetic distribution function respectively. The first four moments are physically meaningful: n is density, u is mean velocity, p is pressure, and q is heat flux. Moments beyond q are not physically meaningful but may need to be calculated to ensure the accuracy of the model.

The Hermite moments in our system are calculated as $f_n = \int f(v)H_n(v)dv$ where $f(v)$ is the gyrokinetic distribution function. Since the Hermite polynomials contain only powers of v, the fluid moments are simply linear combinations of the Hermite moments. The relationship between the first 5 fluid moments and the first 5 Hermite moments is shown below.

$$n = \int dv f(v) = \pi^{1/4} f_0 \quad (41)$$

$$u = \int dv f(v)v = \frac{\pi^{1/4}}{\sqrt{2}} f_1 \quad (42)$$

$$p = \int dv f(v)v^2 = \frac{\pi^{1/4}}{\sqrt{2}} f_2 + \frac{\pi^{1/4}}{2} f_0 \quad (43)$$

$$q = \int dv f(v)v^3 = \frac{\sqrt{3}\pi^{1/4}}{2} f_3 + \frac{3\pi^{1/4}}{2\sqrt{2}} f_1 \quad (44)$$

$$r = \int dv f(v)v^4 = \frac{\sqrt{3}\pi^{1/4}}{\sqrt{2}} f_4 + \frac{3\pi^{1/4}}{\sqrt{2}} f_2 + \frac{3\pi^{1/4}}{4} f_0 \quad (45)$$

D Additional Heat Flux plots

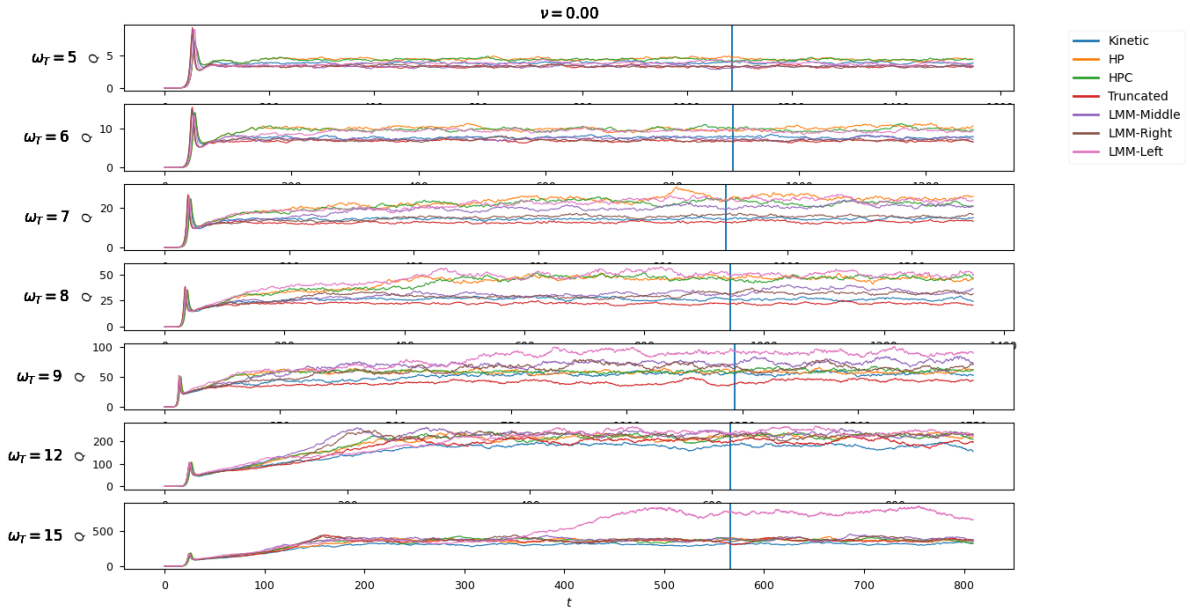

Figure 8: Time traces of the total radial heat flux (Q) for Kinetic (blue), HP (orange), HPC (green), truncated (red), LMM-Middle (purple), LMM-Right (brown), and LMM-Left (pink) simulations for temperature gradient drives (ω_T) ranging from 5 to 15 (increasing downward by panel) and collision frequency $\nu = 0$. The vertical blue lines show the cutoff point - 70% of the simulation time - after which each heat flux curve is averaged to get the final saturation level.

ω_T \ ν	0	0.01	0.05	0.1	0.2	0.5
5	Left	Left	Left	Left	Left	Left
6	Left	Left	Left	Left	Left	Middle
7	Left	Left	Left	Middle	Middle	Middle
8	Left	Left	Middle	Middle	Middle	Right
9	Middle	Middle	Middle	Middle	Middle	Right
12	Middle	Middle	Middle	Middle	Right	Right
15	Middle	Middle	Middle	Middle	Right	Right

Table 1: This table shows which training simulation, equivalently which set of LMM-closure coefficients, was used to produce the LMM-Optimal (3 Point) error plot in Figures 5 and 9. At each grid point in parameter space, the set of coefficient from the training simulation that lies closest to that grid point is used. For example, at $\omega_T, \nu = 5, 0.05$, the LMM-Left closure with coefficients extracted from the $\omega_T, \nu = 6, 0.01$ kinetic simulation is used and at $\omega_T, \nu = 15, 0.5$, the LMM-Right closure with coefficients extracted from the $\omega_T, nu = 12, 0.5$ kinetic simulation is used.

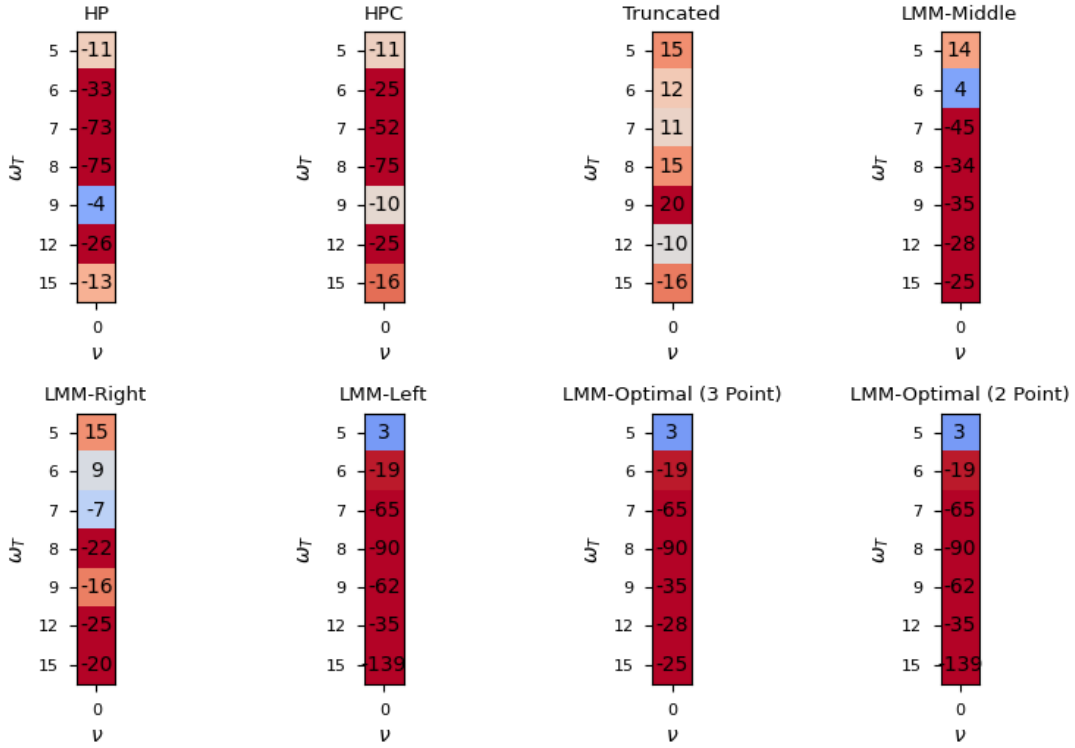

Figure 9: Percent error in saturated heat flux for each closure (HP, HPC, Truncation, LMM-Middle, LMM-Right, LMM-Left, LMM-Optimal (3 Point), LMM-Optimal (2 Point)) as compared against the kinetic simulation for temperature gradient drives (ω_T) ranging from 5 to 15 (increasing downward) and collision frequency $\nu = 0$. Percent errors are calculated as $(Q^{Closed} - Q^{Kinetic})/Q^{Kinetic} \times 100$ where Q^{closed} is calculated by averaging the last 30% of the time trace of the heat flux from the closed simulation and $Q^{Kinetic}$ is calculated by averaging the last 30% of the time trace of the heat flux from the kinetic simulation. Average errors are calculated by averaging the absolute value of the errors.

ω_T \ ν	0	0.01	0.05	0.1	0.2	0.5
5	Left	Left	Left	Left	Left	Left
6	Left	Left	Left	Left	Left	Left
7	Left	Left	Left	Left	Left	Right
8	Left	Left	Left	Left	Right	Right
9	Left	Left	Left	Right	Right	Right
12	Left	Left	Right	Right	Right	Right
15	Left	Right	Right	Right	Right	Right

Table 2: This table shows which training simulation, equivalently which set of LMM-closure coefficients, was used to produce the LMM-Optimal (2 Point) error plot in Figures 5 and 9. At each grid point in parameter space, the set of coefficient from the training simulation that lies closest to that grid point is used. For example, at $\omega_T, \nu = 5, 0.05$, the LMM-Left closure with coefficients extracted from the $\omega_T, \nu = 6, 0.01$ kinetic simulation is used and at $\omega_T, \nu = 15, 0.5$, the LMM-Right closure with coefficients extracted from the $\omega_T, nu = 12, 0.5$ kinetic simulation is used.

References

[1] E. A. Frieman and L. Chen, "Nonlinear gyrokinetic equations for low frequency electromagnetic waves in general plasma equilibria," *The Physics of Fluids*, vol. 25, no. 3, pp. 502–508, 1982.

[2] J. A. Krommes, "The gyrokinetic description of microturbulence in magnetized plasmas," *Annual Review of Fluid Mechanics*, vol. 44, no. 1, pp. 175–201, 2012.

[3] I. G. Abel, G. G. Plunk, E. Wang, M. Barnes, S. C. Cowley, W. Dorland, and A. A. Schekochihin, "Multiscale gyrokinetics for rotating tokamak plasmas: fluctuations, transport and energy flows," *Reports on Progress in Physics*, vol. 76, p. 116201, Nov. 2013.

[4] G. G. PLUNK, S. C. COWLEY, A. A. SCHEKOCHIHIN, and T. TATSUNO, "Two-dimensional gyrokinetic turbulence," *Journal of Fluid Mechanics*, vol. 664, p. 407–435, 2010.

[5] M. J. Pueschel, F. Jenko, D. Told, and J. Büchner, "Gyrokinetic simulations of magnetic reconnection," *Physics of Plasmas*, vol. 18, no. 11, p. 112102, 2011.

[6] D. Told, F. Jenko, J. M. TenBarge, G. G. Howes, and G. W. Hammett, "Multiscale nature of the dissipation range in gyrokinetic simulations of alfvénic turbulence," *Phys. Rev. Lett.*, vol. 115, p. 025003, Jul 2015.

[7] G. G. Howes, S. C. Cowley, W. Dorland, G. W. Hammett, E. Quataert, and A. A. Schekochihin, "A model of turbulence in magnetized plasmas: Implications for the dissipation range in the solar wind," *Journal of Geophysical Research: Space Physics (1978–2012)*, vol. 113, 5 2008.

[8] T. Görler, A. E. White, D. Told, F. Jenko, C. Holland, and T. L. Rhodes, "A flux-matched gyrokinetic analysis of diii-d l-mode turbulence," *Physics of Plasmas*, vol. 21, no. 12, p. 122307, 2014.

[9] D. Hatch, M. Kotschenreuther, S. Mahajan, P. Valanju, F. Jenko, D. Told, T. Görler, and S. Saarelma, "Microtearing turbulence limiting the JET-ILW pedestal," *Nuclear Fusion*, vol. 56, p. 104003, aug 2016.

[10] C. Holland, "Validation metrics for turbulent plasma transport," *Physics of Plasmas*, vol. 23, no. 6, p. 060901, 2016.

[11] G. W. Hammett and F. W. Perkins, "Fluid moment models for landau damping with application to the ion-temperature-gradient instability," *Phys. Rev. Lett.*, vol. 64, pp. 3019–3022, Jun 1990.

[12] W. Dorland and G. W. Hammett, "Gyrofluid turbulence models with kinetic effects," *Physics of Fluids B: Plasma Physics*, vol. 5, no. 3, pp. 812–835, 1993.

[13] G. M. Staebler, J. E. Kinsey, and R. E. Waltz, "Gyro-landau fluid equations for trapped and passing particles," *Physics of Plasmas*, vol. 12, no. 10, p. 102508, 2005.

[14] G. M. Staebler, J. E. Kinsey, and R. E. Waltz, "A theory-based transport model with comprehensive physics," *Physics of Plasmas*, vol. 14, no. 5, p. 055909, 2007.

[15] B. D. Scott, "Tokamak edge turbulence: background theory and computation," *Plasma Physics and Controlled Fusion*, vol. 49, pp. S25–S41, jun 2007.

[16] X. Q. Xu, P. W. Xi, A. Dimits, I. Joseph, M. V. Umansky, T. Y. Xia, B. Gui, S. S. Kim, G. Y. Park, T. Rhee, H. Jhang, P. H. Diamond, B. Dudson, and P. B. Snyder, "Gyro-fluid and two-fluid theory and simulations of edge-localized-modes," *Physics of Plasmas*, vol. 20, no. 5, p. 056113, 2013.

[17] J. Peer, A. Kendl, T. Ribeiro, and B. Scott, "Gyrofluid computation of magnetic perturbation effects on turbulence and edge localized bursts," *Nuclear Fusion*, vol. 57, p. 086026, jul 2017.

[18] P. Hunana, M. L. Goldstein, T. Passot, P. L. Sulem, D. Laveder, and G. P. Zank, "Properties of kinetic alfvén waves: A comparison of fluid models with kinetic theory," *AIP Conference Proceedings*, vol. 1539, no. 1, pp. 179–182, 2013.

[19] P. L. Sulem and T. Passot, "Landau fluid closures with nonlinear large-scale finite larmor radius corrections for collisionless plasmas," *Journal of Plasma Physics*, vol. 81, no. 1, p. 325810103, 2015.

[20] G. G. Plunk, "Landau damping in a turbulent setting," *Physics of Plasmas*, vol. 20, no. 3, p. 032304, 2013.

[21] A. Kanekar, A. A. Schekochihin, W. Dorland, and N. F. Loureiro, "Fluctuation-dissipation relations for a plasma-kinetic langevin equation," *Journal of Plasma Physics*, vol. 81, no. 1, p. 305810104, 2015.

[22] J. T. Parker, E. G. Highcock, A. A. Schekochihin, and P. J. Dellar, "Suppression of phase mixing in drift-kinetic plasma turbulence," *Physics of Plasmas*, vol. 23, no. 7, p. 070703, 2016.

[23] D. R. Hatch, F. Jenko, A. B. Navarro, V. Bratanov, P. W. Terry, and M. J. Pueschel, "Linear signatures in nonlinear gyrokinetics: interpreting turbulence with pseudospectra," *New Journal of Physics*, vol. 18, p. 075018, jul 2016.

[24] A. A. Schekochihin, J. T. Parker, E. G. Highcock, P. J. Dellar, W. Dorland, and G. W. Hammett, "Phase mixing versus nonlinear advection in drift-kinetic plasma turbulence," *Journal of Plasma Physics*, vol. 82, Apr. 2016.

[25] R. Meyrand, A. Kanekar, W. Dorland, and A. A. Schekochihin, "Fluidization of collisionless plasma turbulence," *Proceedings of the National Academy of Sciences*, vol. 116, no. 4, pp. 1185–1194, 2019.

[26] D. R. Hatch, F. Jenko, A. Bañón Navarro, and V. Bratanov, "Transition between saturation regimes of gyrokinetic turbulence," *Phys. Rev. Lett.*, vol. 111, p. 175001, Oct 2013.

[27] D. R. Hatch, F. Jenko, V. Bratanov, and A. B. Navarro, "Phase space scales of free energy dissipation in gradient-driven gyrokinetic turbulence," *Journal of Plasma Physics*, vol. 80, no. 4, pp. 531–551, 2014.

[28] S. I. Braginskii, "Transport Processes in a Plasma," *Reviews of Plasma Physics*, vol. 1, p. 205, Jan. 1965.

[29] N. F. Loureiro, A. A. Schekochihin, and A. Zocco, "Fast collisionless reconnection and electron heating in strongly magnetized plasmas," *Phys. Rev. Lett.*, vol. 111, p. 025002, Jul 2013.

[30] S. A. Smith, *Dissipative closures for statistical moments, fluid moments, and subgrid scales in plasma turbulence*. PhD thesis, PRINCETON UNIVERSITY, Nov. 1997.

[31] P. B. Snyder, G. W. Hammett, and W. Dorland, "Landau fluid models of collisionless magnetohydrodynamics," *Physics of Plasmas*, vol. 4, no. 11, pp. 3974–3985, 1997.

[32] L. SIROVICH, "Turbulence and the dynamics of coherent structures part i: Coherent structures," *Quarterly of Applied Mathematics*, vol. 45, no. 3, pp. 561–571, 1987.

[33] G. Berkooz, P. Holmes, and J. L. Lumley, "The proper orthogonal decomposition in the analysis of turbulent flows," *Annual Review of Fluid Mechanics*, vol. 25, no. 1, pp. 539–575, 1993.

[34] A. P. G. Rozza, D. Huynh, "Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations," *Archives of Computational Methods in Engineering*, vol. 15, p. 229–275, 2008.

[35] R. F. P. Feldmann, "Efficient linear circuit analysis by pade approximation via the lanczos process," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 14, p. 639– 649, 1995.

[36] R. W. Freund, "Model reduction methods based on krylov subspaces," *Acta Numerica*, vol. 12, p. 267–319, 2003.

[37] B. Peherstorfer and K. Willcox, "Data-driven operator inference for nonintrusive projection-based model reduction," *Computer Methods in Applied Mechanics and Engineering*, vol. 306, pp. 196–215, 2016.

[38] H. Sugama, T.-H. Watanabe, and W. Horton, "Collisionless kinetic-fluid closure and its application to the three-mode ion temperature gradient driven system," *Physics of Plasmas*, vol. 8, no. 6, pp. 2617–2628, 2001.

[39] H. Sugama, T.-H. Watanabe, and W. Horton, "Comparison between kinetic and fluid simulations of slab ion temperature gradient driven turbulence," *Physics of Plasmas*, vol. 10, no. 3, pp. 726–736, 2003.

[40] P. W. Terry, D. A. Baver, and S. Gupta, "Role of stable eigenmodes in saturated local plasma turbulence," *Physics of Plasmas*, vol. 13, no. 2, p. 022307, 2006.

[41] D. R. Hatch, P. W. Terry, F. Jenko, F. Merz, and W. M. Nevins, "Saturation of gyrokinetic turbulence through damped eigenmodes," *Phys. Rev. Lett.*, vol. 106, p. 115003, Mar 2011.

[42] D. R. Hatch, P. W. Terry, F. Jenko, F. Merz, M. J. Pueschel, W. M. Nevins, and E. Wang, "Role of subdominant stable modes in plasma microturbulence," *Physics of Plasmas*, vol. 18, no. 5, p. 055706, 2011.

[43] G. G. Whelan, M. J. Pueschel, and P. W. Terry, "Nonlinear electromagnetic stabilization of plasma microturbulence," *Phys. Rev. Lett.*, vol. 120, p. 175002, Apr 2018.

[44] G. H. Golub and C. F. Van Loan, *Matrix computations*, vol. 3. JHU press, 2013.

[45] G. Berkooz, P. Holmes, and J. L. Lumley, "The proper orthogonal decomposition in the analysis of turbulent flows," *Annual Review of Fluid Mechanics*, vol. 25, no. 1, pp. 539–575, 1993.

[46] N. Christen, M. Barnes, M. R. Hardman, and A. A. Schekochihin, "Bistable turbulence in strongly magnetised plasmas with a sheared mean flow," 2021.

[47] R. A. Heinonen and P. H. Diamond, "A closer look at turbulence spreading: How bistability admits intermittent, propagating turbulence fronts," *Physics of Plasmas*, vol. 27, no. 3, p. 032303, 2020.

[48] N. R. Mandell, W. Dorland, and M. Landreman, "Laguerre–hermite pseudo-spectral velocity formulation of gyrokinetics," *Journal of Plasma Physics*, vol. 84, no. 1, p. 905840108, 2018.