

Multiscale model reduction for unsteady fluid flow

Jared Callaham

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2022

Reading Committee:

Steven Brunton, Chair

J. Nathan Kutz

Duane Storti

Program Authorized to Offer Degree:

Mechanical Engineering

©Copyright 2022
Jared Callaham

University of Washington

Abstract

Multiscale model reduction for unsteady fluid flow

Jared Callaham

Chair of the Supervisory Committee:
Professor Steven Brunton
Mechanical Engineering

This dissertation develops methods for constructing simplified models of unsteady fluid flows in regimes ranging from weakly nonlinear to fully turbulent. These models can provide valuable insights into the flow physics, as well as inexpensive surrogate models suitable for analytic study and controller design. The emphasis is on extending traditional methods using recent advances in data-driven modeling in a manner that preserves the interpretability and robustness of classical analysis. Throughout, the proposed methodological developments are critically evaluated against extensive computational fluid dynamics simulations and experimental wind tunnel observations representing a variety of fundamental features of unsteady flows.

This work takes three distinct approaches to model reduction. First, a perspective of the fluid flow as a high-dimensional, dissipative dynamical system with emergent large-scale coherence leads to approximations in terms of low-dimensional nonlinear dynamics. These models can be derived by projection of the governing equations or sparse model discovery; in either case it is crucial to systematically account for the influence of unresolved degrees of freedom. Alternatively, in fully-developed turbulence the evolution of global integral quantities can be viewed as deterministic motion forced by incoherent fluctuations. The analogy with statistical mechanics cannot be made rigorous

for turbulence, but an empirical method is developed to approximate these generalized Brownian motions from limited experimental data. Finally, the observation that the behavior of physical systems is often determined by a dominant balance between a small subset of physical mechanisms motivates the development of an algorithm for objective identification of regions with different active physics. Underlying all of these frameworks is a unifying perspective of the flow as a system with complex nonlinear interactions across a wide range of spatiotemporal scales.

TABLE OF CONTENTS

	Page
List of Figures	iv
Part I: Introduction	1
1.1 Survey of model reduction and related topics	3
1.2 Organization and contributions	8
Part II: Background	11
Chapter 2: Mathematical background	12
2.1 Governing equations for incompressible fluid flow	12
2.2 Motivation for low-dimensional modeling	15
2.3 Fluid flow as a nonlinear dynamical system	22
2.4 Generators of deterministic and stochastic processes	27
2.5 Summary	30
Chapter 3: Modal analysis	32
3.1 Linear stability analysis	33
3.2 Proper orthogonal decomposition (POD)	37
3.3 Dynamic mode decomposition (DMD)	39
3.4 Summary	41
Chapter 4: Reduced-order modeling	43
4.1 Galerkin projection	43
4.2 Sparse identification of nonlinear dynamics	49
4.3 Summary	52
Chapter 5: Case study: the cylinder wake	53
5.1 Flow configuration	53
5.2 Modal analysis	55

5.3	POD-Galerkin model	57
5.4	Invariant manifold model	59
5.5	SINDy model	62
5.6	Summary	63
Part III:	Nonlinear dynamics	65
Chapter 6:	Empirical invariant manifold reduction	66
6.1	Advection and nonlinear correlation	66
6.2	Shear-driven open cavity flow	71
6.3	Reduced-order modeling	74
6.4	Nonlinear correlations analysis	78
6.5	Polynomial manifold approximation	84
6.6	Deep manifold models and the route to chaos	89
6.7	Discussion	95
Chapter 7:	Multiscale closure models via stochastic averaging	98
7.1	Multiscale closure modeling	102
7.2	Prototypical fast effects	114
7.3	Vortex shedding in the cylinder wake	118
7.4	Convective instability and vortex pairing in the mixing layer	119
7.5	Discussion	129
Part IV:	Stochastic system identification	132
Chapter 8:	Stochastic modeling with Langevin regression	133
8.1	Background on stochastic modeling	138
8.2	Proposed Langevin regression method	143
8.3	Results	149
8.4	Discussion	160
Chapter 9:	A mean-field model of the axisymmetric wake	163
9.1	Symmetry-breaking in the bluff body wake	167
9.2	Mean-field theory of symmetry-breaking transitions	168
9.3	Parametric modal expansion	169

9.4	Reduced-order model of symmetry-breaking	176
9.5	Discussion	181
Part V:	Dominant balance	183
Chapter 10:	Data-driven dominant balance modeling	184
10.1	Unsupervised dominant balance identification	187
10.2	The unsteady cylinder wake	190
10.3	A turbulent boundary layer	196
10.4	Discussion	200
Part VI:	Outlook	202

LIST OF FIGURES

Figure Number	Page
<p>1.1 Flow configurations modeled in this thesis. The flows range span more than three orders of magnitude in Reynolds number from DNS of the laminar, two dimensional flow past a cylinder at $Re = 100$ (Chapter 5) to experimental measurements of a three-dimensional fully turbulent bluff body wake at $Re \approx 2 \times 10^5$ (Chapter 9). The shear-driven cavity (Chapter 6) and mixing layer (Chapter 7) are of intermediate dynamical complexity. . . .</p>	2
<p>1.2 Vizualization of the Ruelle-Takens route to chaos. In this scenario the steady flow first loses stability to a periodic limit cycle in a supercritical Hopf bifurcation. A secondary limit cycles leads to quasiperiodic dynamics on a 2-torus. A third bifurcation typically precipitates a breakdown into chaos. The large figure on the left is constructed by plotting the leading three proper orthogonal decomposition coefficients computed from a global analysis of the shear-driven cavity flow presented in Chapter 6 across different Reynolds numbers. The trajectories on the right apply Eq. (6.17) to the leading two pairs of oscillatory modes.</p>	4
<p>2.1 The Richardson-Kolmogorov turbulent energy cascade. For isotropic turbulence, the Fourier space dynamics suggests that energy is transferred from the largest integral scales to the smallest dissipative scales by the nonlinear triadic interactions. Scaling analysis suggests that the energy content should decay with wavenumber like $E \sim k^{-5/3}$.</p>	13
<p>2.2 Rayleigh-Bénard flow The fluid between differentially heated plates conducts heat upwards with zero fluid velocity at low values of the control parameter (Rayleigh number). Above a critical value the steady conducting state is unstable to perturbations, eventually reaching a new nonlinear equilibrium. This flow is a prototypical example of instability and pattern formation.</p>	16
<p>2.3 Stability of the Swift-Hohenberg equation. Above the critical value r_c, the steady solution is unstable to perturbations near the critical wavenumber k_c with maximum growth rate $\epsilon = r - r_c$.</p>	17

2.4	The Swift-Hohenberg model. The approximate amplitude equation (2.20) predicts the qualitative symmetry-breaking and nonlinear stability of the model above the threshold of instability. The numerical simulation was performed with a pseudospectral Fourier method and a second-order Crank-Nicolson/Adams-Bashforth implicit-explicit timestepping scheme.	21
3.1	Types of linear instability in a model of shear flow While disturbances decay exponentially in absolutely stable flow (a), the distinction between convective and absolute instability is a matter of Galilean reference frame. Perturbations are amplified as they travel downstream in convective instability (b), while grow at a fixed point in space for absolute instability (c). The plots are adapted from Bagheri <i>et al.</i> (2009b) and show a modified linear Ginzburg-Landau equation.	35
4.1	Model discovery applied to the Lorenz system. The SINDy algorithm assumes that unknown governing equations for a dynamical system can be approximated by a sparse linear combination of candidate functions. The coefficient vectors can then be identified by solving a sparse regression problem. Note that the matrices Θ and Ξ are the transpose of the usual SINDy formulation. Figure adapted from Brunton <i>et al.</i> (2016a).	50
5.1	Two-dimensional flow past a cylinder. The simulation domain (top left) is the same as Noack <i>et al.</i> (2003). The post-transient flow is characterized by periodic vortex shedding (top right). This limit cycle is stabilized by Reynolds stress-mediated base flow deformation, which can be captured in POD-Galerkin models by augmenting the POD basis with a “shift mode” ψ_Δ capturing the deformation of the base flow (middle). The standard POD modes $\psi_{1,2,\dots,8}$ appear in pairs (bottom); phase-lagged oscillations allow the fixed spatial structures to capture advection of the vortex street.	54
5.2	Transient evolution of the leading POD coefficients. The pair a_1, a_2 describe the periodic vortex shedding, which begins as a linear instability with exponential energy growth. Eventually the base flow deformation (parameterized by the shift mode a_Δ) grows to balance the perturbations, stabilizing the limit cycle. The self-exciting, self-limiting oscillation is characteristic of dynamics above a supercritical Hopf bifurcation.	55
5.3	POD coefficients for the post-transient $Re = 100$ cylinder wake. The POD coefficients are guaranteed to be linearly uncorrelated, but the phase portraits (a) show clear evidence of <i>nonlinear</i> correlation. The Lissajous figures are characteristic of pure harmonics of the fundamental pair (a_1, a_2) . The modes appear in pairs with approximately the same energy (b), typical of oscillations or traveling waves in the flow.	56

5.4	Low-dimensional models of the transient cylinder wake.	A model based on the post-transient limit cycle (grey) is structurally unstable without the addition of a shift mode capturing the base flow deformation (grey-blue). An analytic invariant manifold reduction (blue) is a more accurate approximation, although the lack of resolution for higher harmonics results in an overestimate of post-transient energy. This is corrected by the sparse Galerkin (SINDy) model (orange), which has the same structure as the invariant manifold model and can capture higher harmonics (also see Fig. 5.6). The nine degrees of freedom required for kinematic resolution approximately evolve on a two-dimensional parabolic manifold (b). Even when stable, POD-Galerkin models are not confined to the inertial manifold (left). The invariant manifold and SINDy models restrict the dynamics to the manifold (middle, right).	58
5.5	Linear stability of the cylinder flow.	Both the linearized Navier-Stokes and POD-Galerkin systems have one pair of unstable eigenvalues with growth rate $\sigma > 0$. The Galerkin model is based on an expansion around the marginally stable mean flow, so it tends to underestimate the growth rate relative to the true Navier-Stokes equations linearized around the unstable steady state. This can also be seen in the relatively slow growth of perturbations in the Galerkin models shown in Fig. 5.4a.	60
5.6	Resolving higher harmonics.	The analytic invariant manifold reduction computes a quadratic approximation to the manifold, resolving the second harmonic a_3, a_4 . However, a cubic SINDy model with a fourth-order sparse polynomial manifold equation can capture the higher harmonics. The resulting model achieves significantly better accuracy (Fig. 5.4).	63
6.1	Schematic of the model reduction approach exploiting nonlinear correlations.	The flow fields are first projected onto a linear modal basis Ψ , yielding modal coefficients $a(t)$. The quasiperiodic dynamics can be described by four degrees of freedom; the rest of the modal coefficients can then be reconstructed with polynomial functions consistent with triadic interactions in the frequency domain. The dynamics of the active degrees of freedom can be modeled either by restricting the POD-Galerkin dynamics to the toroidal manifold or by identifying a simple, interpretable dynamical system with the SINDy algorithm.	67

6.2	Linear advection equation with errors $\epsilon_n \sim \mathcal{N}(0, \epsilon^2)$ in the dispersion relation $\omega_n = ck_n$. The Galerkin model (grey) loses coherence with the exact solution (black) over a time scale $1/\epsilon$. If the polynomial correlations implied by the dispersion relation are enforced explicitly, the model is robust to such errors. Nonlinear correlation in the true system, given by (6.5), appears in the Lissajous-type phase portraits of the Fourier coefficients (right). Similar behavior manifests in Galerkin models of nonlinear advection-dominated flows.	69
6.3	Open cavity flow configuration Computational domain and representative instantaneous vorticity field for the shear-driven cavity flow at $Re = 7500$ highlighting the vortical structures developing along the shear layer.	72
6.4	Fourier spectrum $\hat{E}(\omega)$ of the fluctuation's kinetic energy at $Re = 7500$. The high-frequency peak ($\omega_s \simeq 12$) corresponds to the shear layer instability while the low frequency peak ($\omega_c \simeq 3$) are associated to the inner-cavity dynamics. A few other peaks have been labelled based on the quadratic interactions on the two fundamental frequencies for the sake of illustration. Multiple closely-spaced peaks are associated with nearby frequency combinations (e.g. $2\omega_c \approx \omega_s - 2\omega_c \approx 6$). Also shown are the real parts of the DMD modes at ω_s and ω_c	73
6.5	Eigenspectrum of the Navier-Stokes operator at $Re = 7500$. The spectrum is estimated in three ways: linearized in the vicinity of the base flow (BF, black circles), mean flow (MF, blue circles), and from dynamic mode decomposition (DMD, red crosses). For the linear stability analyses, only eigenvalues for which a 10^{-8} convergence has been achieved are plotted. These eigenspectra have been computed using a time-stepper Arnoldi algorithm with a sampling period $\Delta T = 0.1$ and a Krylov subspace dimension of 1024 and 512 for the base and mean flows, respectively.	74
6.6	Singular value spectrum of the quasiperiodic cavity flow. Black dots represent the normalized squared singular values of the snapshot correlation matrix, indicating the fraction of fluctuation kinetic energy resolved by each mode. Red crosses indicate the fraction of residual energy, or normalized cumulative sum of squared singular values. Dashed lines indicate the number of modes retained ($r = 64$).	76
6.7	Harmonic modes identified from POD and DMD analysis. The spatial fields and phase portraits both indicate that certain mode pairs are harmonics arising from the description of wavelike motion in the shear layer and inner cavity. Because DMD is based on both spatial and temporal correlation, this structure is especially pronounced in the DMD coefficients. The vorticity plots are real parts of the DMD modes, but analogous modes exist in the POD basis.	76

6.8	DMD frequencies ω_k and average energy $E_k = \langle \alpha_k ^2 \rangle$ along with vorticity plots for the real part of the most energetic modes. The second mode pair ($k = 3, 4$) is a harmonic of the leading pair ($k = 1, 2$), while the third pair ($k = 5, 6$) represents the low-frequency inner cavity motion. Other modes (e.g. $k = 7, 8$) are either harmonics or indicate nonlinear frequency crosstalk between these leading modes, as in figure 6.4.	77
6.9	POD-Galerkin models for open cavity flow. The time series shows the evolution of the fluctuation kinetic energy predicted by POD-Galerkin reduced-order models of various dimensions along with DNS values. Though all values of r shown here capture sufficient dissipation to remain at finite energy, none resolves the true quasiperiodic dynamics. . . .	78
6.10	Nonlinear correlations in modal coefficients. Subplots show phase portraits of DMD coefficients along with measures of linear and nonlinear correlation (Pearson's ρ and the randomized dependence coefficient, respectively). While α_1 and α_3 are linearly uncorrelated ($\rho = 0$), the clear functional relationship between the two is reflected in the RDC value; physically, α_3 is a pure harmonic of α_1 . On the other hand, α_{17} corresponds to a nonlinear crosstalk mode that has no clear correlation with α_1 , either linear or nonlinear. Nevertheless, it can be accurately approximated by a simple polynomial function of the active degrees of freedom, as shown by Table 6.2 and Fig. 6.12.	80
6.11	Identification of the active degrees of freedom with the randomized dependence coefficient (RDC). The lower triangular portion of the figure shows phase portraits of the real (horizontal axes) and imaginary (vertical axes) DMD coefficients, while the upper triangular portion depicts the RDC values scaled linearly in color and radius. Two approximately independent clusters can be identified: the shear layer dynamics (blue) and inner cavity oscillations (red). Each of these is associated with a dominant mode pair (solid borders) and pure harmonics (dashed borders) that are strongly nonlinearly correlated with the dominant modes. The other modes also have simple polynomial relationships with the active degrees of freedom but include cross terms that break the one-to-one nonlinear correlation (see Table 6.2 and Figs. 6.10 and 6.12).	81
6.12	Example coefficient reconstructions $\alpha \approx \varphi(\hat{\alpha})$ based on the leading DMD coefficients (—). The sparse polynomial approximation (- -) for higher-order modes with pure frequency content (e.g. $\alpha_{17} \approx \varphi_{17}(\hat{\alpha})$) tends to be more accurate than for modes with mixed content (α_{27}).	86
6.13	Evolution of the fluctuation kinetic energy for the reduced-order models compared to DNS. By accounting for nonlinear correlations, both the manifold Galerkin and SINDy models remain at the correct energy level at long times, despite having many fewer degrees of freedom than the standard Galerkin model (top). Similarly, both models resolve the nonlinear interactions leading to the discrete peaks in the power spectrum (bottom). . . .	86

6.14	Phase portraits for the POD coefficients reconstructed from the manifold Galerkin and SINDy reduced-order models (left). Both models accurately capture the shear layer instability and its harmonics (e.g. $\alpha_1, \alpha_3, \alpha_{13}, \alpha_{23}$), though the manifold Galerkin model tends to underestimate the amplitude of the inner cavity motions (e.g. $\alpha_5, \alpha_{22}, \alpha_{50}$). A Poincarè section of the toroidal attractor confirms this discrepancy, but shows clearly that both models are quasiperiodic and remain on the approximate attractor.	88
6.15	Approximation strategies of the slow manifold. For small-amplitude perturbations relatively close to a bifurcation, the classical center manifold approximation is a good choice (left, Chapter 5). For phase-locked or quasiperiodic dynamics, a SINDy-type sparse polynomial approximation to the manifold equation is a straightforward, interpretable alternative (middle, Sec. 6.5). When the flow exhibits broadband chaos, these methods will fail and a deep autoencoder offers a more flexible class of function approximation (right, Sec. 6.6).	90
6.16	Deep manifold models of shear-driven cavity flow. The flow at $Re = 7500$ is quasiperiodic (top), with approximately discrete peaks in the power spectrum, but has transitioned to chaos by $Re = 10^4$ (bottom). Both the phase portraits in the latent space (left) and the power spectra for the reconstructed POD coefficients (right) show qualitative agreement. The quasiperiodic power spectrum is estimated with the FFT, while the chaotic case is estimated with Welch’s method.	93
6.17	Bifurcation diagram for the parametric deep manifold model. The model was trained on data from $Re = 7500$ and 10000 and evaluated on $4500 \leq Re \leq 11000$. The points shown here are a Poincarè section of the reconstructed coefficients corresponding to inner cavity oscillations (see text for details). The appearance of a loop near $Re \simeq 7000$ indicates the onset of quasiperiodicity (dashed lines on $x - y$ plane), while the breakdown into a cloud of points shows chaos (shaded section on $x - y$ plane). Selected trajectories in the latent space show a limit cycle, 2-torus, and chaotic motion.	94
7.1	Illustration of multiscale closure model applied to a mixing layer. The visualization of the network of average energy transfer between the leading harmonic modes shows the cascade of energy to higher-order modes. The multiscale model reduction (MMR) method approximates the effects of unresolved higher-order modes via stochastic averaging, which leads to a generalized Stuart-Landau-type equation with cubic nonlinear interactions. Note that the while the network visualization of energy transfers (left) is computed from the modal coefficients and Galerkin model for the leading harmonics, that of of the multiscale approximation (right) is a notional illustration of the origin of the cubic terms in Eq. (7.24b). See Sec. 7.4 for details on the construction of this figure and the low-dimensional model of the mixing layer.	99

7.2	Schematic of nonlinear interactions in the multiscale closure scheme. The dynamics of one variable (a_1) in a system with two slow variables involve quadratic interactions between fast and slow variables that would be neglected in a standard truncation. Instead, the proposed method averages over the fast scales, ultimately generating effective cubic nonlinearities in the closed equations.	102
7.3	Reduced-order models of the cylinder wake. The standard 9-mode Galerkin model accurately estimates the stable limit cycle (a), but the transient dynamics deviate from the slow manifold, leading to an energy overshoot (b). Both the two-dimensional invariant manifold reduction in Sec. 5.4 and MMR closure models prevent this by eliminating the rapidly-equilibrating variable associated with mean flow deformation, but MMR more accurately estimates the limit cycle amplitude. The full flow field can be reconstructed by approximating the harmonics with polynomial regression (c).	118
7.4	An incompressible mixing layer at $Re = 500$. The hyperbolic tangent base flow is convectively unstable, amplifying small perturbations as they are carried downstream. The primary shear layer instability forms Kelvin-Helmholtz waves that roll up into vortices, which in turn undergo successive vortex pairing. We force the flow at the inlet with eigenfunctions of inviscid flow equations linearized around the base flow (left). The central part of the domain is used for further modeling to avoid boundary effects, while the downstream extent $x \in (250, 300)$ is strongly damped to prevent numerical instability at the outlet.	120
7.5	Proper orthogonal decomposition (POD) applied to the mixing layer. Modes computed on both the short and long domains reveal modes related to the dominant flow features: the shear layer instability and the downstream vortex pairing. Although the vortex pairing is secondary to the upstream instability, on the longer domain it accounts for a larger proportion of the fluctuation kinetic energy. Higher-order modes not pictured here include harmonics, nonlinear crosstalk, and modes related to irregularity in the location of the vortex pairing events.	123
7.6	Standard reduced-order models of the mixing layer. The plots show the fluctuating kinetic energy in the first 16 modes of POD-Galerkin (top) and FTT closure models (bottom) for the mixing layer on the longer domain. Although the FTT closure model does stabilize the energy at a much lower level than the unmodified Galerkin system, it does not reproduce the natural dynamics of the flow over long time horizons.	125
7.7	Models of the mixing layer on a short domain. The phase portraits of the mode pairs (left) show phase-locking between the fundamental mode pair (a_1, a_2), the subharmonic vortex pairing (a_3, a_4), and the first harmonic (a_5, a_6). These phase relationships are preserved by the MMR closure, resulting in physically consistent flow field estimates (right).	126

7.8	Models of the mixing layer on the long domain. While standard POD-Galerkin models are energetically unstable until at least $r = 64$, the multiscale closure stabilizes the model with only 16 variables (left). The model also remains coherent on long integration times (see also Fig. 7.9), producing flow field predictions that are consistent with the large-scale structure of the DNS (right).	127
7.9	Phase portraits of the mixing layer model on the long domain. The MMR closure (orange) closely matches the DNS trajectories (black) for the most energetic modes, even though the structure is more complicated than the Lissajous-type harmonics on the short domain (Fig. 7.7). This figure does not show a Galerkin model for comparison because they are energetically unstable (Fig. 7.8).	128
8.1	Schematic of Langevin regression (top) with example applications (bottom). Given a long time series of a macroscopic variable describing a complex system, we seek to identify an approximate stochastic model. The variable x might represent a reaction coordinate capturing metastable protein configurations or the temporal coefficient of a dominant global hydrodynamic mode, for instance. Langevin regression uses both the forward and adjoint Fokker-Planck operators to optimize free parameters ξ of the model, ensuring consistency with observed statistics such as the finite-time Kramers-Moyal coefficients $m_\tau^{(n)}(x)$ (see Sec. 8.1).	134
8.2	Dual scale separation for stochastic modeling. Even the fastest scales of continuous physical systems are characterized by some finite decorrelation rate α (e.g. Eq. (8.7)). However, if the macroscopic dynamics have a much slower characteristic timescale $\omega \ll \alpha$, we may be able to choose a sampling rate τ^{-1} which can simultaneously resolve the dominant dynamics and allow the unresolved scales to decorrelate. For example, the power spectrum of the radial center of pressure of the turbulent wake is shown at bottom along with the subsampling rate used in Sec. 8.3.	137
8.3	Schematic of Kramers-Moyal coefficient estimation for the first moment (drift). The drift estimate is determined by the conditional mean of the state evolution over the interval, while the diffusion is given by the conditional variance. The conditional moments can be approximated by dividing a long time series into histogram bins and taking the mean and variance within each bin. For example, the Kramers-Moyal drift estimate gives an approximate discretized vector field for the deterministic component of the dynamics (right).	140
8.4	PDF evolution from the Fokker-Planck operator \mathcal{L}. The distribution used to evaluate the conditional finite-time moments (8.4) can be interpreted as the evolution of a delta function initial condition over the sampling interval τ , where the state is known to be x' at time t .	142

8.5	Schematic of the Langevin regression optimization problem. The Fokker-Planck equation can be used to compare the conditional moments and distribution for a proposed model to those observed empirically. The model parameters are chosen to minimize the discrepancy between the model and observations. The form of the model may be simultaneously identified with the model selection procedure outlined in Sec. 8.2	144
8.6	Pitchfork normal form forced by time-correlated (colored) noise. Top: the noise induces random switching between the metastable equilibria. Bottom: if the Kramers-Moyal coefficients are computed with a sampling rate faster than the decorrelation of the noise (\circ), the drift still appears cubic but the amplitude is underestimated by approximately an order of magnitude. On the other hand, if the noise is allowed to decorrelate (Δ), the estimated Kramers-Moyal coefficients are the right order of magnitude, but are distorted from the zero-time value. The diffusion appears multiplicative and quadratic. The adjoint finite-time corrections recover a consistent Langevin model driven by white noise ($---$).	150
8.7	Particle in a one-dimensional double-well potential. Even when the system is far from the pitchfork bifurcation ($\epsilon = 20$), the dynamics are dominated by bistable switching behavior (top). Langevin regression identified a reduced-order approximation to this system (middle), which is consistent with both the state probability distribution (bottom left) and the distribution of metastable dwell times (bottom right).	153
8.8	Finite-time effects in Kramers-Moyal average. Non-Markovian effects from order reduction cause the Kramers-Moyal average to fail at high sampling rates (\circ). The forcing appears uncorrelated when the system is subsampled (Δ). Langevin regression identifies a model ($---$) which is consistent with the Kramers-Moyal average at finite sampling rate ($-$).	155
8.9	Comparison of models far from bifurcation with Monte Carlo simulation. Both the normal form (NF) and the PDF fit (Rigas <i>et al.</i> , 2015) quickly fail to match the metastable dwell time, while the Langevin regression (LR) model continues to be accurate far from the bifurcation. The Kullback-Leibler divergence \mathcal{D}_{KL} measures the difference between the model and system probability distribution.	155
8.10	Experimental configuration for the axisymmetric wake. The bluff body is mounted from the wind tunnel ceiling (left) and the base pressure distribution is measured from 64 pressure taps (middle). The symmetry-breaking instability of the laminar flow persists in the fully turbulent wake, although the center of pressure appears to wander randomly, as seen in the probability distribution $p(x, y)$ (right).	156

8.11	Statistical evaluation of the axisymmetric wake models. The Langevin regression model (blue) better matches both the power spectral density (left) and tails of the PDF (right) compared with the pitchfork normal form (8.34) with coefficients estimated by PDF fitting, as in Rigas <i>et al.</i> (2015). The models are similar, but Langevin regression identifies a quadratic state-dependent noise (Fig. 8.12). The power spectrum is premultiplied by Strouhal number $St = fU/D$, a dimensionless frequency. The large peak at $St \approx 0.2$ corresponds to vortex shedding, which is essentially indistinguishable from the symmetry-breaking instability in the base pressure distribution (Brackston <i>et al.</i> , 2016).	158
8.12	Model selection and Kramers-Moyal coefficients for the axisymmetric wake. The reverse-greedy Sparse Stepwise Regression identifies a hierarchy of candidate models with varying tradeoffs between accuracy and complexity. The optimal model has the fewest terms before the cost function begins to climb, indicating the remaining terms are essential. In this case, the optimal model is a pitchfork bifurcation normal form forced by quadratic multiplicative noise (top). The model to the right of the optimal model includes only additive noise and corresponds to the model proposed by Rigas <i>et al.</i> (2015), while additional terms leads to a higher-order Stuart-Landau equation. The selected model closely matches the empirical finite-time Kramers-Moyal coefficients (bottom), the state PDF, and the power spectral density (Fig. 8.11).	159
9.1	Overview of the model development. Beginning with an order parameter $A(t)$ computed from the base pressure measurements, the spatial modes ψ_m are computed at each wavenumber m by conditional averaging. These modes define an approximate slow manifold $\psi(A)$ that captures the dominant anti-symmetric behavior. To fully resolve the symmetric fluctuations, we introduce a generalized shift mode with amplitude $B(t)$, defined by projection onto the tangent space of this manifold. Finally, we identify a nonlinear stochastic dynamical system model with Langevin regression. We can compare the statistics, including the empirical probability distribution ρ , of the model to the experiment with Monte Carlo evaluation of the stochastic system. Because the proposed model (bottom row) is random in nature, it can only reproduce the experimental data (top row) in a statistical sense; in general neither the time series nor the pressure distributions will match on a point-by-point basis.	164
9.2	Visualization of approximate slow manifolds and the nonlinear interactions responsible for generating them. Weakly nonlinear analyses of laminar flows typically neglect higher-order interactions for small fluctuations, leading to a parabolic manifold for the low-dimensional dynamics (left). Using conditional averaging, we show that these interactions are necessary to explain observations of the turbulent wake (right). Pairs of lines with similar styles indicate the structure of the leading-order triadic interactions. The physical interactions are between velocity components, but here we use the base pressure field as a proxy for coherent structures in the wake.	166

9.3	Exerimental configuration for the axisymmetric wake. The signature of the dominant symmetry-breaking instability is captured by the complex order parameter defined in (9.5). The coherent fluctuations in the pressure distribution can be approximated with the mean-field model (bottom). The proposed model significantly improves the reconstruction over a traditional fixed-mode decomposition.	168
9.4	Illustration of phase-aligned conditional averaging. We assume that the average at small amplitudes is representative of the pressure distribution associated with the unstable steady state q^B . The field at any other point can be approximated with spline interpolation, allowing us to explore the amplitude dependence of the coherent fields. Also shown for reference is the unconditional phase-aligned mean field and the two-dimensional log-probability distribution, which is roughly analogous to a potential field.	171
9.5	Field deformation in the conditional average. The phase-aligned average gives the coherent fields as a function of instability amplitude $ A $ for various azimuthal wavenumbers m . The deformation of the axisymmetric part of the field cannot be explained by a single fixed mode or weakly nonlinear scaling (A), even though the deformation has a smooth dependence on the amplitude (C). We approximate this amplitude dependence with a spline fit (C, dashed lines). On the other hand, the symmetry-breaking fields at $m = \pm 1, \pm 2$ are consistent with the fixed-mode assumption of Stuart-Landau theory (B, D).	173
9.6	Model for axisymmetric fluctuations. The conditional average on the order parameter A defines a slow manifold (left), but including an additional degree of freedom to account for fluctuations significantly improves the resolution of the $m = 0$ component (right). . .	174
9.7	Model selection with Langevin regression. The drift and diffusion functions are sparse linear combinations from a library of candidates. We select the simplest models that are statistically consistent (small cost function). The vertical axes are scaled by the value of the cost function with an identically zero coefficient vector.	178
9.8	Statistical evaluation of the model. The mean-field modal expansion accurately captures most features of the radially averaged premultiplied power spectrum based only on the order parameters $A(t)$ and $B(t)$ (gray). Monte Carlo evaluation of the Langevin model (blue) shows that it reproduces the the empirical probability distributions and dominant frequency content at the leading azimuthal wavenumbers.	180

10.1	Schematic of the dominant balance identification procedure applied to a turbulent boundary layer.	High-resolution direct numerical simulation results (a, visualized with a turbulent kinetic energy isosurface) are averaged to compute the terms in the Reynolds-averaged Navier-Stokes equations (b). The equation space representation of the field enables clustering and sparse approximation methods to extract the distinct geometrical structures in the six-dimensional space corresponding to dominant balance physics (c). Finally, the entire domain can be segmented according to these interpretable balance models, identifying distinct physical regimes (d). A curve fit to the wall-normal extent of the post-transition region of the identified inertial sublayer shows an approximate scaling of $\ell \sim x^{0.81}$, consistent with the theoretical prediction of $x^{4/5}$ from boundary layer theory. The 99 % free-stream velocity (U_∞) contour is also shown for reference.	185
10.2	Dominant balance of the viscous Burgers' equation.	The fields highlight identified dominant balance (a), with constituent terms shown in (b). The viscous term acts to diffuse sharp gradients and prevent formation of a discontinuous shock, but away from the shock front the dynamics are essentially inviscid. Away from the shock front, the field is approximately restricted to the $\nu u_{xx} = 0$ plane (c). This is reflected in the covariance matrices learned by the Gaussian mixture model (d).	189
10.3	Vorticity snapshot for the wake behind a cylinder at $Re = 100$ (a).	A Gaussian mixture model (GMM) assigns field points to clusters by looking for groups with distinct mean and covariance (b). For instance, some clusters vary mainly in the acceleration-advection directions, while others vary principally in the viscous-advection directions. We would expect these to represent the far-field and boundary regions, respectively. This is confirmed by the sparse principal components analysis (SPCA) reduction, where clusters with significant nonzero variance in the same directions are grouped together (c). These directions can be interpreted as active terms in the balance relation (d). As anticipated, the region near the cylinder is dominated by a balance between viscosity and advection and pressure forces, while the far wake is approximately inviscid (e).	193
10.4	The role of nonlinearity in stabilizing a von Kàrmàn vortex street.	The transient flow evolves from an unstable steady-state through the exponential growth of a linear instability mode to the vortex shedding limit cycle (a, visualized with lift coefficient C_L). Only one balance is identified that includes the nonlinear term $(\mathbf{u}' \cdot \nabla)\mathbf{u}'$. This balance does not appear significantly in the linear growth regime (d), consistent with the correspondence between the fluctuations and the instability mode. The fully saturated limit cycle is dominated by this balance, however, confirming the interpretation of the stabilizing feedback loop mediated by the nonlinearity.	194

10.5 **Self-similarity in the transitional boundary layer.** The spatial dominant balance regimes are consistent with known self-similarity in the boundary layer. The laminar balance (Fig. 10.1, purple and green) extends to approximately $x = 200$, a region over which the mean profile approximately matches the Blasius solution (left). Similarly, the turbulent viscous sublayer (Fig. 10.1, red) implies scaling with wall units (right). The mean profile with this scaling collapses until approximately the wall-normal extent of the identified viscous layer, as indicated by dashed lines. 199

ACKNOWLEDGMENTS

First and foremost I would like to thank my advisor Steve Brunton for all of his support throughout my time in graduate school. This would have been a much more difficult and less productive effort without all of the encouragement, guidance, and opportunities he has offered while at the same time remaining patient through the inevitable learning curves and dead ends. I am deeply grateful to both Steve and Nathan Kutz for their mentorship; their care for their students truly extends well beyond anything that could reasonably be asked of an advisor. Likewise, I would like to extend a special thanks to Matthew Munson, who was generous enough to take the time to mentor me, even after leaving the position that originally put us in contact. I would also like to thank Krithika Manohar, Duane Storti, and Armita Nourmohammad for serving on my thesis committee. Finally, my route here was long and not planned “a priori”; I owe a debt of gratitude to Jon Machta and Tom Daniel for introducing me to the world of research and encouraging me to pursue it at the graduate level.

Throughout my thesis work I have been especially grateful for the opportunity to meet and collaborate with amazing researchers during my time at UW, as well as current and former members of the Brunton and Kutz labs. Georgios Rigas and Jean-Christophe Loiseau in particular were instrumental in the development of the core contributions of this work; everything I learned in coursework pales in comparison to what I gained from our conversations, even fighting eight-to-ten-hour time differences. I would also like to thank our other collaborators, including Andres Goza, Aditya Nair, Gavin Portwood, Pino Martin, Hannah Neuenhoff, Joe Bakarji, Alan Kaptanoglu, Kazuki Maeda, Jordan Snyder, Urban Fasel, and the Boeing FT&E team. As is always the case with research, some of these projects never made it to completion, but it was still a valuable experience to work with all of them.

On a personal note, I would like to thank my parents for teaching me to value curiosity and hard work and showing me the worth of doing something for its own sake, as well as my brother for always reminding me that the world is a big place with a lot more in it than math and physics. I would also like to thank my friends, for occasionally listening to me talk about research, but mainly for giving me a reason to step away from it from time to time. Above all I am grateful for my family - my wife Shannon, our daughter Fiadh, and our dog Roo, who all came into my life during graduate school. As much as I've put into this thesis, they will always be the most important thing that happened to me during this time. I am endlessly appreciative of their support, encouragement, and love, always unwavering despite what the reviewers say.

Finally, I gratefully acknowledge the support of this work by the National Defense Science and Engineering Graduate Fellowship, sponsored by the Army Research Office, and the UCLA Institute for Pure and Applied Math.

Part I

INTRODUCTION

For more than two hundred years, the study of fluid mechanics has played a central role in a wide range of applied sciences due to the prevalence of a variety of fluid flows in the natural and engineered world. Once the basic theory of ideal and viscous flows was established, the field rapidly expanded to other domains; modern fluid dynamics encompasses topics as disparate as oceanography, plasma physics, hypersonic aerodynamics, biophysics, combustion, and atmospheric sciences. Despite the diversity of the physical phenomena relevant in each of these regimes, all are characterized by rich, complex structure spanning a wide range of spatiotemporal scales.

Turbulent flow is particularly notorious as Richard Feynman's "most important unsolved problem of classical physics." While that may remain true from a theoretical perspective, in terms of practical engineering the intervening six decades have seen remarkable advances in applied math and scientific computing, many of them inspired by open problems in fluid dynamics. For instance, computational power and advances in numerical methods and turbulence models have made it routine for mechanical and aerospace engineers to run simulations that would not have been imaginable on the world's largest supercomputer in the days of Feynman's Lectures.

Still, the virtuous cycle of theory and application may soon be tested by the staggering complexity characterizing many of the most exciting problems. For example, the majority of standard engineering simulations are designed to approximate a mean flow profile, disregarding any time variation in the flow. This is necessarily the case, since the computational cost of simulating unsteady flow to statistical convergence at full resolution and under realistic conditions strains the limits of even modern computational power, making it impractical for multi-query engineering problems such as uncertainty quantification and design optimization. However, neglecting time variation precludes the analysis and optimization of fundamentally transient phenomena such as acoustics, fluid-structure interaction, unsteady aerodynamics, active flow control, and the transition to turbulence.

There are a number of approaches to bridging this divide between fidelity and feasibility in unsteady fluid flows. As a classic example, in many regimes the equations can be simplified as a matter of dimensional analysis. Nonlinear advection may be negligible in small-scale biological fluid flows so that linear Stokes flow is a good approximation. In the other extreme, an analysis of high-speed external aerodynamics could assume that inertia is dominant and disregard viscous effects.

Numerically, large eddy simulation (LES) avoids the scaling problem of fully-resolved direct numerical simulation (DNS) by filtering out small-scale features of the flow while

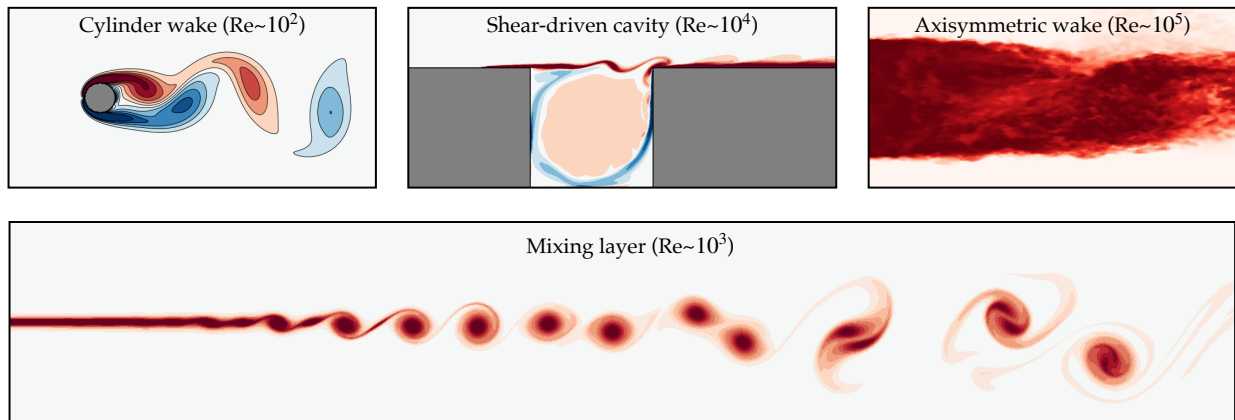


Figure 1.1: **Flow configurations modeled in this thesis.** The flows range span more than three orders of magnitude in Reynolds number from DNS of the laminar, two dimensional flow past a cylinder at $Re = 100$ (Chapter 5) to experimental measurements of a three-dimensional fully turbulent bluff body wake at $Re \approx 2 \times 10^5$ (Chapter 9). The shear-driven cavity (Chapter 6) and mixing layer (Chapter 7) are of intermediate dynamical complexity.

trying to approximate their influence on the resolved scales. Alternatively, the fluid flow can be approximated as perturbations to a fixed profile and analyzed with tools borrowed from linear systems or control theory, such as stability analysis or input/output modeling.

Somewhat more abstractly, fluid flow can be treated as a large-scale, dissipative, non-linear dynamical system. In this case unsteadiness is of primary importance; the flow might be approximated as a set of fixed spatial modes with time-varying coefficients that specify the evolution of the flow. This perspective became especially popular after the discovery that many turbulent flows are dominated by spatiotemporal coherent structures that evolve on relatively long scales in space and time compared to the incoherent background turbulence.

Although this may depart from standard terminology, in a sense all of these methods are a form of model reduction. The “full-order model” is the Navier-Stokes (or equivalent) governing equations derived from first-principles physics. Since it is often impractical and unnecessary to deal with directly with the governing equations, empirical or analytical simplifications must be introduced to reduce their complexity.

Beginning from this point of view, this thesis is devoted to exploring several model reduction strategies that directly exploit the multiscale structure of unsteady fluid flow. These include analytic and data-driven methods for low-dimensional modeling, non-linear stochastic system identification, and an automatic procedure that can determine which terms in the governing equations are locally dominant in space and time. The theoretical development of these methods is paired throughout with their application

to numerical and experimental data from flows spanning a wide range of complexity, as shown in Fig. 1.1. Although all of these flows are idealized, canonical configurations, they collectively exhibit key features found in practical flows, including symmetry-breaking, vortex shedding, shear layer instability, and broadband turbulence.

1.1 Survey of model reduction and related topics

Historical context Arguably the earliest form of model reduction for fluid flows was linear stability analysis; Lord Rayleigh derived the equation governing infinitesimal perturbations to an inviscid, parallel shear flow in 1880, for instance. Hydrodynamic stability theory only grew in importance following Osborne Reynolds' classic 1883 experiment illustrating the transition to turbulence in pipe flow in (Schmid & Henningson, 2001; Drazin & Reid, 2004; Charru, 2011). This transition is triggered by the relative balance of inertial and viscous forces in the flow; remarkably, this can be fully specified by a single dimensionless number now known as the Reynolds number when the governing equations are appropriately scaled. While the linear theory deals with the onset of instability and turbulence at relatively low Reynolds numbers, statistical theories of turbulence became central to the description of flows at higher Reynolds numbers (Taylor, 1935; Kolmogorov, 1941; Monin & Yaglom, 1971; Tennekes & Lumley, 1972; Pope, 2000).

The beginnings of the dynamical systems perspective of turbulence might be traced to two key advances in the latter half of the twentieth century. The first of these was the extension of hydrodynamic stability analysis to the weakly nonlinear regime via multiple-scale expansions (Landau, 1944; Landau & Lifshitz, 1959; Stuart, 1958, 1960, 1967, 1971). This enabled the description of flows near the threshold of instability in terms of a set of fixed spatial modes with time-varying coefficients whose evolution is governed by low-dimensional nonlinear amplitude equations. The second major development was the discovery of the existence and importance of large-scale *coherent structures* in turbulent flows (Brown & Roshko, 1974; Winant & Browand, 1974; Sirovich, 1987; Jiménez, 2018), along with the statistical tools to approximate them from data. The latter has progressed from proper orthogonal decomposition (POD) (Lumley, 1967; Berkooz *et al.*, 1993; Sirovich, 1987; Holmes *et al.*, 1996) and linear stochastic estimation (Adrian, 1975, 1979; Tung & Adrian, 1980) to modern modal analysis (Taira *et al.*, 2017, 2020), including non-modal stability analysis (Schmid, 2007; McKeon & Sharma, 2010), dynamic mode decomposition (DMD) (Schmid, 2010; Rowley *et al.*, 2009; Tu *et al.*, 2014; Kutz *et al.*, 2016), Koopman analysis (Mezić, 2005, 2013; Arbabi & Mezić, 2017; Brunton *et al.*, 2022), and spectral proper orthogonal decomposition (Towne *et al.*, 2018; Towne, 2021; Schmidt, 2020). More recently, it has been shown that in some globally unstable flow configurations the turbulent coherent structures may be closely related to the linear instability modes (Rigas *et al.*, 2014, 2016, 2017), connecting the coherent structures to weakly nonlinear analysis.

In conjunction with the contemporary discovery of deterministic chaos by Lorenz (1963), these developments raised the possibility of conceptualizing turbulence as a rela-

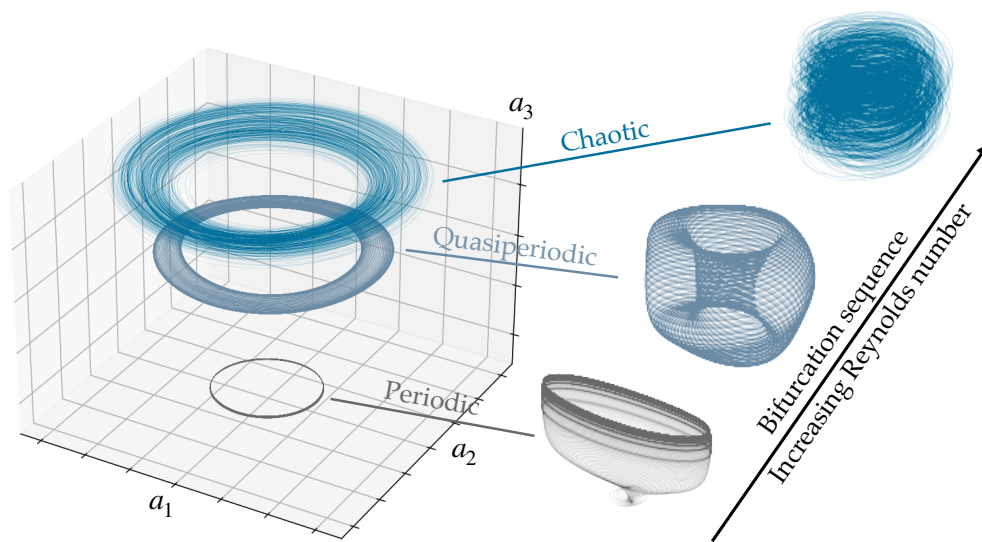


Figure 1.2: Visualization of the Ruelle-Takens route to chaos. In this scenario the steady flow first loses stability to a periodic limit cycle in a supercritical Hopf bifurcation. A secondary limit cycles leads to quasiperiodic dynamics on a 2-torus. A third bifurcation typically precipitates a breakdown into chaos. The large figure on the left is constructed by plotting the leading three proper orthogonal decomposition coefficients computed from a global analysis of the shear-driven cavity flow presented in Chapter 6 across different Reynolds numbers. The trajectories on the right apply Eq. (6.17) to the leading two pairs of oscillatory modes.

tively low-dimensional dynamical system, enabling new types of analysis. For instance, Landau proposed that the transition to turbulence might take place via a series of *Hopf bifurcations*, each of which added a new periodic component with incommensurate frequency until the spectrum appeared continuous (Landau & Lifshitz, 1959). Later, (Ruelle & Takens, 1971) showed that it was far more likely that the third in such a series of bifurcations would tend to destabilize the quasiperiodic flow, leading to a breakdown into chaos (see Fig. 1.1). These breakthroughs spurred the development of sophisticated analytic dynamical systems analyses (Lanford, 1982; Couillet & Spiegel, 1983; Golubitsky & Stewart, 1985; Golubitsky & Langford, 1988; Guckenheimer, 1986), often in an abstract mathematical context or for simplified, highly symmetric flow configurations such as plane Poiseuille, Taylor-Couette, and Rayleigh-Bénard flow. The Ruelle-Takens route to chaos is now thought to be one of the primary scenarios for a transition to turbulence in globally unstable flows (Gollub & Swinney, 1975; Swinney & Gollub, 1981), although other important scenarios include transient energy amplification in convectively unstable shear flows (Monkewitz & Huerre, 1982; Huerre & Monkewitz, 1985, 1990; Trefethen *et al.*, 1993; Reddy *et al.*, 1993; Schmid & Henningson, 2001; Schmid, 2007) and the spatial growth of turbulent patches (Eckhardt *et al.*, 2007; Hof *et al.*, 2010; Avila *et al.*, 2011; Barkley *et al.*, 2015).

Low-dimensional modeling The connection between coherent structures in turbulence and the dynamical systems perspective was further expanded by the seminal work of Aubry *et al.* (1988), who developed a low-dimensional model of a turbulent boundary layer by projecting the Navier-Stokes equations onto a POD basis derived from experimental measurements. The key appeal of this method is that it allows low-dimensional dynamical systems to be derived from the Navier-Stokes equations for arbitrary geometries by means of the empirical POD representation. This framework was soon applied to a variety of other flows, including vortex shedding in the cylinder wake (Deane *et al.*, 1991; Ma & Karniadakis, 2002), three-dimensional coherent structures in a boundary layer (Rempfer & Fasel, 1994), a turbulent mixing layer (Ukeiley *et al.*, 2001), and compressible cavity flow (Rowley *et al.*, 2004; Rowley & Williams, 2006).

Although POD-Galerkin models are derived from combining an energy-optimal modal basis with the first-principles governing equations, all of these early models were unstable at long times unless they incorporated a carefully tuned “turbulent viscosity” term approximating the energy dissipated at scales of the flow that were not resolved by the leading POD modes. Later developments in projection-based modeling methodology for fluid flows have largely centered on stabilization schemes, including accounting for the Stuart-Landau mean flow deformation mechanism (Noack *et al.*, 2003), modeling the projected pressure gradient term (Noack *et al.*, 2005; Caiazzo *et al.*, 2014), and subscale closure models (Noack *et al.*, 2008; Wang *et al.*, 2012; Cordier *et al.*, 2013; Balajewicz *et al.*, 2013; Östh *et al.*, 2014; Mohebujjaman *et al.*, 2017, 2018; Xie *et al.*, 2018; San & Maulik, 2018a,b; Maulik *et al.*, 2019; Menier *et al.*, 2022). See Chapter 4 for further discussion of

POD-Galerkin projection and Holmes *et al.* (1996); Noack *et al.* (2011) for more in-depth presentations.

Whether derived from POD-Galerkin projection, linear model reduction, or system identification, a longstanding objective of low-dimensional modeling has been the design of model-based controllers for closed-loop flow control (Noack *et al.*, 2011; Brunton & Noack, 2015a; Rowley & Dawson, 2017). This has been explored for a variety of flows, including the cylinder wake (Tadmor *et al.*, 2004, 2011), cavity flow (Rowley & Juttijudata, 2005; Rowley & Williams, 2006; Rowley *et al.*, 2006; Barbagallo *et al.*, 2009; Sipp *et al.*, 2010; Leclercq *et al.*, 2019), boundary layer (Bagheri *et al.*, 2009a), high-lift configuration (Luchtenburg *et al.*, 2009), and unsteady aerodynamics (Ma *et al.*, 2011; Brunton *et al.*, 2014; Hemati *et al.*, 2017). Although these studies have shown promising results, significant challenges remain for practical flow control based on empirical models, including capturing parametric variation and predicting the response to unseen actuation (related to sensitivity or adjoint modes of the flow).

Data-driven modeling Just as rapidly expanding computational capabilities have enabled advances in numerical simulation, they have also driven the emergence of the so-called fourth paradigm of data science (after experimental, analytic, and numerical methods). Machine learning has led to recent breakthroughs and even superhuman performance in a variety of computational domains, such as computer vision, natural language processing, and reinforcement learning (Goodfellow *et al.*, 2016). It has become a common refrain that these techniques also stand to revolutionize scientific and engineering disciplines. Despite the obvious promise of machine learning, significant obstacles remain before this renaissance can be realized (Brenner *et al.*, 2019; Brunton *et al.*, 2020).

For example, deep learning is often framed as a high-dimensional nonlinear regression problem which relies on particulars of architecture and training algorithms combined with enormous training sets to *implicitly* constrain and regularize an otherwise non-convex optimization problem. In contrast, many problems of interest in science and engineering are characterized by data that is relatively expensive to acquire (compared to image recognition, for example). Furthermore, problems are often constrained *explicitly* by centuries of domain expertise. Although fully black-box machine learning models have shown promise (Gautier *et al.*, 2015; Verma *et al.*, 2018; Pathak *et al.*, 2018; Hesthaven & Ubbiali, 2018; Wan *et al.*, 2018), it is usually at the cost of discarding significant physical knowledge.

One perspective that is gaining attention and has already produced some remarkable results is “physics-informed machine learning,” where the goal is typically to incorporate insights from domain knowledge as additional constraints to the learning problem (Ling *et al.*, 2016; Thomas *et al.*, 2018; Schütt *et al.*, 2018; Raissi & Karniadakis, 2018; Noé *et al.*, 2019; Bar-Sinai *et al.*, 2019; Cranmer *et al.*, 2020a). The methods discussed in this thesis will generally align with the alternate perspective of “data-informed physics.” Rather than designing machine learning algorithms which incorporate physics constraints, we

will instead focus on the use of *simple* and *targeted* data-driven methods to fill specific gaps which have proven challenging to traditional analysis.

A specific framework that is especially well-suited to this philosophy is the sparse identification of nonlinear dynamics (SINDy) algorithm. Originally proposed by (Brunton *et al.*, 2016a) and later extended to partial differential equations (Rudy *et al.*, 2017; Kaheman *et al.*, 2020), constrained systems (Loiseau & Brunton, 2018; Kaptanoglu *et al.*, 2021a), exogenous inputs (Brunton *et al.*, 2016b), coordinate system discovery (Champion *et al.*, 2019a), stochastic dynamics (Boninsegna *et al.*, 2018; Callaham *et al.*, 2021b), and with several optimization algorithms (Champion *et al.*, 2019b; de Silva *et al.*, 2021; Kaptanoglu *et al.*, 2022), the method is essentially sparse polynomial regression applied to nonlinear dynamical systems. It has been successfully applied to a variety of systems, including the cylinder wake (Loiseau & Brunton, 2018; Loiseau *et al.*, 2018a,b), thermal convection (Loiseau, 2020), the fluidic pinball (Deng *et al.*, 2020), magnetohydrodynamics (Kaptanoglu *et al.*, 2021b), electroconvection (Guan *et al.*, 2021), and shear-driven cavity flow (Callaham *et al.*, 2022). See Sec. 4.2 for a presentation of the basic SINDy algorithm, Chapter 8 for an extension to stochastic systems, and Chapters 5, 6, and 9 for applications.

Operator-based analysis Although they will not be applied in this thesis, there is a final important category of methods in modern fluid dynamics modeling that should be mentioned. These might be described broadly as operator-based methods, since they rely on computational manipulation of the discretized Navier-Stokes operators and are thus neither standard time-marching numerical simulations nor traditional analytic methods.

For example, traditional linear stability analysis relied on assuming there was only one inhomogeneous direction in the flow, thus restricting the method to parallel or weakly nonparallel flows (Schmid & Henningson, 2001; Drazin & Reid, 2004; Charru, 2011). Alongside improvements in computational capacity, the development of matrix-free methods based on Krylov sequences (Edwards *et al.*, 1994; Stewart, 2001; Knoll & Keyes, 2004; Antoulas, 2005) has enabled the efficient solution of the large-scale linear algebra problems arising in two- and three-dimensional stability analysis (Theofilis, 2011; Loiseau *et al.*, 2019).

Resolvent analysis is a closely related topic to global stability analysis; rather than searching for the perturbations with the largest exponential growth rate, resolvent analysis uses a singular value decomposition to identify the most-amplified perturbations over a *finite* time horizon (Trefethen *et al.*, 1993; Reddy *et al.*, 1993; Schmid, 2007; McKeon & Sharma, 2010; McKeon *et al.*, 2013; Sharma & McKeon, 2013). The distinction is subtle but important in shear-driven flows, for which the linearized operator may be globally stable but still exhibit significant transient energy amplification as a result of operator non-normality. Resolvent analysis has proven to be a powerful framework for analysis and control of turbulent flows (Luhar *et al.*, 2014; McKeon, 2017; Saxton-Fox & McKeon, 2017; Symon *et al.*, 2018; Yeh & Taira, 2019; Rosenberg & McKeon, 2019; Rosenberg *et al.*, 2019; Jacobi *et al.*, 2021; Bae *et al.*, 2021; Ahmed *et al.*, 2021; Pickering *et al.*, 2021; Liu *et al.*,

2021).

As with stability analysis, the classical weakly nonlinear analyses presented by Stuart and others relied on highly symmetric flows, typically with only one inhomogeneous direction. The method requires solving a series of high-dimensional linear systems and both direct and adjoint eigenvalue problems. Since the extension to fully inhomogeneous flows in a finite element framework by Sipp & Lebedev (2007), weakly nonlinear analysis has been used to study bluff body wakes (Meliga *et al.*, 2009; Rigas *et al.*, 2017), vortex-induced vibrations (Meliga & Chomaz, 2011), and thermoacoustic instabilities (Orchini *et al.*, 2016; Ghirardo *et al.*, 2016). Although weakly nonlinear analysis is a rigorous and accurate method close to the bifurcation point, its validity quickly decays in the non-asymptotic regime (Gallaire *et al.*, 2016). Alternatively, Mantič-Lugo *et al.* (2014); Mantič-Lugo & Gallaire (2016) recently proposed the self-consistent mean flow modeling framework, which is able to accurately predict both the mean flow and instability modes by iteratively solving a sequence of linear problems until the two are in a Reynolds stress-mediated equilibrium. The self-consistent method has also been extended to second order to examine the generation of harmonics in cavity flow (Meliga, 2017).

Although operator-based methods are often more technically and computationally demanding than other approaches to modeling, they benefit from a lack of dependence on an ensemble of observations. In principle, this suggests that they may generalize to different geometries, controller designs, or parametric regimes more readily than methods that rely heavily on empirical approximation. Although many of the operator-based methods are still early in their development, it is a promising line of research, particularly as computational capabilities and linear algebra algorithms continue to advance.

1.2 Organization and contributions

The objective of this thesis is to develop semi-empirical model reduction methods that leverage the multiscale structure of unsteady fluid flows. The contributions can be broadly separated into three parts: approximating unresolved variables in reduced-order models, nonlinear stochastic system identification, and dominant balance modeling. Before presenting these contributions, Part II provides a fairly comprehensive introduction to reduced-order modeling. The objective of these early chapters is to establish terminology and conceptual background that will be common throughout the later chapters, making the document as self-contained as possible.

Chapter 2 contains an overview of the key mathematical concepts that will be used throughout the development of the proposed modeling methods. This chapter begins with a brief summary of relevant properties of the incompressible Navier-Stokes equations (Sec. 2.1). Section 2.2 continues with a simplified weakly nonlinear analysis that both illustrates a mechanism by which low-dimensional structure emerges in dissipative partial differential equations (PDEs) and introduces important mathematical concepts in the context of a toy problem. In the same vein, Sec. 2.3 presents some relevant aspects

of dynamical systems theory, in particular the center manifold reduction. Finally, Sec. 2.4 provides an introduction to the representation of nonlinear dynamics in terms of linear PDEs, likely the most abstract content in this thesis.

Chapters 3 and 4 review standard results in modal analysis and low-dimensional modeling, respectively. These chapters do not introduce novel contributions and are intended mainly for reference and completeness, covering topics such as proper orthogonal decomposition (Sec. 3.2), Galerkin projection (Sec. 4.1), and the SINDy framework (Sec. 4.2). Finally, Chapter 5 demonstrates the application of the model reduction machinery developed in Chapters 2, 3, 4 to vortex shedding in the flow past a cylinder. This chapter has the structure of an abridged review article and serves as a more intuitive presentation of the more technical material in the earlier chapters of Part II. The remaining chapters represent the core contributions of this thesis.

Part III explores improvements in accuracy and stability of low-dimensional models of fluid flow based on dynamical systems analysis. Chapter 6 generalizes the analytic center manifold reduction to flows that are not close to a bifurcation. The concept is first developed in terms of nonlinear correlations analysis and model problems, before a SINDy-type polynomial regression is used to reduce quasiperiodic shear-driven cavity flow to a highly accurate 4-dimensional model. The chapter concludes with an even more flexible representation of the manifold equation in terms of a deep autoencoder, which is able to model the cavity flow throughout the full Ruelle-Takens route to chaos.

Chapter 7 builds on the variable-elimination strategy of the previous chapter. However, rather than approximating the unresolved variables directly as algebraic functions of the remaining degrees of freedom as in the manifold model, this chapter introduces a stochastic averaging strategy that only attempts to capture the average influence of the subscale variables. This approach transforms the linear-quadratic Galerkin system into a cubic generalized Stuart-Landau equation that can be viewed as a system of coupled nonlinear oscillators. This method is shown to be consistent with the invariant manifold model of the cylinder wake presented in Chapter 5 before it is applied to more complex flows, including the mixing layer shown in Fig. 1.1.

Part IV shifts focus from physics-based reduced-order modeling to empirical nonlinear stochastic system identification. The scale separation is now implied in the model structure, where any unobserved variables are treated as process noise. Chapter 8 first presents the Langevin regression algorithm, a stochastic variant of SINDy, including its application to several test problems. Chapter 9 then applies Langevin regression to experimental measurements of the turbulent axisymmetric wake shown in Fig. 1.1. In conjunction with a novel parametric modal decomposition that generalizes the mean flow model presented in Chapter 5 for the cylinder wake, Langevin regression identifies a compact stochastic model that captures key aspects of the symmetry-breaking wake deflection in the turbulent flow.

Finally, Part V offers a different perspective on model reduction in terms of scaling analysis. Rather than modeling the flow as a dynamical system, each point in the domain

is treated as a vector in an “equation space” in which the dominant balance between terms in the governing PDE emerges naturally as geometric structure. The subset of active local terms can then be identified with simple clustering and classification tools from machine learning. The proposed dominant balance identification method is demonstrated on the mean flow profile of a boundary layer in transition to turbulence, as well as the unsteady cylinder wake.

Part VI concludes with a summary and outlook, including a critical assessment of the viability of these strategies for the engineering objectives of analysis, design, and control.

Part II
BACKGROUND

Chapter 2

MATHEMATICAL BACKGROUND

2.1 Governing equations for incompressible fluid flow

We will consider the incompressible Navier-Stokes equations as our prototype physical system, although many of the techniques described here will generalize to similar linear-quadratic dynamics (Rowley *et al.*, 2004; Qian *et al.*, 2020; Kaptanoglu *et al.*, 2021b). For an incompressible flow the state is given by the velocity $\mathbf{u}(\mathbf{x}, t) = (u, v, w)$ and pressure field $p(\mathbf{x}, t)$. If the fluid has density ρ and dynamic viscosity μ and the flow is characterized by length scale U and velocity scale L , the equations of motion can be expressed in dimensionless form, parameterized by the Reynolds number $Re = \rho UL/\mu$:

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot (\mathbf{u} \otimes \mathbf{u}) = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u} \quad (2.1a)$$

$$\nabla \cdot \mathbf{u} = 0. \quad (2.1b)$$

For a viscous flow, solid boundaries have no-slip Dirichlet boundary conditions $\mathbf{u} = 0$, while other appropriate boundary conditions depend on the flow configuration and numerical schemes. For incompressible flows, the divergence form of the nonlinearity $\nabla \cdot (\mathbf{u} \otimes \mathbf{u})$ is equivalent to the more common advection form $(\mathbf{u} \cdot \nabla) \mathbf{u}$ via the tensor divergence identity $\nabla \cdot (\mathbf{u} \otimes \mathbf{u}) = (\mathbf{u} \cdot \nabla) \mathbf{u} + \mathbf{u} (\nabla \cdot \mathbf{u})$ and incompressibility condition (2.1b).

The Reynolds number is essentially a balance parameter that captures the relative importance of advection and diffusion in the flow. High Reynolds-number configurations like external aerodynamic flows are characterized by strong advection and weak diffusion away from thin viscous layers near boundaries. In the opposite regime advection is negligible and the flow is a balance between viscosity and the pressure gradient (also known as Stokes flow). Chapter 10 will return to the idea of dominant balance and introduce an algorithm for automatically identifying regions of a flow with different effective balances.

The Reynolds number also often acts as a bifurcation parameter; low Reynolds number flows are typically smooth and laminar, while high Reynolds number turbulence is chaotic and irregular, with broadband spectral content in both space and time. Scaling arguments show that the ratio between the largest and smallest scales of a flow should scale with $Re^{9/4}$ (Tennekes & Lumley, 1972). The transition between the two is a subject of long-standing debate (Burgers, 1948; Hopf, 1948; Hof *et al.*, 2003; Eckhardt *et al.*, 2007; Avila *et al.*, 2011; Barkley *et al.*, 2015). It may occur via several scenarios, includ-

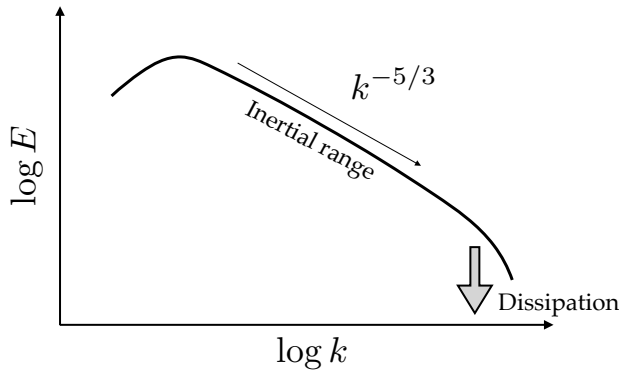


Figure 2.1: **The Richardson-Kolmogorov turbulent energy cascade.** For isotropic turbulence, the Fourier space dynamics suggests that energy is transferred from the largest integral scales to the smallest dissipative scales by the nonlinear triadic interactions. Scaling analysis suggests that the energy content should decay with wavenumber like $E \sim k^{-5/3}$.

ing a sequence of Hopf bifurcations (Fig. 1.2, see also (Ruelle & Takens, 1971; Gollub & Swinney, 1975)), intermittency, or bypass transitions (Barkley, 2016). One of the goals of low-dimensional modeling is to enable analysis of fluid flows as large-scale dynamical systems and reach a better understanding of phenomena such as the transition to chaos and turbulence.

Nonlinearity: spectral dynamics and the Richardson-Kolmogorov cascade There are many challenges in dealing with the Navier-Stokes equations analytically and numerically, but one of the more important from a modeling standpoint is the effect of the quadratic nonlinearity. For a broad class of boundary conditions, the nonlinearity is energy-conserving in the sense that its role is to transfer energy between scales, rather than add or remove it from the system (Kraichnan & Chen, 1989; Holmes *et al.*, 1996; Schlegel & Noack, 2015). This can be most easily seen for homogeneous isotropic turbulence in the Fourier domain.

First, assuming the spatial domain is either infinite or periodic, the velocity and pressure fields can be expanded in Fourier series, e.g.

$$\mathbf{u}(\mathbf{x}, t) = \sum_{\mathbf{k}} \hat{\mathbf{u}}_{\mathbf{k}}(t) e^{i\mathbf{k} \cdot \mathbf{x}}. \quad (2.2)$$

Substituting the Fourier expansion into the momentum equation and making use of the convolution theorem and orthogonality properties of Fourier series (Tennekes & Lumley, 1972),

$$\frac{d}{dt} \hat{\mathbf{u}}_{\mathbf{k}} = - \sum_{\ell} (i\mathbf{k} \cdot \hat{\mathbf{u}}_{\mathbf{k}-\ell}) \hat{\mathbf{u}}_{\ell} - ik\hat{p}_{\mathbf{k}} - |\mathbf{k}|^2 \hat{\mathbf{u}}_{\mathbf{k}}. \quad (2.3)$$

The advective nonlinearity becomes a convolutional sum in the spectral domain, where the flow at one spatial scale can be influenced by pairwise interactions between *all other* scales. This phenomenon is known as triadic interaction. Note that the pressure may be further eliminated by means of the continuity equation (2.1b), leading to an additional quadratic term, but this does not substantially change the nature of the nonlinearity.

The evolution equation in Fourier space also highlights the scale-dependent nature of viscous effects. The dissipation in the spectral domain scales with $|\mathbf{k}|^2$, so it is strongest at the smallest scales. These observations lead to the idea of a “turbulent cascade” (Richardson, 1922). If energy is transferred between scales by triadic interactions which are predominantly local in wavenumber space and dissipated at the smallest scales, turbulence can be conceptualized as a cascade of energy from the largest energy-containing scales to the smallest dissipative scales (Fig. 2.1).

Based on considerations of spectral dynamics, Kolmogorov (1941) proposed a universal equilibrium theory of turbulence. In this theory, large-scale anisotropy introduced by boundary conditions or body forcing is destroyed at smaller scales by the isotropic action of pressure gradient forces. If the Reynolds number is large enough that there are several orders of magnitude between the large energy-containing scales and the scales at which viscosity is dominant, there should be an intermediate region characterized by neither significant anisotropy nor dissipation. Kolmogorov argued that this “inertial range” should be close to statistical equilibrium and used scaling analysis to show that the spectral energy content $E(\mathbf{k})$ should scale fall off with $E(\mathbf{k}) \propto |\mathbf{k}|^{-5/3}$.

This theory was developed for high-Reynolds number turbulence, but many of the qualitative conclusions are thought to hold for a wide range of flows. In particular, the importance of the triadic interactions introduces a “closure problem” to many modeling efforts. Since the full range of scales is dynamically significant, unresolved degrees of freedom must be accounted for in order for the models to be physically relevant. For example, heavily truncated reduced order models can become unstable if accurate dissipation is not ensured (Aubry *et al.*, 1988; Rempfer & Fasel, 1994; Noack *et al.*, 2008; Cordier *et al.*, 2013; Östh *et al.*, 2014).

Reynolds averaging and the closure problem In many applications it is neither necessary nor desirable to solve the full unsteady Navier-Stokes equations. For instance, an engineer interested in the aerodynamic performance of a car or aircraft wing would typically only need average profiles of surface pressure, skin friction, etc., but not detailed time-resolved measurements of the turbulent flow field.

A straightforward approach is to neglect the time derivative in the momentum equation (2.1a), which yields the steady-state equations

$$\nabla \cdot (\mathbf{u}_0 \otimes \mathbf{u}_0) = -\nabla p_0 + \frac{1}{\text{Re}} \nabla^2 \mathbf{u}_0, \quad (2.4a)$$

$$\nabla \cdot \mathbf{u}_0 = 0 \quad (2.4b)$$

for the stationary flow fields $\mathbf{u}_0(\mathbf{x})$ and $p_0(\mathbf{x})$. The problem with this approach is that at even moderate Reynolds numbers the steady state is often strongly unstable, meaning that a typical unsteady flow field may not resemble the steady solution in any useful way.

The reason for this is that the mean flow $\bar{\mathbf{u}}(\mathbf{x}), \bar{p}(\mathbf{x})$ is not necessarily the same as the

steady state¹. This can be seen by defining the *Reynolds decomposition* for velocity

$$\mathbf{u}(\mathbf{x}, t) = \bar{\mathbf{u}}(\mathbf{x}) + \mathbf{u}'(\mathbf{x}, t), \quad (2.5)$$

and likewise for pressure. Substituting into the Navier-Stokes equations and averaging, with use of the fact that the fluctuations \mathbf{u}', p' have zero mean by definition,

$$\nabla \cdot (\bar{\mathbf{u}} \otimes \bar{\mathbf{u}}) + \nabla \cdot \overline{(\mathbf{u}' \otimes \mathbf{u}')} = -\nabla \bar{p} + \frac{1}{\text{Re}} \nabla^2 \bar{\mathbf{u}} \quad (2.6a)$$

$$\nabla \cdot \bar{\mathbf{u}} = 0. \quad (2.6b)$$

Equations (2.6) are known as the Reynolds-averaged Navier-Stokes (RANS) equations. Comparing to the steady-state equations (2.4), the effect of the fluctuations is to deform the base flow via the Reynolds stress tensor $\nabla \cdot \overline{(\mathbf{u}' \otimes \mathbf{u}')}$. The mean flow need not even be a solution to the regular Navier-Stokes equations. Clearly, this discrepancy between the mean and steady-state is a result of the nonlinear interactions, giving a complementary mean-flow view of the role of nonlinearity to the energy cascade.

Unfortunately, since the Reynolds stress is not a fundamental quantity of the flow but a second moment, the RANS equations are not algebraically *closed* in the sense that there are more unknowns than equations. While it is possible to develop an equation for the Reynolds stresses, it includes terms depending on the third moment (Tennekes & Lumley, 1972; Pope, 2000). Although a significant amount of work has gone into modeling the Reynolds stresses and approximations with varying degrees of accuracy are widely used, the closure problem remains open. Turbulence methods for the RANS equations are not directly connected to reduced-order modeling, but Reynolds stress deformation and the importance of the closure problem remain essential to modern perspectives on low-dimensional modeling in fluid dynamics (Noack *et al.*, 2003; Sipp & Lebedev, 2007; Noack *et al.*, 2008; Meliga *et al.*, 2009).

2.2 Motivation for low-dimensional modeling

Many systems with complex, multiscale structure are nevertheless characterized by emergent large-scale coherence (Haken, 1983; Cross & Hohenberg, 1993), generating low-dimensional structure often conceptualized as an attracting or slow manifold. This phenomenon is especially relevant in fluid dynamics, where successive bifurcations can lead to increasingly complex behavior and eventually the transition to turbulence (Landau, 1944; Stuart, 1958; Lorenz, 1963; Ruelle & Takens, 1971; Swinney & Gollub, 1981). Dynamical models that capture this intrinsic low-dimensional structure can improve our physical understanding and are critical for real-time optimization and control objectives (Noack *et al.*, 2011;

¹The ensemble average will typically be used interchangeably with the time-averaged flow here, assuming ergodicity when necessary




Figure 2.2: Rayleigh-Bénard flow The fluid between differentially heated plates conducts heat upwards with zero fluid velocity at low values of the control parameter (Rayleigh number). Above a critical value the steady conducting state is unstable to perturbations, eventually reaching a new nonlinear equilibrium. This flow is a prototypical example of instability and pattern formation.

Brunton & Noack, 2015*b*; Rowley & Dawson, 2017).

Close to a bifurcation, the dynamics are approximately restricted to the invariant manifold described by the amplitudes of the unstable eigenmodes. The evolution equations for these effective coordinates are given by the normal form for the bifurcation (Guckenheimer & Holmes, 1983), the form of which can be deduced with symmetry arguments (Golubitsky & Langford, 1988; Glaz *et al.*, 2017; Deng *et al.*, 2020), weakly nonlinear analysis (Stuart, 1958; Sipp & Lebedev, 2007; Meliga *et al.*, 2009), or a center manifold reduction (Carini *et al.*, 2015). Normal forms can describe a wide range of stereotypical dynamics, including bistability, self-sustained oscillations, and chaos.

Before delving into the details of constructing stable and accurate reduced-order models, it may be helpful to examine some of the mechanisms by which this reduction in effective dimension occurs. Historically, much progress was made in the study of instability, bifurcations, and the nonlinear dynamics in large-scale dissipative systems by studying canonical configurations such as Rayleigh-Bénard convection, Taylor-Couette flow, and Poiseuille (pipe or channel) flow (Cross & Hohenberg, 1993). These flows are geometrically simple and highly symmetric, making them ideal candidates both for experimental design and theoretical analysis. In all of these cases, the complexity of the flow is governed by one or more control parameters; at low values the solutions are steady, with instability and pattern formation occurring spontaneously above a critical value.

Although these flows are simplified relative to those of practical significance in engineering applications, they are still three-dimensional and involve the intricacies of treating interactions between primitive variables (buoyancy, pressure Poisson equations, etc). As an even more simplified model problem that exhibits many of the same features as these canonical problems, Swift & Hohenberg (1977) introduced a model that approximates the Boussinesq equations governing Rayleigh-Bénard convection. For a one-dimensional scalar field $q(x, t)$ (representing velocity in the full fluid flow equations, for instance), the




Figure 2.3: **Stability of the Swift-Hohenberg equation.** Above the critical value r_c , the steady solution is unstable to perturbations near the critical wavenumber k_c with maximum growth rate $\epsilon = r - r_c$.

Swift-Hohenberg equation is

$$\partial_t q = (r - r_c)q - (\partial_x^2 + k_c^2)^2 q + \frac{1}{2} \partial_x q^2, \quad x \in (-\infty, \infty), \quad (2.7)$$

where r is the control parameter, r_c and k_c are real constants, and ∂_t , ∂_x indicate partial derivatives with respect to the subscript variable. The linear part of this equation can be derived from the Boussinesq equations, while the Burgers nonlinearity $\partial_x q^2$ is a simplified one-dimensional analogue of the quadratic nonlinearity in the Navier-Stokes equations.

As will be seen, below $r = r_c$ the steady solution $q = 0$ is stable, while above this critical value a linear instability corresponding to steady convection rolls in Rayleigh-Bénard flow emerges (see Fig. 2.2). The linear theory cannot predict the behavior of the flow once it reaches finite amplitude, but a multiple-scale expansion shows that it in fact reaches a new equilibrium state that is stabilized by the nonlinear self-interaction of the fluctuations. Although this type of analytic treatment cannot be applied to even the slightly more complicated geometries that will be introduced in subsequent chapters, the proposed methods for model reduction are motivated by this view of the role of nonlinear interactions. Moreover, the following procedure is nearly identical (though simpler) to that used for weakly nonlinear analysis of fluid flows (Sipp & Lebedev, 2007).

It can easily be seen from Eq. (2.7) that $q_0 = 0$ is a steady solution. To gauge its stability, we will perturb q_0 with wavelike fluctuations proportional to $e^{ik(x-ct)}$ for real wavenumber k , and complex wave speed $c = c_r + ic_i$. If $c_i > 0$, the perturbation will grow exponentially with growth rate $\sigma = kc_i$. Since $c_r \neq 0$ is not admissible in the linearized form of Eq. (2.7), this leads to the dispersion relation

$$\sigma = (r - r_c) - (k_c^2 - k^2)^2. \quad (2.8)$$

Figure 2.3 shows the stability profile of growth rate as a function of wavenumber above and below the critical value. For r slightly above r_c the growth rate at $k = k_c$ is $\epsilon = r - r_c$. Smaller and larger wavenumbers are both stable, with $|k| > k_c$ subject to espe-

cially large damping. We therefore expect that the behavior on a slow timescale $\tau = \epsilon t$ will be primarily driven by the pair of marginally stable modes with $k = \pm k_c$. These modes are said to span the *center manifold* of the system at the bifurcation point, in contrast to the stable manifold spanned by the remaining eigenfunctions.

Introducing the small parameter $\epsilon = r - r_c$, we will look for a perturbation series solution valid for $0\epsilon \ll 1$.

$$q_\epsilon(x, t, \tau) = q_0(x) + \sqrt{\epsilon}q_1(x, t, \tau) + \epsilon q_2(x, t, \tau) + \epsilon\sqrt{\epsilon}q_3(x, t, \tau) + \dots \quad (2.9)$$

The time derivative must now be taken with respect to both timescales, so that $\partial_t \rightarrow \partial_t + \epsilon\partial_\tau$. Substituting into the governing equation (2.7), we can isolate each order in $\sqrt{\epsilon}$.

ϵ^0 : The zeroth-order equation gives the condition for the steady-state solution $q_0(x)$. In this case $q_0(x) = 0$ and can be dropped from further analysis, but more generally the steady solution is nonzero and must be carried through to higher orders.

$\epsilon^{1/2}$: The first-order equation is the linear problem

$$\partial_t q_1 = -(\partial_x^2 + k_c^2)q_1 \equiv \mathcal{L}q_1. \quad (2.10)$$

Based on the previous discussion, we will take as the first-order solution the neutrally-stable mode modified by a slowly-varying complex amplitude $A(\tau)$:

$$q_1(x, \tau) = A(\tau)e^{ik_c x} + \text{c.c.} \quad (2.11)$$

If we could find an evolution equation for the amplitude $A(\tau)$, which would give a closed ODE approximation to the PDE (2.7) close to the bifurcation. However, this will require solving for higher-orders in ϵ until the slow dynamics for $A(\tau)$ can be identified.

Before proceeding to higher orders in ϵ , note that because q_1 is not a function of the fast time scale t , the linear problem can also be written as a homogeneous equation in terms of the operator $\mathcal{L} = -(\partial_x^2 + k_c^2)$:

$$\mathcal{L}q_1 = 0. \quad (2.12)$$

Since q_1 and its conjugate are the only eigenfunctions of \mathcal{L} with zero eigenvalue according to the dispersion relation (2.8), they span its nullspace. Note that in general the frequency of the marginally stable mode will not be zero, so the relevant linear operator would be frequency-dependent (Sipp & Lebedev, 2007).

This likewise implies that the adjoint operator \mathcal{L}^\dagger , defined for any functions $p(x), q(x)$ by $\langle p|\mathcal{L}q \rangle = \langle \mathcal{L}^\dagger p|q \rangle$ for the inner product

$$\langle p|q \rangle = \int p^*(x)q(x)dx, \quad (2.13)$$

has a two-dimensional nullspace spanned by the adjoint mode p_1 and its conjugate, which solves the adjoint equation

$$\mathcal{L}^\dagger p_1 = 0. \quad (2.14)$$

In this case \mathcal{L} is self-adjoint, so $p_1 = q_1$ up to some complex normalization \hat{p}_1 . Although \hat{p}_1 could be a function of the slow time τ , for the purposes of deriving an amplitude equation it can be treated as a constant without loss of generality.

The utility of this digression is the *Fredholm alternative*, which can be roughly stated as the condition that, for an inhomogeneous linear operator equation $\mathcal{L}q = f$, one and only one of the following is true:

1. $\mathcal{L}q = f$ has a unique solution $q(x)$
2. $\mathcal{L}^\dagger p = 0$ has a solution with $\langle p|f \rangle \neq 0$.

The Fredholm alternative can be used to establish solvability conditions on the linear equation. In particular, in order for $\mathcal{L}q = f$ to have a unique solution (case 1), the forcing function $f(x)$ must be orthogonal to the nullspace of \mathcal{L}^\dagger (thereby ruling out case 2). Since the nullspace of \mathcal{L}^\dagger is spanned by p_1 and its conjugate, the solvability condition for $\mathcal{L}q = f$ is simply that $\langle p_1|f \rangle = 0$, and likewise for its conjugate p_1^* . This is a more general form of the resonance condition used to eliminate secular terms in classical analyses of weakly nonlinear oscillators (Bender & Orszag, 1999).

ϵ^1 : The second-order equation is the inhomogeneous linear problem

$$(\partial_t - \mathcal{L})q_2 = \frac{1}{2}\partial_x q_1^2 \quad (2.15a)$$

$$= ik_c A^2 e^{2ik_c x} + \text{c.c.} \quad (2.15b)$$

Based on the right-hand side forcing, q_2 can be identified as the response of the linear operator to the self-interaction of the neutrally stable mode $q_1(x, \tau)$. In other words, q_2 is the harmonic of q_1 .

The solvability condition from the Fredholm alternative is trivially satisfied in this case due to the orthogonality property of Fourier modes, since the right-hand side forcing is at wavenumber $2k_c$ while the adjoint mode p_1 has wavenumber k_c . That is, q_2 is not an eigenvector of \mathcal{L} . The equation may then be solved directly by assuming a form of q_2 that matches the forcing, i.e.

$$q_2(x, \tau) = \hat{q}_2 A^2 e^{2ik_c x} + \text{c.c.} \quad (2.16)$$

Substituting this ansatz, Eq. (2.15a) can be solved for the complex constant $\hat{q}_2 = i/9k_c^3$.

$\epsilon^{3/2}$: The third-order equation is another inhomogeneous linear problem:

$$(\partial_t - \mathcal{L})q_3 = (1 - \partial_\tau)q_1 + \frac{1}{2}\partial_x q_1 q_2 \equiv f_3(x, \tau). \quad (2.17)$$

The q_1 forcing term at this order comes from the $\mathcal{O}(\epsilon)$ linear term in the full PDE acting on the $\mathcal{O}(\sqrt{\epsilon})$ term q_1 . As with the second-order forcing, the nonlinear term in f_3 can be interpreted as the interaction of the perturbation q_1 with the harmonic q_2 .

Expanding the lower-order terms q_1 and q_2 , the forcing is

$$f_3(x, \tau) = -\partial_\tau A e^{ik_c x} + \{ A e^{ik_c x} + 3ik_c \hat{q}_2 A^3 e^{3ik_c x} + ik_c \hat{q}_2 A |A|^2 e^{ik_c x} + \text{c.c.} \}. \quad (2.18)$$

Since the forcing once again does not depend on t , neither does q_3 , and Eq. (2.17) reduces to $-\mathcal{L}q_3 = f_3$.

It is not necessary to actually solve this third-order problem in order to derive an amplitude equation at leading order. Instead, the solvability condition from the Fredholm alternative finally gives an equation governing $A(\tau)$. In particular, the condition that the forcing is orthogonal to the nullspace of the adjoint is that

$$0 = \langle p_1 | f_3 \rangle = \int \hat{p}_1^* e^{-ik_c x} f_3(x, \tau) dx. \quad (2.19)$$


While this condition was automatically satisfied for the second-order equation since the forcing term was at a different wavenumber than p_1 , in this case the solvability condition is nontrivial for components of f_3 proportional to $e^{ik_c x}$. Retaining these components and dropping the arbitrary multiplicative constant, Eq. (2.19) reduces to

$$\frac{dA}{d\tau} = A - \mu A |A|^2 \quad (2.20)$$

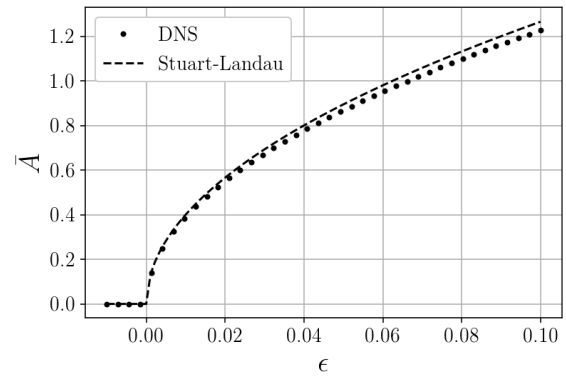
with $\mu = 1/9k_c^2$ after substituting the value of \hat{q}_2 from the second-order equation. The cubic term is a result of the nonlinear interaction between the fluctuation q_1 and the harmonic q_2 in the forcing term.

For small perturbations to the steady state $A = 0$, the cubic term is negligible and the perturbation grows exponentially. However, once A reaches finite amplitude, the nonlinearity becomes active and stabilizes the solution in a nonlinear equilibrium at $\bar{A} = 1/\sqrt{\mu} = 3k_c$, scaling with $\sqrt{\epsilon}$ (since this is the coefficient of A in the perturbation series solution q_ϵ). A typical interpretation of this behavior is that the amplitude A^2 of second-order harmonic is “slaved” to that of the fundamental mode. The quadratic interaction between the two will on average transfer energy from the unstable mode q_1 (for $\epsilon > 0$) to the dissipative harmonic until the two are in balance in the nonlinear equilibrium (see Fig. 2.4).

Equation (2.20) is known as a *Stuart-Landau equation* after the seminal work by Landau



(a) Solution of the supercritical Swift-Hohenberg equation, modeling convection rolls in Rayleigh-Bénard flow (see Fig. 2.2b)



(b) Bifurcation diagram for the primary instability comparing numerical simulation to the simplified amplitude equation (2.20)

Figure 2.4: The Swift-Hohenberg model. The approximate amplitude equation (2.20) predicts the qualitative symmetry-breaking and nonlinear stability of the model above the threshold of instability. The numerical simulation was performed with a pseudospectral Fourier method and a second-order Crank-Nicolson/Adams-Bashforth implicit-explicit timestepping scheme.

(1944); Landau & Lifshitz (1959) and Stuart (1958, 1960, 1967, 1971) on nonlinear stability theory. Although it was derived here from a simplified model equation, Stuart-Landau-type equations are ubiquitous in low-dimensional modeling of fluid flows and will play a central role through this work in different forms.

More broadly, the Swift-Hohenberg example illustrates many key features of fluid flows that are central to low-dimensional modeling, most importantly the general picture of large-scale instability, small-scale dissipation, and nonlinear interactions mediating energy transfer between them. In this example the stabilizing mechanism was the transfer of energy to higher harmonics, but there can be other interactions that contribute to the amplitude equations. For instance, when the steady solution is nonzero there can be a stabilizing contribution from the interaction of the fundamental mode and its conjugate. This effect is often described as a deformation of the base flow from an unstable steady state to a neutrally stable mean flow. Mean flow deformation will be discussed at length in the following, particularly in Chapters 5 and 9.

Although the effect of the cubic nonlinearity is stabilizing in this example, this is not necessarily the case (i.e. μ can be negative). In some important scenarios, including plane Poiseuille flow and thermoacoustic instabilities, the nonlinear term acts to further amplify fluctuations (Stewartson & Stuart, 1971; Orchini *et al.*, 2016; Ducimetière *et al.*, 2022). In this case, the bifurcation is termed *subcritical* (as opposed to the *supercritical* case discussed here) and finite amplitude perturbations can be significantly amplified even below the

threshold of linear instability.

For example, the steady convection rolls that result from the instability modeled here in Rayleigh-Bénard flow (e.g. Fig. 2.2b) are unstable to perturbations of sufficiently long wavelength. These fluctuations are amplified until the convection roll pattern breaks down completely and a new pattern of stable rolls eventually emerges. This phenomenon is known as the Eckhaus secondary instability (Charru, 2011). Although subcritical bifurcations are important in many flows, this work will primarily deal with the supercritical case in which effective cubic terms are stabilizing.

The method of solution via a perturbation series applied in this section is also known as weakly nonlinear analysis. It was pioneered in fluid dynamics by Stuart (1958, 1960) for highly symmetric geometries such as plane Poiseuille and Taylor-Couette flow, but more recently has been extended to more complex geometries and bifurcation scenarios with the aid of modern matrix-free computations (Sipp & Lebedev, 2007; Meliga *et al.*, 2009; Meliga & Chomaz, 2011; Orchini *et al.*, 2016; Rigas *et al.*, 2017; Ducimetière *et al.*, 2022). A similar asymptotic series expansion will also be used to develop a multiscale closure for reduced-order models in Chapter 7.

Finally, it is worth noting that besides reducing the partial differential equation (2.7) to the Stuart-Landau ordinary differential equation (2.20), the solution could also be expanded in multiple scales in both space and time. This approach, which is particularly important in the study of pattern formation, often leads to coarse-grained PDEs called Ginzburg-Landau equations, in which cubic terms appear by a similar process as outlined here (Stewartson & Stuart, 1971; Cross & Hohenberg, 1993; Charru, 2011).

2.3 Fluid flow as a nonlinear dynamical system

Section 2.2 presented an example of low-dimensional behavior arising in a partial differential equation near the onset of linear instability. The weakly nonlinear analysis reveals a picture of the nonlinear interactions transferring energy from the linearly unstable large scales to the smaller, dissipative scales. Since the small scales equilibrate much faster than the weakly unstable degrees of freedom, they can be eliminated from the dynamics, resulting in an effective reduction in dimensionality of the system. In this example, the PDE could be analytically reduced to a scalar-valued Stuart-Landau ODE by means of a perturbation series expansion.

Further from the threshold of instability this type of analysis tends to break down, but it suggests an alternative view of unsteady fluid flow as a large-scale dissipative dynamical system. When expanded about a fixed point \mathbf{x}_0 , such a system can be written in the general form

$$\dot{\mathbf{x}} = \mathbf{L}(\mathbf{x}_0; \boldsymbol{\mu})\mathbf{x} + \mathbf{N}(\mathbf{x}) \equiv \mathbf{f}(\mathbf{x}), \quad (2.21)$$

where $\mathbf{x} \in \mathbb{R}^N$ and in the case of the Navier-Stokes equations and similar fluid equations the dependence on the set of parameters $\boldsymbol{\mu}$ (for simplicity, we will take $\boldsymbol{\mu} = Re$) is restricted to the linear term \mathbf{L} . The discrete set variables $\mathbf{x}(t)$ (not to be confused with the spatial

coordinates) might come from a numerical discretization of the PDE system or projection onto a linear subspace, for example.

Again, for low values of the Reynolds number (e.g. $Re = \mathcal{O}(1)$), the fixed point \mathbf{x}_0 is stable to arbitrary perturbations. That is, solutions (λ, \mathbf{v}) to the eigenvalue problem

$$\lambda \mathbf{v} = \mathbf{L} \mathbf{v} \tag{2.22}$$

will have $\text{Re}\{\lambda\} < 0$. Since eigenvectors evolve with $e^{\lambda t}$, this corresponds to exponential decay in amplitude.

Above some critical Reynolds number Re_c (which may vary widely depending on the flow but is typically $\mathcal{O}(10^2 - 10^3)$), the system undergoes a *bifurcation*, or qualitative change in dynamical behavior. This is accompanied by a change in sign of the real part of some subset of the eigenvalues, and the bifurcation is categorized by the number of eigenvalues that simultaneously change sign and whether they are real or complex.

A wide variety of bifurcations have been identified since the study of nonlinear dynamics rose to prominence, but a few are particularly prevalent in fluid dynamics. One of these is the pitchfork bifurcation, in which a single real eigenvalue changes sign. This often accompanies a steady symmetry-breaking transition such as a wake deflection (Meliga *et al.*, 2009; Rigas *et al.*, 2015; Deng *et al.*, 2020) or Taylor vortices (Taylor, 1923). Both the onset of steady convection rolls in Rayleigh-Bénard flow (Fig. 2.2) and the simplified Swift-Hohenberg model of this instability (Fig. 2.4) are also examples of pitchfork bifurcations. Another common scenario is the Hopf bifurcation, in which a complex-conjugate pair of eigenvalues becomes unstable. This often indicates the onset of oscillations such as vortex shedding in the wake of a bluff body (Jackson, 1987) or spiral vortices in Taylor-Couette flow (Taylor, 1923).

The new equilibria reached above a bifurcation point generally often lose stability themselves in a secondary bifurcation when the control parameter is further increased. For example, if a Hopf bifurcation is supercritical, meaning that the dynamics settle on a stable, periodic limit cycle, this limit cycle can lose stability in a Neimark-Sacker bifurcation which generates quasiperiodic dynamics on a stable torus (Fig. 1.2). An early theory of the transition to turbulence proposed by Landau & Lifshitz (1959) was that an increasingly frequent series of these bifurcations would add incommensurate frequencies to the flow until a crisis point caused a breakdown into turbulence. However, Ruelle & Takens (1971) later argued on geometric grounds that following Hopf and Neimark-Sacker bifurcations in a generic dynamical system, another Hopf-type bifurcation would most likely *not* result in a new stable torus. Instead, the dynamics would typically break down into chaos, with the post-transient trajectories evolving on a complicated geometric structure called a *strange attractor* (of which the Lorenz “butterfly” attractor (Lorenz, 1963) is likely the most famous example). See Chapter 6 for an example of these secondary bifurcations and the route to chaos in shear-driven cavity flow.

The center manifold reduction Returning to the primary bifurcation, above the critical point the set of eigenvectors can be partitioned based on the growth rate into the *stable manifold*, spanned by damped eigenvectors, and the *unstable manifold*, spanned by the instability modes. At the critical point there is also a *center manifold*, spanned by eigenvectors with zero real value, as briefly discussed in the previous section. We can write the eigenvector bases for these manifolds as \mathbf{V}^s , \mathbf{V}^u , and \mathbf{V}^c , respectively.

Above the critical point, the system tends to be dominated by instability modes, since dynamics along the stable manifold are exponentially damped. Naively, one might attempt to reduce the order of the system by projecting the dynamics onto the span of the unstable eigenvectors. For instance, with the full eigenvector basis the linear term can be diagonalized with

$$\dot{z} = \Lambda z + \mathbf{V}^{-1} \mathbf{N}(\mathbf{V}z) \equiv \mathbf{f}^\Lambda(z), \quad (2.23)$$

where $\Lambda = \mathbf{V}^{-1} \mathbf{L} \mathbf{V}$ is the diagonal matrix of eigenvalues, sorted in order of decreasing real part and $z \in \mathbb{C}^N$ are generalized coordinates defined by $x = \mathbf{V}z$.

This system could then be truncated by discarding any elements z_j of z for which $\text{Re}\{\lambda_j\} < 0$. However, note that in the case of general nonlinearity, the full state x must be reconstructed in order to evaluate the nonlinear term. This motivates advanced model reduction methods like least-squares Petrov-Galerkin projection (Carlberg *et al.*, 2011), hyperreduction (Everson & Sirovich, 1995; Carlberg *et al.*, 2013), and “lifting” the variables to a space of observables in which the dynamics are bilinear (Qian *et al.*, 2020). Since we will primarily consider the quadratically nonlinear incompressible Navier-Stokes equation, these topics fall beyond the present scope.

Even for purely quadratic nonlinearity, the fundamental problem with this approach is that the projected dynamics do not resolve the critical mechanism of nonlinear energy transfer to higher orders, as illustrated by the weakly nonlinear expansion in Sec. 2.2. In fact, the standard POD-Galerkin modeling method, introduced in Chapter 4, is closely related to the truncation of Eq. (2.23). As will be shown, this approach often fails unless the model reduction is constructed to systematically account for the effect of the unresolved degrees of freedom.

One classical method for achieving this close to the threshold of bifurcation is the center manifold reduction (Coullet & Spiegel, 1983; Guckenheimer & Holmes, 1983; Wiggins, 1990; Carini *et al.*, 2015). In this case the system is marginally stable, so the dynamics are approximately restricted to the center manifold. We can then partition the state into the “active” and “stable” variables with $z = [z^c \ z^s]^T$ and $z \in \mathbb{C}^r$. However, rather than discarding the stable degrees of freedom, they are approximated with an algebraic *manifold equation* $z \approx \tilde{z} = \varphi(z^c)$. Substituting this into the partitioned dynamics for the stable variables and applying the chain rule gives the consistency condition (Guckenheimer & Holmes, 1983; Pavliotis & Stuart, 2012)

$$\mathbf{J}_\varphi(z^c) \dot{z}^c = \mathbf{f}^\Lambda(\varphi(z^c)), \quad (2.24)$$

where $\mathbf{J}_\varphi(\mathbf{z}^c)$ is the Jacobian of φ at \mathbf{z}^c . The original N -dimensional system has been reduced to order r by retaining only the “slow”, or marginally stable degrees of freedom.

In the case where the manifold equation is only used to approximate the fast variables $\tilde{\mathbf{z}} = [\mathbf{z}^c \ \varphi(\mathbf{z}^c)]^T$ and the the first r rows of the consistency condition are trivially satisfied, since the leading r rows of the Jacobian are the simply the identity matrix. The conditions corresponding to the slow variables become the $(N - r)$ -dimensional system of nonlinear partial differential equations for φ in the slow variables \mathbf{z}^c

$$(\varphi \cdot \nabla_c) \mathbf{f}^c(\tilde{\mathbf{z}}) = \mathbf{f}^s(\tilde{\mathbf{z}}), \quad (2.25)$$

where $\mathbf{f}^\Lambda = [\mathbf{f}^c \ \mathbf{f}^s]^T$ and the subscript on the del operator indicates that the gradient is taken with respect to the slow variables \mathbf{z}^c . This is not an obvious gain over the original system, but there are at least two approaches to rendering the consistency condition tractable.

The first is the classical method of solution in dynamical systems theory. Close to the bifurcation the state variables are typically small, so the manifold equation can be expanded in a power series. Substituting the power series expansion into Eq. (2.25) and equating powers of \mathbf{z}^c reduces the PDE to an algebraic linear system for the expansion coefficients. Once this is solved, the reduced-order model can be calculated explicitly (See the example below and Sec. 5.4 for examples of this computation). An example of this procedure for a three-dimensional model problem is given below; for more details see standard dynamical systems references such as Guckenheimer & Holmes (1983); Wiggins (1990). An alternative perspective on the consistency condition in terms of operator theory will also be presented in Section 2.4.

Even for systems that are already very low-dimensional this procedure is not trivial, and the algebra involved quickly becomes tedious as the dimension of the state space grows. Alternatively, if the Jacobian \mathbf{J}_φ is available, the consistency condition (2.24) can be solved as an algebraic system of equations for $\dot{\mathbf{z}}^c$ at each timestep, i.e.

$$\mathbf{J}_\varphi(\mathbf{z}^c) \dot{\mathbf{z}}^c = \mathbf{f}^\Lambda(\varphi(\mathbf{z}^c)). \quad (2.26)$$

With modern matrix-free iterative solvers the Jacobian itself need not even be formed to solve for $\dot{\mathbf{z}}^c$, provided the Jacobian-vector product can be efficiently computed. See for example Sec. 6.6 for an application of this approach using a deep neural network to approximate φ and automatic differentiation to compute the Jacobian.

Although this approach is automated, a significant drawback to this solution method is that it requires forming the N -dimensional right-hand side \mathbf{f}^Λ at each time step, so that it still scales with the dimension of the original state. It is therefore more appropriately used as a stabilization of moderately-sized systems rather than a primary model reduction method. For instance, Chapter 6 proposes applying this technique as a secondary model reduction of POD-Galerkin systems, so that N is already low-dimensional.

Example manifold reduction in a toy problem As an instructive example of the power series approach to center manifold reduction, consider the three-dimensional “mean-field” model modified from the example introduced by Noack *et al.* (2003) to model mean flow deformation in the cylinder wake (see also Chapter 5):

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} \epsilon & -\omega & 0 \\ \omega & \epsilon & 0 \\ 0 & 0 & -1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + \begin{bmatrix} -(\mu x_1 - \nu x_2)x_3 \\ -(\mu x_2 + \nu x_1)x_3 \\ x_1^2 + x_2^2 \end{bmatrix}. \quad (2.27)$$

The linear operator has eigenvalues $\lambda_{1,2} = \epsilon \pm i\omega$ and $\lambda_3 = -1$, so in this case the system is already partitioned into the slow variables x_1, x_2 and fast variable x_3 and diagonalization is not necessary.

For $\epsilon < 0$ all three eigenvalues have negative real part, so the origin is stable. As ϵ crosses zero, the leading pair of eigenvalues loses stability with complex-conjugate frequency, indicating a Hopf bifurcation. The preceding discussion suggests that near $\epsilon = 0$ the stable variable x_3 can be approximated as an algebraic function $x_3 \approx \varphi(x_1, x_2)$. Expanding φ in a power series with unknown constants c ,

$$\varphi(x_1, x_2) \approx c^{(0)} + c_1^{(1)} x_1 + c_2^{(1)} x_2 + c_1^{(2)} x_1^2 + c_2^{(2)} x_1 x_2 + c_3^{(2)} x_2^2 + \dots \quad (2.28)$$

The constant and linear terms $c^{(0)}$ and $c_{1,2}^{(1)}$ can be eliminated based on requiring that $\varphi(\mathbf{0}) = 0$ and $\mathbf{J}_\varphi(\mathbf{0}) = \mathbf{0}$, or in other words that the manifold equation is tangent to the center manifold at the origin.

Since we have not modified the slow variables x_1, x_2 , the consistency condition only applies to the fast variable, i.e. $\dot{\varphi}(x_1, x_2) = f_3(x_1, x_2, \varphi(x_1, x_2))$. Applying the chain rule to the left-hand side time derivative at $\epsilon = 0$ and dropping the superscript (2) since only quadratic terms remain,

$$\frac{\partial \varphi}{\partial x_1} f_1(x_1, x_2, \varphi(x_1, x_2)) + \frac{\partial \varphi}{\partial x_2} f_2(x_1, x_2, \varphi(x_1, x_2)) = f_3(x_1, x_2, \varphi(x_1, x_2)) \quad (2.29)$$

$$\begin{aligned} & (2c_1 x_1 + c_2 x_2) (-\omega x_2 + \mathcal{O}(x^3)) + (c_2 x_1 + 2c_3 x_2) (\omega x_1 + \mathcal{O}(x^3)) \\ & = x_1^2 + x_2^2 - (c_1 x_1^2 + c_2 x_1 x_2 + c_3 x_2^2). \end{aligned} \quad (2.30)$$

The first equation is the consistency condition, which reduces to an algebraic equation, from which terms proportional to x_1^2 , $x_1 x_2$, and x_2^2 can be isolated to form a system of three equations in three unknowns:

$$\begin{bmatrix} 1 & \omega & 0 \\ -2\omega & 1 & 2\omega \\ 0 & -\omega & 1 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}. \quad (2.31)$$

The solution $\mathbf{c} = [1 \ 0 \ 1]^T$ may be determined by inspection, leading to the parabolic

manifold equation $x_3 \approx x_1^2 + x_2^2$. Finally, the reduced system for small ϵ is the linear-cubic nonlinear oscillator

$$\dot{x}_1 = \epsilon x_1 - \omega x_2 - (\mu x_1 - \nu x_2)(x_1^2 + x_2^2) \quad (2.32a)$$

$$\dot{x}_2 = \epsilon x_2 + \omega x_1 - (\mu x_2 + \nu x_1)(x_1^2 + x_2^2). \quad (2.32b)$$

In terms of the leading eigenvector $z = (x_1 + ix_2)/2$, this simplifies further to the Stuart-Landau equation

$$\dot{z} = (\epsilon + i\omega)z - (\mu + i\nu)z|z|^2. \quad (2.33)$$

This equation is also known as the *normal form* for a Hopf bifurcation, since it is the simplest generic model for any system near the threshold of such a bifurcation, including for instance oscillatory vortex shedding instabilities in fluid flow (see Chapter 5). If $\mu > 0$, the nonlinear term will stabilize the energy of the system in an equilibrium similar to the case derived via weakly nonlinear analysis in Sec. 2.2, although μ may also be negative in the case of a subcritical bifurcation. The imaginary term ν in the nonlinearity indicates amplitude-dependent modulation of the base frequency ω .

The algebra involved in the center manifold reduction scales rapidly with the system size, even for this highly simplified case. Although the full derivation is given here to illustrate the concept, the original presentation by Noack *et al.* (2003) gives a much more straightforward approach to eliminating x_3 by a method known as *adiabatic elimination*. If we note that the nonlinear forcing $x_1^2 + x_2^2$ in the x_3 dynamics is independent of the phase and that the linear damping acts on a much faster time scale than the instability ϵ , we may directly derive the manifold equation by setting $\dot{x}_3 = 0$ on the slow time scale so that $x_3 = x_1^2 + x_2^2$. This is closely related to the stochastic averaging procedure introduced in Chapter 7; the two will be compared at greater length in Sec. 7.1.

Finally, the center manifold reduction will be examined further in the model of a cylinder wake presented in Chapter 5, including a higher-order data-driven manifold model based on sparse polynomial regression that does not require solving the consistency condition. The nonlinear correlations analysis and deep learning stabilization methods introduced in Chapter 6 can also be viewed as generalizations of the center manifold reduction; these concepts will be revisited frequently throughout this work.

2.4 Generators of deterministic and stochastic processes

In the study of mechanics there are often several equivalent representations of the same system (e.g. Lagrangian, Hamiltonian, evolution equations, etc.) that may variously be useful depending on the application. One of these, which will be particularly helpful in the development of multiscale closure modeling in Chapter 7 and nonlinear stochastic system identification in Chapter 8, is the operator theoretic perspective. This section gives a brief overview of topics necessary for the development of these closure models; for a more in-depth presentation of these topics see references such as Risken (1996); Pavliotis

& Stuart (2012); Klus *et al.* (2020); Brunton *et al.* (2022). Perhaps the most widely-discussed of the operator-theoretic methods in recent work on the analysis of fluid flows and other large-scale nonlinear dynamics is Koopman theory, making it a convenient place to begin the discussion.

As in Sec. 2.3, suppose the state $\mathbf{x}(t)$ of a system is governed by the autonomous (ODE)

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}), \quad \mathbf{x} \in \mathcal{X}. \quad (2.34)$$

We will assume unless otherwise specified that states are real-valued, e.g. $\mathcal{X} = \mathbb{R}^N$, and that all functions (e.g. p, q) are L^2 -integrable with inner product

$$\langle p|q \rangle = \int_{\mathcal{X}} p^*(\mathbf{x})q(\mathbf{x})d\mathbf{x}. \quad (2.35)$$

As usual, the adjoint of a linear operator \mathcal{A} is denoted \mathcal{A}^\dagger and defined by the relationship

$$\langle p|\mathcal{A}q \rangle = \langle \mathcal{A}^\dagger p|q \rangle. \quad (2.36)$$

The Koopman operator \mathcal{K}^t acts on the L^2 -integrable space of scalar observables $g(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$, advancing them forward time t . More precisely, let \mathbf{x} be the solution to the initial value problem of (2.34) with $0 \leq t \leq T$ and $\mathbf{x}(0) = \mathbf{x}_0$. Then the Koopman operator is defined as

$$\mathcal{K}^T g(\mathbf{x}_0) = g(\mathbf{x}(T)). \quad (2.37)$$

Although the Koopman operator is generally difficult to either represent explicitly or approximate in a useful finite-dimensional subspace, analysis of its spectral properties has drawn great interest in recent work. For our purposes, the infinitesimal generator \mathcal{L} , defined by

$$\mathcal{L}g = \lim_{t \rightarrow 0} \frac{\mathcal{K}^t g - g}{t} \quad (2.38)$$

and sometimes called the Lie operator, is more theoretically useful. If we consider $g(\mathbf{x}(t))$ to be an explicit function of time, then \mathcal{L} can be derived by applying the chain rule to g . With a slight abuse of notation, if $g(\mathbf{x}, t)$ is taken to be a function of both time and initial state \mathbf{x} (i.e. in the Lagrangian frame of reference), then this can be extended to a PDE over all of state space:

$$\frac{\partial g}{\partial t} = \mathcal{L}g = \mathbf{f}(\mathbf{x}) \cdot \nabla g. \quad (2.39)$$

Thus the generator \mathcal{L} is a linear advection operator governing the evolution of scalar observables of the system described by (2.34). Comparing Eqs. (2.39) and (2.24), each entry in the consistency condition (2.24) can be viewed as a Lie operator governing the evolution of the fast variables as approximated by the manifold approximation φ . In other words, the approximated fast variables are the observables g whose evolution is determined by the slow dynamics \mathbf{f} according to Eq. (2.39).

The adjoint of Eq. (2.39),

$$\frac{\partial \rho}{\partial t} = \mathcal{L}^\dagger \rho = -\nabla \cdot (\rho \mathbf{f}(\mathbf{x})), \quad (2.40)$$

is a continuity equation governing the evolution of densities $\rho(\mathbf{x}, t)$ in phase space. As a point of reference, equation (2.40) reduces to the Liouville equation from classical mechanics under the incompressibility condition $\nabla \cdot \mathbf{f} = 0$, and the adjoint \mathcal{L}^\dagger is also known as the generator of the Perron-Frobenius operator.

This description of dynamics can readily be extended to systems governed by stochastic differential equations (SDEs) of the form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \Sigma \mathbf{w}(t), \quad (2.41)$$

where the deterministic component $\mathbf{f}(\mathbf{x})$ is known as the drift function and the diffusion Σ modifies a vector-valued Gaussian white noise process $\mathbf{w}(t)$ ². In this case the evolution of the probability distribution $\rho(\mathbf{x}, t)$ is governed by the stochastic analog of (2.40), known as the forward Kolmogorov, or Fokker-Planck, equation:

$$\frac{\partial \rho}{\partial t} = \mathcal{L}^\dagger \rho = -\nabla \cdot (\rho \mathbf{f}) + \nabla \nabla^T : (\rho \mathbf{D}), \quad (2.42)$$

where $\mathbf{D}(\mathbf{x}) = \Sigma \Sigma^T / 2$ is the diffusion tensor and the colon indicates tensor contraction. The associated backwards Kolmogorov equation is the adjoint of Eq. (2.42):

$$\frac{\partial g}{\partial t} = \mathcal{L} g = \mathbf{f} \cdot \nabla g + \mathbf{D} : \nabla \nabla^T g. \quad (2.43)$$

To interpret Eq. (2.43), consider the expectation of a scalar observable $g(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$ defined by the inner product of g with the probability distribution $\rho_0(\mathbf{x})$:

$$\mathbb{E}[g] = \langle \rho_0 | g \rangle = \int_{\mathcal{X}} \rho_0(\mathbf{x}) g(\mathbf{x}) d\mathbf{x}. \quad (2.44)$$

The probability distribution is advanced in time by $\rho(\mathbf{x}, t) = e^{\mathcal{L}^\dagger t} \rho_0(\mathbf{x})$. Using the definition of the adjoint from Eq. (2.36), the expectation of g at time t can be written equivalently as

$$\mathbb{E}[g](t) = \langle \rho | g \rangle = \int_{\mathcal{X}} \rho_0(\mathbf{x}) e^{\mathcal{L} t} g(\mathbf{x}) d\mathbf{x}. \quad (2.45)$$

Therefore, just as the Koopman operator advances an observable in time, the backwards Kolmogorov equation describes the evolution of the expectation of an observable.

²State-dependent diffusion is also possible; see Chapter 8

The Ornstein-Uhlenbeck process As a simple example relevant for the stochastic modeling methods in 7 and Chapters 8, consider the linear stochastic *Ornstein-Uhlenbeck* process defined by the SDE

$$\dot{x} = \nu(\mu - x) + \sigma w(t) \quad (2.46)$$

for positive constants ν , μ , and σ . The stationary distribution $\rho^\infty(x)$ is given by the steady-state Fokker-Planck equation

$$0 = \mathcal{L}^\dagger \rho^\infty = -\frac{d}{dx} \rho^\infty(x) \nu(\mu - x) + \frac{\sigma^2}{2} \frac{d^2}{dx^2} \rho^\infty(x), \quad (2.47)$$

along with the usual normalization condition on ρ^∞ . This can be solved analytically with a Gaussian distribution with mean μ and variance $\sigma^2/2\nu$. Defining an observable that is the state itself $g(x) = x$, the evolution is given by the backwards Kolmogorov equation (2.43), which in this case reduces to the initial value problem

$$\frac{dg}{dt} = \nu(\mu - g), \quad g(0) = x. \quad (2.48)$$

The expectation at time t is the solution $g(t)$, which in this case is exponential decay towards the mean μ :

$$\mathbb{E}[x](t) = g(t) = \mu - (x - \mu)e^{-\nu t}. \quad (2.49)$$

Given that the model reduction procedures described throughout this work generally aim to reduce the physics model from an infinite-dimensional partial differential equation to a finite-dimensional system of ordinary differential equations, it may be counterintuitive that it would be helpful to return to an infinite-dimensional function space and represent the dynamics with a partial differential equation. The primary advantage in doing so is that these generators are linear operators, which in some cases can enable approaches that are unavailable for the nonlinear dynamics of the ODE (2.34) or SDE (2.41). Nevertheless, aside from simple one-dimensional or linear examples like the Ornstein-Uhlenbeck process, it is typically very difficult to solve these PDEs analytically, motivating the approximation in terms of an asymptotic expansion in Chapter 7.

2.5 Summary

Unsteady low-speed fluid flow is governed by the incompressible Navier-Stokes equations, a set of nonlinear partial differential equations that is notoriously challenging to approach analytically or computationally. The main purpose of this chapter, besides introducing mathematical background that will be useful throughout, is to motivate the treatment of unsteady fluid flow as a high-dimensional, dissipative, nonlinear dynamical system with latent low-dimensional structure.

Close to the threshold of instability, the effective dimensionality of the system can be

reduced considerably by approximating all but the marginally stable spatial structures as asymptotically strongly damped and modeling the evolution of the instability modes on a slow time scale. This analytic approach begins to lose validity further from the critical point, but experimental and theoretical evidence suggests that even high-Reynolds number turbulence continues to be dominated by relatively low-dimensional coherent structures that evolve relatively slowly and regularly compared to the fast turbulent fluctuations.

Roughly speaking, the working picture of fluid dynamics presented here is that of large-scale instability and small-scale dissipation, with energy transfers between the scales mediated by nonlinear interactions. The objective of the model reduction methods discussed in this work is to model the highly energetic large-scale motions while approximating the effect of the unresolved small scales in a consistent manner, with ultimate goals including describing phenomena such as the routes to chaos and turbulence, and developing compact approximations suitable for real-time feedback control.

The mathematical tools presented in this chapter, including dynamical systems analysis, perturbation series expansions, and operator theory, are central to these objectives. The following three chapters will introduce further background material necessary for the applications in the latter half of this work, although for the most part these will be significantly less abstract and technical than the current chapter. Chapter 3 will introduce several relevant topics in modal analysis that enable identification of the dominant large-scale structures in the case that intrusive methods such as weakly nonlinear analysis are not practical. Once the large-scale modes are identified, Chapters 4 will deal with the problem of approximating their dynamics, either with physics-based projection methods or data-driven model discovery.

Chapter 3

MODAL ANALYSIS

Linear modal analysis is a powerful tool for extracting low-dimensional coherent structure in flows, even those characterized by strong nonlinearity; see for instance Taira *et al.* (2017, 2020) for a comprehensive survey. We will approximate the flow fields \mathbf{q} with truncated (rank- r) affine decompositions of the form

$$\mathbf{q}(\mathbf{x}, t) \simeq \mathbf{q}_0(\mathbf{x}) + \sum_{k=1}^r \boldsymbol{\psi}_k(\mathbf{x}) a_k(t), \quad (3.1)$$

including for example modal stability analysis (Theofilis, 2011), proper orthogonal decomposition (Lumley, 1967; Berkooz *et al.*, 1993; Holmes *et al.*, 1996), and dynamic mode decomposition (Schmid, 2010; Rowley *et al.*, 2009), but excluding approaches such as non-modal stability analysis (Schmid, 2007; McKeon & Sharma, 2010) and spectral proper orthogonal decomposition (Towne *et al.*, 2018). Broadly speaking, the goal of modal analysis is to identify a suitable basis $\{\boldsymbol{\psi}_k\}_{k=1}^r$ in which the flow kinematics can be represented, while the reduced-order dynamical systems models discussed in Chapter 4 treat the time evolution of the coefficients $\mathbf{a}(t)$. Since the approximate state is specified by the r -dimensional coefficient vector, equation (3.1) is a linear dimensionality reduction.

Depending on the context, the state \mathbf{q} may either consist of only the velocity field, or the concatenated velocity and pressure fields, i.e. $\mathbf{q} = [\mathbf{u} \ p]^T$. Usually the inner product between states is based on the velocity:

$$\langle \mathbf{q}_1 | \mathbf{q}_2 \rangle = \int_{\Omega} \mathbf{u}_1^* \cdot \mathbf{u}_2 \, d\mathbf{x}, \quad \mathbf{x} \in \Omega, \quad (3.2)$$

so that if $\mathbf{q}(\mathbf{x}, t) = \bar{\mathbf{q}}(\mathbf{x}) + \mathbf{q}'(\mathbf{x}, t)$ for mean flow $\bar{\mathbf{q}}(\mathbf{x})$ and the density ρ is constant, then $\langle \mathbf{q}' | \mathbf{q}' \rangle$ is proportional to the fluctuation kinetic energy

$$K(t) = \frac{\rho}{2} \int_{\Omega} \mathbf{u}'(\mathbf{x}, t) \cdot \mathbf{u}'(\mathbf{x}, t) \, d\mathbf{x}. \quad (3.3)$$

Since the data is often on a nonuniform mesh, Eq. (3.2) can be approximated with a weighted sum

$$\langle \mathbf{q}_2, \mathbf{q}_1 \rangle \approx \mathbf{u}_1^T \mathbf{W} \mathbf{u}_2, \quad (3.4)$$

where \mathbf{u} is the discrete approximation to \mathbf{u} , and \mathbf{W} is a diagonal weight matrix containing

the cell volumes (or lumped mass matrix, in a finite element context). The transpose in Eq. (3.4) is understood to include complex conjugation in the case that the fields or modes are complex-valued.

Many of the numerical algorithms for modal analysis assume that the state can be written as a vector, or one-dimensional array. If the state includes the pressure, the energy inner product can be implemented by defining the discretized state vector $\mathbf{q} = [\mathbf{u} \ p]^T$ and setting the diagonal entries of \mathbf{W} corresponding to pressure to zero. Then $\mathbf{q}_1^T \mathbf{W} \mathbf{q}_2$ is equivalent to Eq. (3.4) and the inner product has the same interpretation in terms of kinetic energy, but the associated pressure fields are automatically carried through the computations.

Modal analysis methods can generally be classified into two groups: empirical and operator-based. Operator-based methods are rooted in the analysis of the governing equations, while empirical methods are based on statistical analysis of some collection of experimental or numerical data. The two categories are complementary; operator-based methods tend to be more detailed, mechanistic, and predictive, but require discretized approximations of the governing equations and are often restricted to a limited regime of validity (e.g. linearity or proximity to a bifurcation). On the other hand, empirical methods give a “true” statistical picture of the flow in any condition and are non-intrusive, but may not generalize under modifications of the geometry or flow conditions.

This chapter gives a brief overview of three modal analysis techniques that will play major roles in this work. Section 3.1 introduces linear stability analysis, the only operator-based method discussed at length here, followed by proper orthogonal decomposition (POD) in Section 3.2 and the more recent dynamic mode decomposition (DMD) in Section 3.3.

3.1 Linear stability analysis

Due to the difficulties in treating nonlinearity analytically, one of the foundational tools in the study of dynamics across physical domains is linear stability analysis. Assuming an equilibrium solution $\mathbf{q}_0(\mathbf{x})$ is known that satisfies the steady-state Navier-Stokes equations (2.4), linear stability theory is concerned with the fate of perturbations $\mathbf{q}'(\mathbf{x}, t)$ to the base flow $\mathbf{q}_0(\mathbf{x})$. Such a steady state is often difficult to determine analytically but Eq. (2.4) can be solved numerically, for instance with a Newton-Krylov (Knoll & Keyes, 2004) or selective frequency damping (Åkervik *et al.*, 2006) method.

Substituting the perturbed base flow into the Navier-Stokes momentum equations (2.1) and neglecting the quadratic term $\nabla \cdot (\mathbf{u}' \otimes \mathbf{u}')$, we find the linearized Navier-Stokes equations

$$\frac{\partial \mathbf{u}'}{\partial t} + \nabla \cdot (\mathbf{u}_0 \otimes \mathbf{u}' + \mathbf{u}' \otimes \mathbf{u}_0) = -\nabla p' + \frac{1}{Re} \nabla^2 \mathbf{u}' \quad (3.5a)$$

$$\nabla \cdot \mathbf{u}' = 0. \quad (3.5b)$$

The linearized Navier-Stokes equations are parameterized by Reynolds number, and can be used to investigate the stability properties of various configurations; if any of the eigenvalues λ_n have positive real part then the energy of the corresponding modes is predicted to grow exponentially. This exponential growth is likely the reason that linearly unstable structures such as Tollmien-Schlichting waves, Rossiter modes, Rayleigh-Bénard rolls, and von Kármán vortex streets are often commonly observed, even in the fully non-linear regime. This is also one reason why linear theory can be successfully used in flow control schemes (Rowley & Williams, 2006; Bagheri *et al.*, 2009a; Sipp *et al.*, 2010; Brunton & Noack, 2015b).

Introducing a normal mode ansatz

$$\mathbf{q}'(\mathbf{x}, t) = \sum_{n=1}^{\infty} \hat{\mathbf{q}}_n(\mathbf{x}) e^{\lambda_n t},$$

the linearized Navier-Stokes equations can be recast as a generalized eigenvalue problem

$$\lambda_n \begin{bmatrix} \mathcal{I} & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{u}}_n \\ \hat{p}_n \end{bmatrix} = \begin{bmatrix} -\nabla \cdot (\mathbf{u}_0 \otimes (\cdot) + (\cdot) \otimes \mathbf{u}_0) + Re^{-1} \nabla^2 & -\nabla \\ \nabla^T & 0 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{u}}_n \\ \hat{p}_n \end{bmatrix}. \quad (3.6)$$

Under most two- or three-dimensional discretizations, this becomes a very large, sparse system of equations that is difficult to solve with standard (dense) linear algebra methods. Stability analyses are therefore usually classified as either *local* if the base flow is parallel and Eq. (3.6) can be simplified to a one-dimensional problem (discussed further below), or *global* if the full multidimensional problem is solved. Historically, the field of hydrodynamic stability analysis developed around local or weakly nonparallel analyses (Schmid & Henningson, 2001; Drazin & Reid, 2004; Charru, 2011) since the global problem was computationally intractable, although the global problem is likely more recognizable to those with a background in linear systems theory.

However, recent decades have seen not only the development of significantly greater computational power, but also of advanced matrix-free linear algebra methods for solving Eq. (3.6) and similar problems (Edwards *et al.*, 1994; Stewart, 2001; Antoulas, 2005). When global stability analysis is presented in this work, the results have been computed using a Krylov-Schur time-stepping algorithm implemented in the spectral element solver Nek5000 (Fischer *et al.*, 2008). The details of these matrix-free eigenvalue solvers are not important for present purposes, but for more details see Antoulas (2005); Bagheri *et al.* (2009b); Sipp *et al.* (2010), or the recent review chapter Loiseau *et al.* (2019).

Convective instability and transient energy growth Although recent algorithmic developments have made fully three-dimensional global stability analysis possible, the framework of local stability analysis is still theoretically important for reduced-order modeling, not least because of the distinction it makes between two fundamental types of hydrody-

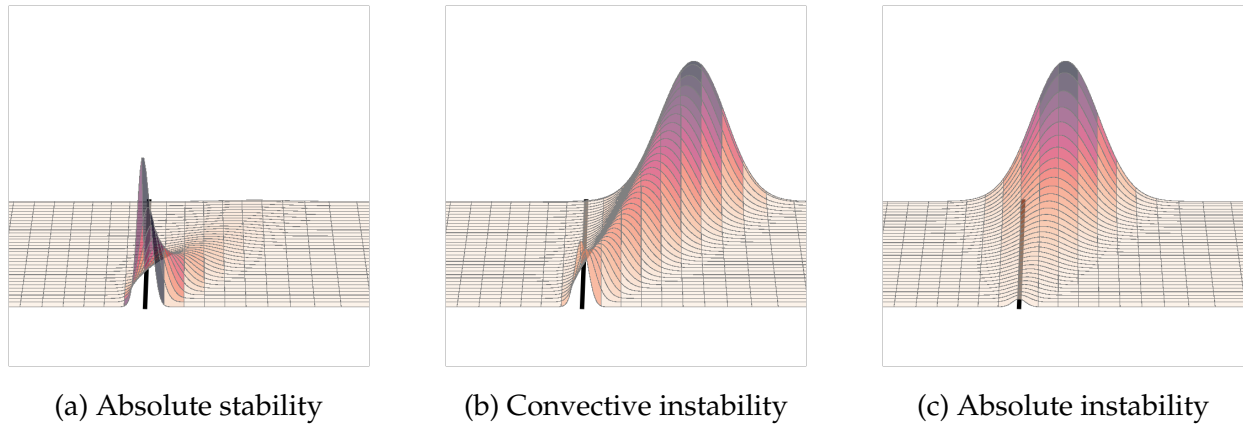


Figure 3.1: Types of linear instability in a model of shear flow While disturbances decay exponentially in absolutely stable flow (a), the distinction between convective and absolute instability is a matter of Galilean reference frame. Perturbations are amplified as they travel downstream in convective instability (b), while grow at a fixed point in space for absolute instability (c). The plots are adapted from Bagheri *et al.* (2009b) and show a modified linear Ginzburg-Landau equation.

namic instability.

Consider a two-dimensional flow with a *parallel* base flow, i.e. $\mathbf{u}_0(x, y) = U(y)\hat{\mathbf{i}}$, where $\hat{\mathbf{i}}$ is the unit vector in the x -direction. This includes boundary layers, free shear layers, channel flows, and, in radial coordinates, pipe flow and Taylor-Couette flow between counter-rotating cylinders. If the x -coordinate extends to infinity, streamwise variations can be represented with Fourier modes. The infinitesimal disturbances $\mathbf{q}'(x, t)$ then take the wave-like form

$$\mathbf{q}'(x, y, t) = \hat{\mathbf{q}}(y)e^{ik(x-ct)} \quad (3.7)$$

for real wavenumber k and complex speed $c = c_r + ic_i$ ¹. There will be a different eigenvalue problem for each wavenumber k , each with eigenvalue c . Comparing to the modal ansatz for the global problem, the temporal growth rate is then $\sigma = kc_i$, with oscillation frequency $\omega = kc_r$. As in the global case, when $\sigma < 0$ any perturbations will decay exponentially.

The critical distinction is between *absolute* and *convective* instability, as introduced by Monkewitz & Huerre (1982); Huerre & Monkewitz (1985) and illustrated in Fig. 3.1 for a one-dimensional model of shear flow. In absolute instability $\sigma > 0$ for a mode with zero group velocity c_r and the instability will grow at a fixed point in space (Fig. 3.1c).

¹Alternatively, a stability analysis of a parallel flow can be framed a spatial problem with modes growing or decaying in space rather than in time, in which case the Fourier modes are constructed so that the wavenumber is complex and the oscillation frequency is real.

If $c_r \neq 0$ for all modes with $\sigma > 0$, then the instability is convected downstream as it is amplified, but it will at any fixed spatial location (Fig. 3.1b). Absolutely and convectively unstable flows are also sometimes known as “oscillator” and “amplifier” flows, respectively, due to their tendency to either reach an oscillatory nonlinear equilibrium state or exhibit transient amplification of small disturbances.

For reduced-order models based on global modes (whether or not the modes were derived from a global stability analysis), this distinction is also important. Since perturbations eventually decay to zero at all fixed points in space for convectively unstable flows, these flows will be identified as linearly stable in a global analysis. Instead, the convective instability manifests as strong non-normality in the linear operator, which implies that perturbations may have significant algebraic energy growth before decaying exponentially (Trefethen *et al.*, 1993; Reddy *et al.*, 1993; Chomaz, 2005). In this case a nonmodal analysis based on the “pseudospectrum” of the operator (also known as a resolvent analysis) may be more appropriate than the traditional modal stability analysis (Schmid & Henningson, 2001; Schmid, 2007; McKeon & Sharma, 2010; McKeon *et al.*, 2013; Symon *et al.*, 2018, 2019).

Mean flow stability Although linear stability analysis is often conducted in an expansion around a steady solution of the Navier-Stokes equations, this does not need to be the case. In fact, there is increasing interest in the role that the *mean* flow can play in stability analysis.

Although it is not a steady-state solution of the Navier-Stokes equations, several studies have shown that linearizing about the mean flow $\bar{u}(\mathbf{x})$ nonetheless provides valuable insights into the dynamics of coherent structures existing in the nonlinear flow (Malkus, 1956; Barkley, 2006; Mantič-Lugo *et al.*, 2014; Beneddine *et al.*, 2016; Meliga, 2017). The analysis is the same as before, replacing the base flow $\mathbf{q}_0(\mathbf{x})$ with the mean flow $\bar{\mathbf{q}}(\mathbf{x})$. In some cases, a mean flow expansion can predict the structure and frequencies of the dominant modes in the nonlinear flow much more accurately than a steady-state expansion.

This improvement of the stability analysis about the mean flow accounts for its increasing popularity in both modal (Sipp & Lebedev, 2007; Beneddine *et al.*, 2016) and nonmodal (McKeon & Sharma, 2010) analysis. It is also appealing experimentally, since the mean flow can be estimated practically for statistically stationary flows, while unstable steady states are difficult to produce. However, from a numerical perspective standard mean flow analysis is not predictive in the sense that fully-converged statistics are necessary to compute the mean flow prior to the stability analysis. Predictive mean flow analysis is the subject of ongoing work, for example with eddy viscosity-based Reynolds-averaged Navier-Stokes mean flow estimates (Pickering *et al.*, 2021) or self-consistent modeling (Mantič-Lugo *et al.*, 2014; Mantič-Lugo & Gallaire, 2016; Meliga, 2017).

From the perspective of reduced-order modeling, this has several consequences. In particular, even in an empirical analysis the expansion could be conducted about either

the steady base flow or the mean. This can have ramifications for the accuracy of the model; while a base flow expansion can allow the model to resolve the stabilizing mechanism of mean flow deformation (Noack *et al.*, 2003), the base flow is not necessarily relevant for strongly nonlinear flows. If the Reynolds stresses in Eq. (2.6) are strong enough to significantly deform the base flow and the unsteady solution rarely passes near the unstable steady state, a mean flow expansion may be a better choice.

3.2 Proper orthogonal decomposition (POD)

One of the most widely used techniques for dimensionality reduction and modal analysis is proper orthogonal decomposition (POD), which solves the optimization problem

$$\min_{\{\psi_j\}} \mathbb{E} \left\| \mathbf{q}' - \sum_{j=1}^r \langle \psi_j | \mathbf{q}' \rangle \psi_j \right\|^2 \quad \text{subject to } \langle \psi_j | \psi_k \rangle = \delta_{jk} \quad (3.8)$$

assuming an expansion of the form (3.1), where the expectation is an ensemble average, approximated in practice by a time average, and δ_{jk} is the Kronecker delta. The objective is to minimize the residual energy in a linear subspace of r orthonormal modes, providing an optimal low-rank representation of the flow.

This problem can be solved with the calculus of variations as follows. Replacing the ensemble average with a time average under the assumption of ergodicity and expanding the norm in Eq. (3.8), the problem becomes

$$\min_{\{\psi_j\}} \frac{1}{T} \int_0^T \left\| \mathbf{q}'(\mathbf{x}, t) - \sum_{j=1}^N \langle \psi_j | \mathbf{q}' \rangle \psi_j \right\|^2 dt \quad (3.9)$$

$$= \min_{\{\psi_j\}} \frac{1}{T} \int_0^T \left\{ \langle \mathbf{q}' | \mathbf{q}' \rangle - 2 \sum_{j=1}^N |\langle \psi_j | \mathbf{q}' \rangle|^2 + \sum_{j,k=1}^N \langle \psi_j | \mathbf{q}' \rangle \langle \psi_k | \mathbf{q}' \rangle \langle \psi_j | \psi_k \rangle \right\} dt \quad (3.10)$$

$$= \min_{\{\psi_j\}} - \sum_{j=1}^N \frac{1}{T} \int_0^T |\langle \psi_j | \mathbf{q}' \rangle|^2 dt \quad \text{subject to } \langle \psi_j | \psi_k \rangle = \delta_{jk}, \quad (3.11)$$

where the last step takes advantage of orthonormality of the basis functions, and removes the norm $\langle \mathbf{q}' | \mathbf{q}' \rangle$ (which doesn't depend on the basis functions).

This constrained optimization problem can be solved to find some particular mode ψ by introducing the Lagrangian functional

$$\mathcal{L}[\psi, \lambda] = -\frac{1}{T} \int_0^T |\langle \psi | \mathbf{q}' \rangle|^2 dt + \lambda (||\psi||^2 - 1) \quad (3.12)$$

with Lagrange multiplier λ . Minima of this cost functional satisfy

$$\left. \frac{d}{d\epsilon} \mathcal{L}[\boldsymbol{\psi} + \epsilon \boldsymbol{\varphi}, \lambda] \right|_{\epsilon=0} = 0 \quad (3.13)$$

for any function $\boldsymbol{\varphi}(\boldsymbol{x})$ in the same space as $\boldsymbol{\psi}(\boldsymbol{x})$. For real-valued data $\boldsymbol{q}'(\boldsymbol{x}, t)$, modes $\{\boldsymbol{\psi}(\boldsymbol{x})\}$, and test function $\boldsymbol{\varphi}(\boldsymbol{x})$, this optimality condition evaluates to

$$\frac{1}{T} \int_0^T \langle \boldsymbol{\varphi} | \boldsymbol{q}' \rangle \langle \boldsymbol{q}' | \boldsymbol{\psi} \rangle dt - \lambda \langle \boldsymbol{\varphi} | \boldsymbol{\psi} \rangle = 0. \quad (3.14)$$

The inner products are just velocity integrals over the fixed spatial domain, so the order of integration can be exchanged with the time average. Then, since the test function $\boldsymbol{\varphi}(\boldsymbol{x})$ is arbitrary, the factors in the integrands can be equated:

$$\int_{\Omega} \boldsymbol{\psi}(\boldsymbol{x}') \left(\frac{1}{T} \int_0^T \boldsymbol{u}'(\boldsymbol{x}, t) \cdot \boldsymbol{u}'(\boldsymbol{x}', t) dt \right) d\boldsymbol{x}' = \lambda \boldsymbol{\psi}(\boldsymbol{x}). \quad (3.15)$$

This optimality condition is an eigenvalue problem for the eigenfunctions $\{\boldsymbol{\psi}_j\}$ of the two-point spatial correlation tensor of the velocity fluctuations $\boldsymbol{C}(\boldsymbol{x}, \boldsymbol{x}')$:

$$\int_{\Omega} \boldsymbol{C}(\boldsymbol{x}, \boldsymbol{x}') \boldsymbol{\psi}_j(\boldsymbol{x}') d\Omega' = \sigma_j^2 \boldsymbol{\psi}_j(\boldsymbol{x}), \quad (3.16)$$

where $\boldsymbol{C}(\boldsymbol{x}, \boldsymbol{x}') = \mathbb{E}[\boldsymbol{u}(\boldsymbol{x}, t) \boldsymbol{u}^*(\boldsymbol{x}', t)]$ and $\{\sigma_j\}$ are the POD eigenvalues, representing the average fluctuation kinetic energy captured by each mode. The coefficients $\boldsymbol{a}(t)$ can be extracted with the projection $a_j(t) = \langle \boldsymbol{\psi}_j(\boldsymbol{x}) | \boldsymbol{q}'(\boldsymbol{x}, t) \rangle$.

In practice the correlation tensor is often not feasible to construct, since it scales with the square of the discretized state dimension. Instead it is approximated numerically with either a singular value decomposition (SVD) or the snapshot method (Sirovich, 1987; Holmes *et al.*, 1996; Taira *et al.*, 2017). In this work we use the snapshot method as implemented in the `modred` library (Belson *et al.*, 2014), since it does not require storing the entire time series of high-dimensional discretized velocity fields in memory.

The method of snapshots is based on simple linear algebraic manipulations of the discretized form of the eigenvalue problem (3.16); see Holmes *et al.* (1996) for a derivation. Rather than form the spatial correlation tensor $\boldsymbol{C}(\boldsymbol{x}, \boldsymbol{x}')$, we compute a temporal correlation matrix \boldsymbol{R} with entries defined by

$$R_{jk} = \frac{1}{M} \langle \boldsymbol{q}'(t_j) | \boldsymbol{q}'(t_k) \rangle, \quad j, k = 1, 2, \dots, M. \quad (3.17)$$

The temporal correlation matrix \boldsymbol{R} has dimensions $M \times M$, and is typically much smaller than the discretized spatial correlation tensor. The eigenvalues of \boldsymbol{R} approximate those of

\mathcal{C} , and the modes that solve the discretized form of (3.16) are given by

$$\psi_k = \frac{1}{\sigma_k \sqrt{M}} \sum_{j=1}^M \mathbf{q}'(t_j) V_{jk}, \quad (3.18)$$

where $\mathbf{R}\mathbf{V} = \mathbf{V}\Sigma^2$ is the eigendecomposition of \mathbf{R} .

The proper orthogonal decomposition has the following useful properties:

1. The spatial modes form an orthonormal set: $\langle \psi_j | \psi_k \rangle = \delta_{jk}$.
2. The temporal coefficients are linearly uncorrelated: $\mathbb{E}[a_j a_k] = \sigma_k^2 \delta_{jk}$.
3. The modes can be ranked hierarchically by average energy content σ_k^2 .

Since POD can be viewed as a continuous form of the SVD, these properties are analogous to unitarity of the matrices of left and right singular vectors. The singular values quantify the statistical variance captured by the low-rank SVD approximation. As a consequence of the hierarchical ordering, the POD can be computed without *a priori* specification of the rank r , with the truncation determined later by a threshold based on the residual energy. Finally, since the modes are a linear combination of DNS snapshots, the reconstruction (3.1) automatically satisfies the incompressibility constraint and boundary conditions.

Although POD analysis is a powerful tool, it suffers from several limitations. First, it assumes a space-time separation of variables, which is typically appropriate for closed flows, but is often a poor approximation of inhomogeneous open flows dominated by traveling waves. Second, the condition that the temporal coefficients be linearly uncorrelated does not rule out nonlinear correlations. Finally, the optimality guarantee of the POD approximation is purely based on the instantaneous flow field and does not account for any time dynamics. The relationship between the first two issues will be examined at length in Chapter 6; nonlinear correlations will be shown to stabilize reduced-order models of advection-dominated flows while reducing the dimensionality of the model. The third issue motivates alternative modal decompositions like dynamic mode decomposition (Rowley *et al.*, 2009; Schmid, 2010; Mezić, 2013) and spectral POD (Towne *et al.*, 2018).

3.3 Dynamic mode decomposition (DMD)

Although proper orthogonal decomposition is guaranteed to provide an energy-optimal spatial reconstruction of the flow field, it sacrifices all temporal information in the computation of the correlation tensor. The POD basis is therefore purely kinematic and contains no dynamic information. An alternative approach is to compute the discrete Fourier

transform of the fields, which suffers from the opposite issue: frequency information is perfectly resolved, but the result is not necessarily associated with a useful reduced-order linear subspace for kinematic representation. Dynamic mode decomposition (DMD), introduced by Schmid (2010), is a useful compromise between these extremes.

DMD seeks to approximate a discrete-time linear evolution operator defined by

$$\min_{\mathcal{A}} \mathbb{E} \|\mathbf{u}(\mathbf{x}, t_{n+1}) - \mathcal{A}\mathbf{u}(\mathbf{x}, t_n)\|^2. \quad (3.19)$$

The description of nonlinear dynamics in terms a linear evolution operator acting on observables has a deep connection to Koopman theory (Rowley *et al.*, 2009; Mezić, 2013; Brunton *et al.*, 2022), but in terms of modal analysis it is often more useful to think of DMD as solving an alternative optimization to equation (3.8).

Given a series of snapshots, a least-squares solution to equation (3.19) could be found in terms of the pseudoinverse of the snapshot matrix. However, this calculation is typically computationally prohibitive, ill-conditioned, and disregards low-dimensional structure in the flow. Instead, DMD seeks to approximate the spectral properties of the operator \mathcal{A} without explicitly forming it. There are a variety of algorithms to compute DMD in conditions with limited or noisy data, but beginning with high-fidelity DNS snapshots we follow the simple *exact DMD* algorithm introduced by Tu *et al.* (2014).

Beginning with a truncated POD basis and associated coefficients $\{\mathbf{a}(t_n)\}_{n=1}^M$, the coefficients are arranged into time-shifted matrices $\mathbf{X} = [\mathbf{a}(t_1) \ \mathbf{a}(t_2) \ \cdots \ \mathbf{a}(t_{M-1})]$ and $\mathbf{X}' = [\mathbf{a}(t_2) \ \mathbf{a}(t_3) \ \cdots \ \mathbf{a}(t_M)]$. Here we assume the $\{t_n\}$ are evenly sampled in time, but it is possible to account for situations where this is not the case. A least-squares solution to (3.19) in the POD subspace is $\tilde{\mathbf{A}} = \mathbf{X}'\mathbf{X}^+$, where \mathbf{X}^+ is the pseudoinverse. With some assumptions, the spectrum of \mathcal{A} is approximated by the spectrum of the low-dimensional approximation $\tilde{\mathbf{A}}$, which can now be easily computed via an eigendecomposition:

$$\tilde{\mathbf{A}} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}, \quad (3.20)$$

where $\mathbf{V}, \mathbf{\Lambda} \in \mathbb{C}^{r \times r}$ and \mathbf{V} is the eigendecomposition of $\tilde{\mathbf{A}}$ (different from Eq. (3.18)). For details on theory and algorithms of dynamic mode decomposition, see Tu *et al.* (2014); Kutz *et al.* (2016). Based on this eigendecomposition, complex-valued DMD modes $\{\phi_k(\mathbf{x})\}$ and associated projection coefficients $\alpha(t)$ are linear combinations of the POD modes and coefficients, given by

$$\phi_k(\mathbf{x}) = \sum_{j=1}^r \psi_j(\mathbf{x})V_{jk} \quad \alpha(t) = \mathbf{V}^{-1}\mathbf{a}(t). \quad (3.21)$$

In principle the approximate time evolution is specified by the DMD eigenvalues $\{\lambda_k\} = \text{diag}(\mathbf{\Lambda})$, but in terms of reduced-order modeling the decomposition can also be viewed

as an alternative expansion to (3.1):

$$\mathbf{u}(\mathbf{x}, t) \simeq \bar{\mathbf{u}}(\mathbf{x}) + \sum_{k=1}^r \phi_k(\mathbf{x}) \alpha_k(t). \quad (3.22)$$

The DMD coefficient vector $\alpha(t)$ may then be modeled as a time series, as with the POD coefficients $\mathbf{a}(t)$.

This representation is essentially a similarity transformation of the POD basis; the two encode the same information and span the same subspace. However, the time-dependence in the optimization problem leads DMD to transform the POD basis to modes that tend to have similar frequency content (see Fig. 6.7, for instance). This can be leveraged in modeling strategies that exploit triadic interactions in the frequency domain, for instance; whereas POD happens to identify modes that are roughly coherent in time by coincidence only, the DMD modes are closer to pure harmonics.

3.4 Summary

Modal analysis of fluid flows is a deep field in its own right that has led to important advances in our understanding of flow physics; see for instance Taira *et al.* (2017); Towne *et al.* (2018); Taira *et al.* (2020); Barthel *et al.* (2021, 2022); Brunton *et al.* (2022) for recent advances. This chapter only introduces selected topics that will be important for later discussion of low-dimensional modeling. From that limited perspective, modal analysis seeks to address the spatial component of the space-time separation of variables in standard model reduction. This is made explicit in the generic modal decomposition (3.1).

With that in mind, arguably the two most important techniques for deriving a set of spatial modes are linear stability analysis and proper orthogonal decomposition (POD). Whereas linear stability analysis is an operator-based method that describes the evolution of infinitesimal perturbations, POD is a fully empirical approach that approximates the most energetic coherent structures in a statistical ensemble of fully nonlinear transient flow data. On the other hand, dynamic mode decomposition is a more recent alternative that might be thought of as an intermediate between stability analysis and POD. As with POD, it is an empirical modal decomposition, but like linear stability analysis, it produces modes that are coherent in both space and time, along with approximate eigenvalues describing their evolution.

While these will be the primary techniques used to develop modal bases in this work, it is worth bearing in mind at least one key limitation of the modal analysis framework. As suggested by the discussion of local stability analysis and transient energy growth, this problem is the artificial separation of space and time variables assumed by all of these methods. While this can be justified in the analysis of many PDEs, including diffusion-dominated fluid flow, it is ill-suited for the description of advection, or traveling waves. Along with the complications of nonlinearity introduced in Chapter 2, this will be a crit-

ical challenge faced by low-dimensional models constructed in this setting, with some strategies presented in Chapters 6 and 9.

Nevertheless, POD, DMD, and stability analysis are powerful and robust techniques for identifying coherent structures in unsteady fluid flows and will form the foundation of the low-dimensional modeling methods in Chapter 4.

Chapter 4

REDUCED-ORDER MODELING

Broadly speaking, reduced-order modeling is an interdisciplinary field of applied mathematics concerned with developing compressed approximations of high-dimensional dynamical systems, including PDEs (Antoulas, 2005; Noack *et al.*, 2011; Benner *et al.*, 2015; Rowley & Dawson, 2017). The modal analyses discussed in Chapter 3 may be viewed as linear dimensionality reduction methods that transform the system to a compact coordinate system in which low-dimensional dynamical systems models can be developed. In addition to acting as an inexpensive surrogate for the flow, such models can provide valuable insight into latent structure of the physical solutions. This chapter introduces the two key methods for developing nonlinear reduced-order models that will be applied throughout the remainder of this work: projection-based modeling (Sec. 4.1) and the SINDy framework for system identification (Sec. 4.2). The discussion of model reduction methods will continue in Chapter 8 with a proposed method for stochastic system identification.

4.1 Galerkin projection

We have already motivated model reduction from the perspective of dissipative nonlinear dynamics in Chapter 2, but there are even more basic reasons for such an approach that can be succinctly framed in terms of the spectral properties of linear operators. In a sense, a standard eigendecomposition is the simplest form of model reduction. Suppose for instance a discrete state vector $\mathbf{x}(t) \in \mathbb{R}^N$ evolves according to the linear system

$$\dot{\mathbf{x}} = \mathbf{L}\mathbf{x} \tag{4.1}$$

and the dynamics matrix \mathbf{L} has eigenvalue decomposition $\mathbf{L}\mathbf{V} = \mathbf{V}\mathbf{\Lambda}$ for the set of eigenvectors \mathbf{V} and diagonal eigenvalue matrix $\mathbf{\Lambda}$ in order of descending real part. The left eigenvectors \mathbf{W} solve the adjoint eigenvalue problem and may be normalized so that $\mathbf{W}^T\mathbf{V} = \mathbf{I}$.

If there is a *spectral gap* in the eigenvalues, so that the first r eigenvalues have real part $\mathcal{O}(\epsilon)$ for some $\epsilon \ll 1$ and the remaining $N - r$ eigenvalues are damped with negative real part $\mathcal{O}(1)$ or greater, then after a time t/ϵ , any component of \mathbf{x} in a direction of the higher-order modes will be exponentially damped. In this case the solution $\mathbf{x}(t)$ will be approximately restricted to the linear subspace spanned by the leading r modes. That is, $\mathbf{x}(t) \approx \mathbf{V}_r\mathbf{z}(t)$, where \mathbf{V}_r consists of the leading r eigenvectors and $\mathbf{z}(t)$ is the vector of

their coefficients. Then a reduced-order model of the dynamics $\dot{\mathbf{x}} \approx \mathbf{V}_r \dot{\mathbf{z}}$ is the projected model

$$\dot{\mathbf{z}} = \mathbf{W}_r^T \mathbf{L} \mathbf{V}_r \mathbf{z} = \mathbf{\Lambda}_r \mathbf{z}, \quad (4.2)$$

where the \mathbf{W}_r is the matrix of truncated adjoint eigenvectors.

The basis vectors in \mathbf{V}_r and \mathbf{W}_r make up the “trial” and “test” bases, respectively. If they happen to coincide, the projection is called a *Galerkin projection*; otherwise it is often referred to as a Petrov-Galerkin projection. The eigenvector transformation is particularly simple since it diagonalizes the dynamics, but different choices of basis may be appropriate in different circumstances. For example, linear model reduction is especially important for actuated systems, where low-latency models are at a premium; see for instance Antoulas (2005); Benner *et al.* (2015). In this case it is often more accurate to use methods such as balanced truncation (Moore, 1981) or the eigensystem realization algorithm (Juang & Pappa, 1985), which construct projection subspaces that account for the controllability and observability of a state-space system.

Alternatively, the projection basis can be constructed using an empirical method such as POD or DMD (see for instance Taira *et al.* (2017) or Chapter 3). POD is a convenient choice since the basis is orthonormal by construction. If \mathbf{V} is an r -dimensional basis of POD modes and $\mathbf{x} \approx \mathbf{V} \mathbf{z}$, then the Galerkin projection of the dynamics (4.1) onto this subspace is

$$\dot{\mathbf{z}} = \mathbf{V}^T \mathbf{L} \mathbf{V} \mathbf{z}, \quad (4.3)$$

where the reduced-order dynamics matrix will not generally be diagonal. Assuming that the flow is statistically stationary and the ensemble is sufficiently resolved, the POD approximation provides optimal kinematic resolution in an orthonormal basis.

Controllability, observability, and the proper orthogonal decomposition are connected via the Gramian matrices, which allows for analysis in terms of linear systems theory (Bagheri *et al.*, 2009b) and hybrid methods such as balanced POD (Rowley, 2005). In the context of fluid dynamics, linear models have been particularly successful for shear flows (Bagheri *et al.*, 2009a; McKeon & Sharma, 2010; Luhar *et al.*, 2014; Zare *et al.*, 2017; Rowley & Dawson, 2017), since these are often dominated by the linear mechanism of non-normal transient energy amplification.

However, in many cases nonlinearity is a key feature of the dynamics. The Galerkin projection framework can be readily extended to nonlinear dynamics as follows. Suppose that for a general autonomous dynamical system

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^N, \quad (4.4)$$

we wish to approximate the state in an r -dimensional linear subspace $\mathbf{x} \approx \mathbf{V} \mathbf{z}$ with $\mathbf{V} \in \mathbb{R}^{N \times r}$. For simplicity (and because it will be the case for the POD basis), we will assume that the modes are orthonormal, so that $\mathbf{V}^T \mathbf{V} = \mathbf{I}$, although the extension to Petrov-Galerkin projection for different test and trial bases is straightforward (Benner *et al.*, 2015).

Writing the ODE (4.4) in the implicit form $\mathbf{N}(\mathbf{x}) = \dot{\mathbf{x}} - \mathbf{f}(\mathbf{x}) = 0$, there will inevitably be

some residual error $\mathbf{r}(\mathbf{z}) = \mathbf{N}(\mathbf{V}\mathbf{z})$ in the dynamics approximation. To realize the optimal approximation of the dynamics in this subspace, this residual should be orthogonal to the subspace, so that $\mathbf{V}^T \mathbf{r}(\mathbf{z}) = 0$. Expanding the implicit form, this condition implies that the reduced-order dynamics are

$$\dot{\mathbf{z}} = \mathbf{V}^T \mathbf{f}(\mathbf{V}\mathbf{z}). \quad (4.5)$$

If $\mathbf{f}(\mathbf{x})$ is linear, this simplifies immediately to the linear reduced-order model (4.3). However, for arbitrary nonlinearity Eq. (4.5) cannot be reduced further. This means that the N -dimensional right-hand side $\mathbf{f}(\mathbf{V}\mathbf{z})$ must be evaluated before projecting back to the low-dimensional subspace. There are a number of strategies for avoiding this costly step (Carlberg *et al.*, 2011; Amsallem & Farhat, 2012; Carlberg *et al.*, 2013; Carlberg, 2015; Carlberg *et al.*, 2017), but as will be seen below, when the nonlinearity is polynomial the projection can be computed offline.

Galerkin models of fluid flow Linear model reduction is an appealing and practical framework for a number of reasons (not least the breadth of analytic tools available for linear systems), but it disregards the critical role played by nonlinear interactions in multiscale fluid flows. As suggested by the discussion in Chapter 2, this is particularly important for flows with absolute instability. Fortunately, the nonlinear extension of Galerkin projection can be calculated in closed form for the Navier-Stokes equations (2.1), resulting in a linear-quadratic system of equations that parallels the spectral dynamics equations (2.3) for homogeneous flows.

As with the general case above, let $\mathcal{N}[\mathbf{q}] = 0$ be the Navier-Stokes equations written in implicit form with $\mathbf{q}(\mathbf{x}, t) = [\mathbf{u}^T \ p]^T$. Note that \mathbf{x} here is the spatial variable, which should not be confused with the state vector \mathbf{x} in the dynamical systems notation above. By approximating the flow field with a POD basis, we expect some residual error $\mathbf{r}(\mathbf{x}, t)$ in the approximated dynamics defined by

$$\mathbf{r}(\mathbf{x}, t) = \mathcal{N} \left[\mathbf{q}_0(\mathbf{x}) + \sum_{\ell=1}^r \boldsymbol{\psi}_\ell(\mathbf{x}) a_\ell(t) \right], \quad (4.6)$$

where $\mathbf{q}_0(\mathbf{x})$ might be the mean flow, a stationary solution, or any other time-invariant divergence-free base flow.

In order to minimize the residual in this basis, the Galerkin projection condition is that that the residual be orthogonal to each mode, so that $\langle \boldsymbol{\psi}_k | \mathbf{r} \rangle = 0$ for any mode $\boldsymbol{\psi}_k$. The mass equation (2.1b) may be ignored if all spatial modes are assumed to be divergence-free, while the projected the momentum equation gives the low-dimensional dynamics $\dot{\mathbf{a}}$. The derivation is straightforward and is given in standard references (Holmes *et al.*, 1996;

Noack *et al.*, 2011), leading to the linear-quadratic system of ODEs

$$\dot{a}_k = F_k + \sum_{\ell=1}^r L_{k\ell} a_\ell + \sum_{\ell=1}^r \sum_{m=1}^r Q_{k\ell m} a_\ell a_m, \quad (4.7)$$

with constant, linear, and quadratic terms given by

$$F_k = \left\langle \boldsymbol{\psi}_k \left| -\nabla \cdot (\mathbf{u}_0 \otimes \mathbf{u}_0) + \frac{1}{Re} \nabla^2 \mathbf{u}_0 \right. \right\rangle \quad (4.8a)$$

$$L_{k\ell} = \left\langle \boldsymbol{\psi}_k \left| -\nabla \cdot (\mathbf{u}_0 \otimes \mathbf{u}_\ell + \mathbf{u}_\ell \otimes \mathbf{u}_0) - \nabla p_\ell + \frac{1}{Re} \nabla^2 \mathbf{u}_\ell \right. \right\rangle \quad (4.8b)$$

$$Q_{k\ell m} = \left\langle \boldsymbol{\psi}_k \left| -\nabla \cdot (\mathbf{u}_m \otimes \mathbf{u}_\ell + \mathbf{u}_\ell \otimes \mathbf{u}_m) \right. \right\rangle. \quad (4.8c)$$

Here the modes consist of both velocity and pressure components, so that $\boldsymbol{\psi}_k = [\mathbf{u}_k^T \ p_k]^T$. The pressure gradient term can often be neglected, depending on the flow configuration. In this work the only case where it is not negligible is the mixing layer in Chapter 7. Here we assume that each POD mode was calculated with an associated pressure field as described in Sec. 3.2, but there are several strategies for modeling the pressure term if it is retained; for a detailed discussion of the pressure term see Noack *et al.* (2005).

Also note that the constant term vanishes if the flow is expanded about a steady-state solution of the governing equations. However, if \mathbf{q}_0 is the mean flow it is typically not a stationary solution, in which case this term represents important mean-flow forcing and is not negligible.

As discussed above, for general nonlinear dynamics a closed-form low-dimensional system cannot be computed. In this case, because the nonlinearity in the incompressible Navier-Stokes equations is purely quadratic, the projection of this term can be computed offline, leading to the quadratic tensor \mathbf{Q} . Under certain assumptions, the same can be done for isentropic compressible flow (Rowley *et al.*, 2004) and magnetohydrodynamics equations (Kaptanoglu *et al.*, 2021b).

In principle, we might expect that the POD-Galerkin system (4.7) leads to approximate solutions with comparable accuracy to the resolution of the expansion basis. However, for reasons that will be discussed below, the long-time behavior of the reduced-order model may deviate significantly from that of the underlying physical system. The accuracy may not even necessarily improve with increasing subspace dimension r . A central issue is that solutions of the model are not constrained to lie on an invariant manifold of the flow. This is a particular problem for advection-dominated flows, where coefficients associated with harmonics evolve independently from the fundamental modes, eventually leading to an unphysical loss of coherence. See Chapter 6 for a detailed investigation and proposed solution to this issue.

A model similar to (4.7) may also be derived beginning with the DMD expansion (3.22). In this case, since the DMD basis is not orthonormal, the Galerkin projection must be re-

placed with the more general oblique projection of Petrov-Galerkin methods (Benner *et al.*, 2015). We implement this with a coordinate transform of the standard POD-Galerkin model based on the DMD eigenvector matrix. This can be useful for instance if the stabilization strategy benefits from modes with pure frequency content, as will be the case in Chapter 6.

Symmetries and conservation properties The combination of proper orthogonal decomposition and Galerkin projection is a popular choice for model reduction of incompressible flows, in part because it provides a semi-empirical “gray-box” framework that combines some of the advantages of data-driven analysis with physics-based modeling. One consequence of this is that the dynamical system Eq. (4.7) inherits certain properties from the Navier-Stokes equations.

For example, although we will usually take the base flow $\psi_0(\boldsymbol{x})$ to be the time-averaged mean flow, if the flow is instead expanded about the steady-state solution the constant forcing term \boldsymbol{F} vanishes, ensuring that the origin is a fixed point. If, in addition, an energy inner product such as Eq. (3.2) is used for projection, the reduced-order model also preserves Lyapunov stability of the origin, although not necessarily the stability of limit cycles or other fixed points (Rowley *et al.*, 2004).

A key feature of the quadratic nonlinearity of the Navier-Stokes equations is that it is energy-preserving, in the sense that it has no net contribution to the evolution equation for kinetic energy in the spectral domain (with some restrictions on boundary conditions) (Kraichnan & Chen, 1989). This is the foundation of the energy cascade picture of turbulence, in which the role of the nonlinearity is to transfer energy from the large scales to the small, dissipative scales. For the Galerkin system, the energy preservation condition implies that (Schlegel & Noack, 2015)

$$Q_{ijk} + Q_{ikj} + Q_{jik} + Q_{jki} + Q_{kij} + Q_{kji} = 0. \quad (4.9)$$

This is often approximately true numerically for POD-Galerkin systems, but enforcing it explicitly can improve the stability of the model (Cordier *et al.*, 2013). Based on the structure of the linear-quadratic system (4.7), the quadratic tensor can also be symmetrized in the last two indices:

$$Q_{ijk} = Q_{ikj}. \quad (4.10)$$

Finally, Schlegel & Noack (2015) showed that the long-term boundedness of the Galerkin system can be guaranteed if a criterion based on generalized Lyapunov functions is satisfied. Although this is a dynamical systems analysis and not an *a priori* theoretical guarantee, it can also be used as a constraint in sparse model identification to construct empirical Galerkin-type models with guaranteed stability (Kaptanoglu *et al.*, 2021a).

Shortcomings of the Galerkin method Although the POD-Galerkin framework is a popular and theoretically convenient starting place for low-dimensional modeling, it suf-

fers from several key limitations. First, linear representations of advection-dominated flows fundamentally suffer from a large Kolmogorov n -width, or a linear subspace that slowly approaches full kinematic resolution with increasing dimension (Grimberg *et al.*, 2020). In this case, even when enough modes are retained to reconstruct the flow field, the Galerkin model may not faithfully represent the underlying physics.

Reduced-order models based on heavily truncated linear representations are therefore known to suffer from severe instabilities without careful closure modeling (Aubry *et al.*, 1988; Noack *et al.*, 2008; Wang *et al.*, 2012; Maulik *et al.*, 2019). From a numerical perspective, part of the problem is the Galerkin formulation commonly used to derive a set of time-continuous ordinary differential equations (Carlberg *et al.*, 2017; Grimberg *et al.*, 2020), but there are also at least two physical reasons for the instability of projection-based models:

1. The higher-order modes tend to represent smaller scales of the flow, which are responsible for the bulk of energy dissipation, so that truncated models may not accurately capture the energy cascade in the flow. This motivates eddy viscosity-type modifications by analogy with classical turbulence closure models (Aubry *et al.*, 1988; Wang *et al.*, 2012) as well as alternative Galerkin schemes that explicitly target energy balance (Balajewicz *et al.*, 2013; Mohebujjaman *et al.*, 2017).
2. The dimensionality of the linear subspace required to reconstruct the flow field may significantly exceed the intrinsic dimension of the attractor of the system. Since traditional model reduction methods have one state variable per mode, the projected dynamics may have many more degrees of freedom than the physical system. For instance, the traveling wave solution to a linear advection equation on a periodic domain may require arbitrarily many Fourier modes for a linear reconstruction, yet with the method of characteristics the state is determined by a single degree of freedom: the scalar phase. This intuition can be used to develop empirical manifold models that eliminate spurious variables; see Chapter 6 and Callahan *et al.* (2022).

The competition between these two effects tends to lead to fragile Galerkin systems without further modeling. Enough modes must be retained to sufficiently resolve dissipation, but this large number of kinematic modes may be considerably larger than the number of dynamic degrees of freedom. Therefore, the dynamics of models that include a large number of modes may not resemble those of the underlying flow. A case study of these considerations is the pioneering work of Noack *et al.* (2003) modeling the two-dimensional flow past a cylinder. With an augmented POD basis and a dynamical systems analysis, they reduce a structurally unstable eight-dimensional Galerkin system to a two-dimensional cubic model that reproduces the dominant flow physics; see also Chapter 5 for an updated presentation of this model.

The issues of stability and validity are intimately connected to the question of correlation. The temporal coefficients of POD modes are linearly uncorrelated on aver-

age (Holmes *et al.*, 1996), but no such guarantee is available for nonlinear correlation. For example, one mode may be a harmonic of another; in this case their temporal coefficients are linearly uncorrelated but the harmonic is a perfect algebraic function of the fundamental. If these coefficients are modeled independently, as in a classical Galerkin system, slight inaccuracies can lead to decoherence and unphysical solutions. Later chapters will explore nonlinear dimensionality reduction (Chapter 6) and stochastic averaging (Chapter 7) as systematic solutions for stabilizing these models while respecting their natural dynamics.

4.2 Sparse identification of nonlinear dynamics

As an alternative to projection-based reduced-order modeling, a low-dimensional system can be approximated directly from the data in a procedure typically called system identification. In a continuous-time setting, this is done by estimating parameters Ξ for a function $\mathbf{f}(\mathbf{a}; \Xi) \approx \dot{\mathbf{a}}$ that solve the approximation problem

$$\underset{\Xi}{\text{minimize}} \mathbb{E} \|\dot{\mathbf{a}} - \mathbf{f}(\mathbf{a}; \Xi)\|^2. \quad (4.11)$$

This is similar to the residual minimization in Galerkin projection, except that knowledge of the governing equations is neither required nor assumed. Instead, some intuition about the structure of the dynamics is typically encoded in the parameterization of \mathbf{f} . Different parameterizations lead to symbolic regression (Schmidt & Lipson, 2009), operator inference (Peherstorfer & Willcox, 2016), or deep learning (Vlachas *et al.*, 2018). In the discrete-time counterpart to (4.11), the NARMAX framework provides a powerful approach that can incorporate time delays, stochastic forcing, and exogenous inputs (Billings, 2013).

In this work we apply the *sparse identification of nonlinear dynamics* (SINDy) approach to system identification (Brunton *et al.*, 2016a), illustrated in Fig. 4.1. Let $\Theta(\mathbf{a})$ denote a library of candidate functions of the time series $\mathbf{a}(t)$, e.g.

$$\Theta \left([a_1 \ a_2]^T \right) = [a_1 \ a_2 \ a_1^2 \ a_1 a_2 \ a_2^2 \ \dots]^T. \quad (4.12)$$

We seek a sparse approximation $\dot{\mathbf{a}} \approx \mathbf{f}(\mathbf{a}; \Xi) = \Xi \Theta(\mathbf{a})$ in the range of these candidate functions. We can frame this as a linear algebra problem by forming the $r \times M$ data matrix \mathbf{X} as in the DMD method of Sec. 3.3, where each column is a snapshot of modal coefficients in time. Similarly, we estimate the time derivative $\dot{\mathbf{X}}$, for instance with second-order finite differencing. The SINDy formulation of the optimization problem (4.11) is then

$$\min_{\Xi} \left\| \dot{\mathbf{X}} - \Xi \Theta(\mathbf{X}) \right\|_2^2 + \gamma \|\Xi\|_0, \quad (4.13)$$

where $\|\cdot\|_p$ indicates the p -norm and γ is some regularization weight. This optimization problem is non-convex and requires a combinatorial search over function combinations.

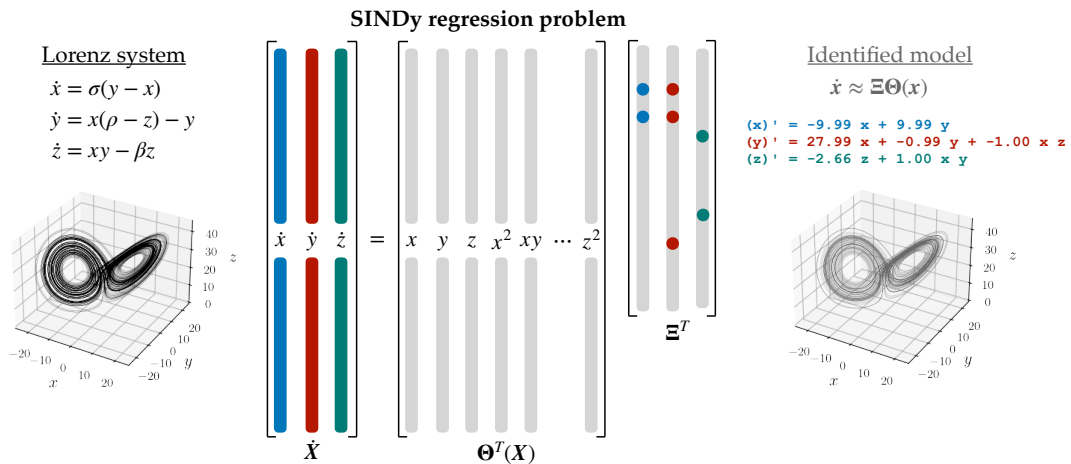


Figure 4.1: Model discovery applied to the Lorenz system. The SINDy algorithm assumes that unknown governing equations for a dynamical system can be approximated by a sparse linear combination of candidate functions. The coefficient vectors can then be identified by solving a sparse regression problem. Note that the matrices Θ and Ξ are the transpose of the usual SINDy formulation. Figure adapted from Brunton *et al.* (2016a).

To avoid this, we approximate the solution to (4.13) with the greedy forward regression orthogonal least squares (FROLS) algorithm used in NARMAX analysis (Billings, 2013). Other possible optimization algorithms include sequential thresholding (Brunton *et al.*, 2016a), SR3 (Champion *et al.*, 2019b), and convex relaxation (Loiseau, 2020). See also Brunton *et al.* (2016a); Loiseau & Brunton (2018); Loiseau (2020); Callaham *et al.* (2022) for details on SINDy reduced-order modeling of fluid flows.

Although the systems obtained from POD-Galerkin projection will be dense in general, we expect that the dynamics can be closely approximated by a sparse combination of candidate functions. This may be justified intuitively by the sparse structure of the triadic interactions in isotropic flow (see Section 2.1), where only r^2 out of r^3 possible interactions are admissible (Rubini *et al.*, 2020). Moreover, as observed in Section 3.3, DMD approximates a diagonalization of the evolution operator, so that by analogy with normal form theory we may reasonably hope for a minimal representation of the dynamics if we work with DMD coefficients α rather than POD coefficients a (Loiseau, 2020; Callaham *et al.*, 2022). More generally, sparsity promotion reflects the inductive bias of Occam’s razor or Pareto analysis, where we expect that the most important features of the dynamics will be due to a small subset of terms. For incompressible flows, it has been repeatedly demonstrated that a library of low-order polynomials provides a good basis of functions. Quadratic terms are clearly necessary to capture the advective nonlinearity of the Navier-Stokes equations, but cubic terms allow the model to resolve Stuart-Landau-type nonlinear stability mechanisms (Loiseau & Brunton, 2018).

For PDEs with more general nonlinearity, low-order polynomials still present an attractive basis for SINDy. In many interesting regimes the effect of the nonlinearity may be relatively weak, so that quadratic and cubic polynomials can be seen as second- or third-order Taylor series approximations to the underlying nonlinearity. Moreover, even strongly nonlinear systems can be *lifted* with a change of variables to a coordinate system wherein the dynamics are linear-quadratic (Rowley *et al.*, 2004; Qian *et al.*, 2020; Kaptanoglu *et al.*, 2021b).

Enforcing physical constraints Although the SINDy method presented above is unconstrained, Loiseau & Brunton (2018); Loiseau (2020) and Kaptanoglu *et al.* (2021a) showed that an effective alternative to sparse regression is regularized, constrained regression. This approach replaces the optimization problem (4.13) with

$$\begin{aligned} \min_{\Xi} \left\| \dot{\mathbf{X}} - \Xi \Theta(\mathbf{X}) \right\|_2^2 + \gamma \|\Xi\|_2 \\ \text{subject to } \mathbf{C}\Xi = 0, \quad \mathbf{H}\Xi \leq 0. \end{aligned} \quad (4.14)$$

For example, if the library only contains linear and quadratic terms in the state variables, the symmetries and conservation properties (4.9) and (4.10) of the Galerkin system can be used to derive equality constraints on the relationships between the coefficients of the nonlinear terms (Loiseau & Brunton, 2018). See Loiseau (2020) for examples of other physically-motivated equality and inequality constraints, including equivariance, dissipative dynamics, and preservation of the stable manifold.

Besides enforcing hard constraints on the coefficients, Kaptanoglu *et al.* (2021a) derived a stability-promoting optimization scheme based on the trapping theorem due to Schlegel & Noack (2015). This method is based on a generalized form of Lyapunov stability analysis; although the derivation and implementation is somewhat technical, results show a dramatic improvement in long-time stability of the SINDy models, even when the input data is highly corrupted by measurement noise.

The energy-preserving and trapping conditions are based on analysis of linear-quadratic dynamics and it is still currently not clear how they might be applied to cubic models. Since both constraints and cubic terms have been shown to significantly improve SINDy models of fluid flow (Loiseau & Brunton, 2018), combining the two could be a powerful advance. One possibility is to first identify a constrained linear-quadratic system and then perform a secondary reduction to a cubic model by averaging over fast variables in the quadratic system; see Chapter 7 for further discussion.

Other extensions Since its introduction, the SINDy framework has been extended in a number of directions. These include generalizations to actuated systems (Brunton *et al.*, 2016b; Kaiser *et al.*, 2018), stochastic differential equations (Boninsegna *et al.*, 2018; Callahan *et al.*, 2021b), partial differential equations (Rudy *et al.*, 2017; Schaeffer *et al.*, 2017),

weak-form equations (Reinbold *et al.*, 2020; Messenger & Bortz, 2021), and implicit equations (Mangan *et al.*, 2016; Kaheman *et al.*, 2020). The basic sparse regression algorithm has also been refined with more robust numerical optimization (Champion *et al.*, 2019b) and other techniques for noisy observational data (Delahunt & Kutz, 2022; Fasel *et al.*, 2022). Many of these advances, including the constrained and trapping optimizations discussed above, are available in the open-source PySINDy library (de Silva *et al.*, 2020; Kaptanoglu *et al.*, 2022)¹.

4.3 Summary

This chapter reviews the two reduced-order modeling methods that will be central to the remainder of this thesis: POD-Galerkin projection and the SINDy model discovery framework. Of course, there are many important alternative model reduction methods. These include physics-based methods such as balanced POD (Rowley, 2005), resolvent analysis (McKeon & Sharma, 2010; McKeon *et al.*, 2013; Sharma & McKeon, 2013), center manifold reduction (Carini *et al.*, 2015), and weakly nonlinear analysis (Sipp & Lebedev, 2007; Meliga *et al.*, 2009; Meliga & Chomaz, 2011; Rigas *et al.*, 2017). Empirical system identification is also crucial when analytic approaches are not practical; these methods include the eigensystem realization algorithm (Juang & Pappa, 1985; Ma *et al.*, 2011; Brunton *et al.*, 2014), the NARMAX framework (Billings, 2013), symbolic regression (Schmidt & Lipson, 2009; Cranmer *et al.*, 2020b), black-box forecasting (Gautier *et al.*, 2015; Vlachas *et al.*, 2018; Pathak *et al.*, 2018; Hesthaven & Ubbiali, 2018; Wan & Sapsis, 2018), and other time-series analysis methods (Crutchfield & McNamara, 1987; Kantz & Schreiber, 2003; Bradley & Kantz, 2015; Tabar, 2019).

As always, the appropriate algorithm will typically be application-dependent. For the purposes of deriving compact, nonlinear dynamical systems approximations to unsteady fluid flow, Galerkin projection and SINDy are particularly useful because they are interpretable, relatively simple to implement, and balance inputs from observational data and physical models. After Chapter 5 illustrates the application of both methods to the cylinder wake, Chapters 6 and 7 introduce extensions to Galerkin and SINDy modeling based on multiscale dynamical systems analysis. Chapter 8 returns to the topic of system identification to present the Langevin regression algorithm, an extension of SINDy to stochastic differential equations.

¹<https://github.com/dynamicslab/pysindy>

Chapter 5

CASE STUDY: THE CYLINDER WAKE

For all that it bears no obvious resemblance to any practical problems, the flow past a cylinder has played an outsized role in the history of fluid dynamics, including d’Alembert’s paradox, the Kutta-Joukowski theorem connecting circulation and lift, and airfoil theory (via conformal mapping and the Joukowski transform). More recently, the stable vortex shedding limit cycle at moderate Reynolds numbers has generated renewed interest in the cylinder wake, including stability analysis (Zebib, 1987; Barkley & Henderson, 1996; Barkley, 2006; Sipp & Lebedev, 2007), low-dimensional modeling (Deane *et al.*, 1991; Ma & Karniadakis, 2002; Noack *et al.*, 2003; Brunton *et al.*, 2016a; Loiseau *et al.*, 2018a), and closed-loop flow control (Tadmor *et al.*, 2011).

Alongside seminal works such as Lorenz (1963); Ruelle & Takens (1971); Aubry *et al.* (1988); Holmes *et al.* (1996) the reduced-order model for the cylinder wake developed by Noack *et al.* (2003) is one of the centerpiece examples of the dynamical systems perspective on fluid dynamics. Because it highlights much of the important physics of model reduction reviewed in Chapters 2, 3, and 4, this chapter briefly summarizes relevant results of that paper in Sections 5.2-5.4. Section 5.5 concludes with an example of a data-driven manifold closure using sparse regression which yields a highly accurate two-dimensional Galerkin-type model.

5.1 Flow configuration

We model the two-dimensional incompressible flow past a cylinder at Reynolds number $Re = UL/\nu = 100$. As in Noack *et al.* (2003), the simulation domain is $x \in (-5, 15)$, $y \in (-5, 5)$; this configuration is shown in Fig. 5.1. At low Reynolds numbers, the equilibrium solution is a stable, symmetric, laminar profile similar to that predicted by potential flow theory but modified by a shear-induced recirculation bubble.

At $Re_c \approx 46$, the flow undergoes a supercritical Hopf bifurcation where the wake becomes unstable to perturbations and tends toward a stable vortex shedding limit cycle. The post-transient flow at $Re = 100$ is therefore characterized by a von Kármán vortex street, as shown in Fig. 5.1. The transient evolution of the flow is captured by the time history of the leading POD coefficients in Fig. 5.2; the instability grows from the initial unstable steady state until the flow reaches the stable limit cycle. This transient growth/saturation behavior is determined by the Reynolds stress deformations (Noack *et al.*, 2003; Sipp & Lebedev, 2007). Consequently, the flow can be accurately described by a cubic Stuart-Landau equation, as shown below.

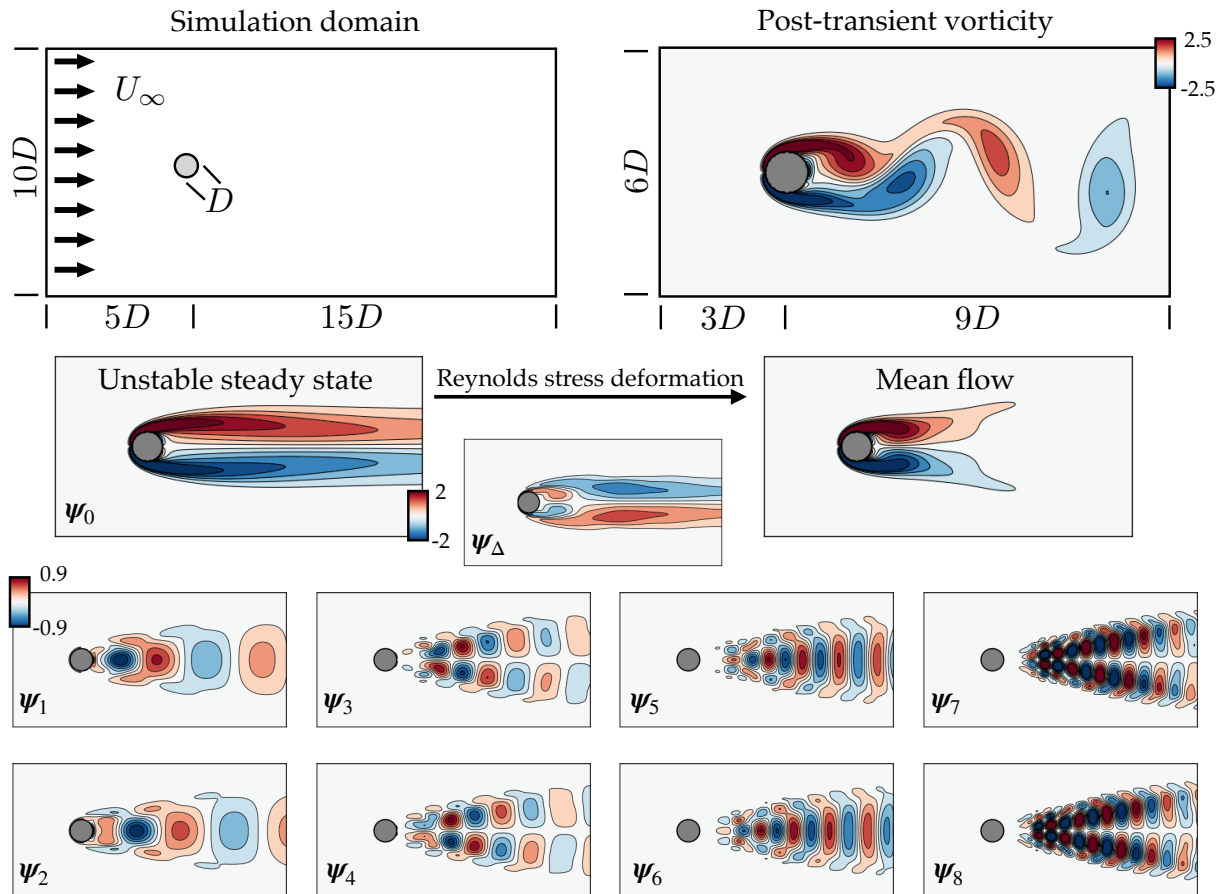


Figure 5.1: **Two-dimensional flow past a cylinder.** The simulation domain (top left) is the same as Noack *et al.* (2003). The post-transient flow is characterized by periodic vortex shedding (top right). This limit cycle is stabilized by Reynolds stress-mediated base flow deformation, which can be captured in POD-Galerkin models by augmenting the POD basis with a “shift mode” ψ_Δ capturing the deformation of the base flow (middle). The standard POD modes $\psi_{1,2,\dots,8}$ appear in pairs (bottom); phase-lagged oscillations allow the fixed spatial structures to capture advection of the vortex street.




Figure 5.2: **Transient evolution of the leading POD coefficients.** The pair a_1, a_2 describe the periodic vortex shedding, which begins as a linear instability with exponential energy growth. Eventually the base flow deformation (parameterized by the shift mode a_Δ) grows to balance the perturbations, stabilizing the limit cycle. The self-exciting, self-limiting oscillation is characteristic of dynamics above a supercritical Hopf bifurcation.


We solve the Navier-Stokes equations with direct numerical simulation (DNS) using the spectral element solver Nek5000 (Fischer *et al.*, 2008). The domain is discretized with 2622 sixth-order spectral elements, refined close to the boundaries and centerline. Diffusive terms are integrated with third-order backwards differentiation, while convective terms are advanced with a third-order extrapolation.

5.2 Modal analysis

The localized, periodic nature of the post-transient evolution suggests that modal analysis should be able to successfully describe the main features of this flow. We compute the mean-subtracted proper orthogonal decomposition (POD) using the method of snapshots (Sirovich, 1987) from 100 fields on the vortex shedding limit cycle. The fields are separated by an interval $\Delta t = 0.056$ so that the integration time corresponds to one full period. As shown in Fig. 5.3b, the leading 8 POD modes provide near-perfect kinematic resolution, with a relative fluctuation kinetic energy residual of order 10^{-4} .

The leading modes themselves are visualized by vorticity in Fig. 5.1. Each pair of modes (first and second, third and fourth, etc.) is analogous to the complex-conjugate eigenvectors describing propagating waves; they capture the same spatial scales, but are phase-shifted slightly. In fact, this illustrates one of the main challenges of modal analysis for advection-diffusion problems. Although the limit cycle flow is purely oscillatory, the modal decomposition introduces an artificial space/time separation and the modal basis must span several times the intrinsic dimensionality of the system in order to resolve the advecting structures.

Another useful property of POD analysis is that the orthogonality of right singular vectors guarantees that the time evolution of the coefficients $a_i(t)$ are linearly uncorrelated in time. However, this does not exclude *nonlinear* correlations. In fact, as explored by Loiseau *et al.* (2018a); Callaham *et al.* (2022) and Chapter 6, we might well expect that for an oscillating flow, the triadic interactions arising from projecting a modal decomposition of the form of Eq. (3.1) onto the dynamics would generate higher harmonics in the higher modes. That is, if the leading modes evolve at the natural frequency, $a_{1,2} \sim e^{\pm i\omega t}$, the triadic terms proportional to a_1^2 and a_2^2 will excite higher-order modes at $e^{\pm 2i\omega t}$. Mean-



(a) Phase portraits for the POD coefficients.


(b) Normalized POD eigenvalues. The residual energy of an 8-dimensional basis is $\mathcal{O}(10^{-4})$.

Figure 5.3: POD coefficients for the post-transient $Re = 100$ cylinder wake. The POD coefficients are guaranteed to be linearly uncorrelated, but the phase portraits (a) show clear evidence of *nonlinear* correlation. The Lissajous figures are characteristic of pure harmonics of the fundamental pair (a_1, a_2) . The modes appear in pairs with approximately the same energy (b), typical of oscillations or traveling waves in the flow.

while, nonlinear interactions such as $a_1 a_2$ have zero frequency and instead modulate the amplitude. This is borne out by a weakly nonlinear analysis of the flow near the bifurcation point; see Sipp & Lebedev (2007) and Chapter 2.

We can observe this effect for the cylinder wake by plotting phase portraits for the temporal coefficients, as in Fig. 5.3a. These Lissajous-type figures show clear evidence of higher harmonics in the higher-order modes. Therefore, although the coefficients are linearly uncorrelated, the post-transient evolution of the coefficients are nonlinearly correlated in the sense that the higher-order coefficients can be expressed as algebraic functions of the pair (a_1, a_2) (see Sec. 5.5). This evidence also supports the idea that the slow dynamics are inherently two-dimensional; if nonlinear correlations can explain all the higher-order terms then we only need a phase and amplitude of the leading pair of modes to fully characterize the state of the system.

5.3 POD-Galerkin model


The first successful POD-Galerkin model of the vortex shedding past the cylinder was developed by Deane *et al.* (1991), who modeled the evolution of the leading eight POD modes in the post-transient regime. Although this set of modes can reconstruct the flow with high accuracy and the model dynamics were shown to match the DNS over one period of the flow, this model is in fact structurally unstable. That is, the model may perfectly describe the limit cycle in theory, but small errors in the estimates of the Galerkin operators and lead to asymptotic instabilities in the model. In other words, there is no mechanism to stabilize the limit cycle; Fig. 5.4a (grey) shows that from an infinitesimal perturbation the energy grows exponentially.

The reason for this is that the model based on an expansion about the *mean* flow does not resolve the Stuart-Landau nonlinear stability mechanism. On the other hand, the model naturally stabilizes if the POD basis is augmented with a “shift mode” that accounts for the base flow deformation (Noack *et al.*, 2003). The shift mode $\mathbf{q}_\Delta(\mathbf{x})$ is computed from the the difference $\bar{\mathbf{q}} - \mathbf{q}_0$ between the mean flow $\bar{\mathbf{q}}$ and the linearly unstable steady-state base flow \mathbf{q}_0 , orthonormalized with respect to the rest of the POD basis using a Gram-Schmidt procedure. Figure 5.1 shows the computed shift mode.


The modal ansatz (3.1) is then replaced by

$$\mathbf{q}(\mathbf{x}, t) = \mathbf{q}_0(\mathbf{x}) + a_\Delta(t)\mathbf{q}_\Delta(\mathbf{x}) + \sum_{j=1} a_j(t)\psi_j(\mathbf{x}), \quad (5.1)$$

where the new temporal coefficient $a_\Delta(t)$ parameterizes the mean flow deformation en route to the limit cycle. For notational convenience, we will treat the Δ index of the shift mode as the final integer index, so that if we use a basis of eight POD modes, $a_\Delta \equiv a_9$. The linear and quadratic operators are the same as in Eqs. (4.7) and (4.8), except that the constant term vanishes for an expansion around the steady-state base flow \mathbf{q}_0 . The



(a) Energy comparison of the reduced-order models.



(b) Parabolic manifold for the transient cylinder wake.

Figure 5.4: Low-dimensional models of the transient cylinder wake. A model based on the post-transient limit cycle (grey) is structurally unstable without the addition of a shift mode capturing the base flow deformation (grey-blue). An analytic invariant manifold reduction (blue) is a more accurate approximation, although the lack of resolution for higher harmonics results in an overestimate of post-transient energy. This is corrected by the sparse Galerkin (SINDy) model (orange), which has the same structure as the invariant manifold model and can capture higher harmonics (also see Fig. 5.6). The nine degrees of freedom required for kinematic resolution approximately evolve on a two-dimensional parabolic manifold (b). Even when stable, POD-Galerkin models are not confined to the inertial manifold (left). The invariant manifold and SINDy models restrict the dynamics to the manifold (middle, right).

pressure gradient term is also negligible for this flow (Noack *et al.*, 2003, 2005).

Visualizing the transient evolution with the additional degree of freedom $a_\Delta(t)$ for the shift mode makes clear the parabolic structure of the underlying manifold (Fig. 5.4b). The deformation grows in proportion to the energy of the perturbations until the two are in balance in the limit cycle. Noack *et al.* (2003) show that this correction stabilizes the model, even when the POD basis only includes the leading pair of modes (Fig. 5.4a, red).

5.4 Invariant manifold model

Although the shift mode correction effectively stabilizes the model on a limit cycle, this approach has two main physical flaws. First, since the leading order terms tend to donate energy to the higher-order terms on average (see Fig. 9 in Noack *et al.* (2003)), truncation of the modal basis at two terms also truncates the “cascade”. Without higher-order terms to receive and dissipate the energy, the final limit cycle amplitude is overestimated.

This can be rectified by retaining the additional POD modes in the model (Fig. 5.4a, grey-blue), at the cost of increasing the model dimensionality to $r = 9$. A second problem is that the transfer of energy between modes happens more slowly in the model than in the flow. This manifests in a transient overshoot in the fluctuation amplitude before stabilizing to the limit cycle. Physically, this may be because the local interactions in the PDE system “equilibrate” to the manifold more rapidly than the global ODE model. Dynamically, the problem is that the system is not restricted to the invariant manifold, as shown in Fig. 5.4b

This can be corrected using methods from dynamical systems theory, in particular a center manifold reduction (see Guckenheimer & Holmes (1983); Wiggins (1990), and the simple example in Sec. 2.3). As shown in Fig. 5.5, stability analysis of the fixed point of the $r = 9$ Galerkin system shows that there are only two unstable eigenvectors (consistent with stability of the flow itself). The fundamental assumption of the invariant manifold reduction is that the stable directions quickly decay towards a manifold on which the dynamics can be expressed in terms of the two “active” coordinates. The other degrees of freedom may be kinematically relevant (i.e. to reconstruct the full flow field), but are “slaved” to the active coordinates.

In other words, the full state is decomposed into active coordinates a_1, a_2 with the remaining stable directions a_i an algebraic function of the active coordinates, so that $a_j \approx \varphi_j(a_1, a_2)$. This equation determines the geometry of the manifold. Assuming that the deformation is relatively small, a power series expansion of the manifold equation implies it is parabolic to leading order. With the further assumption that the dynamics are restricted to the manifold, we can derive the expansion coefficients for the manifold equation as follows.

Since the base flow in the expansion (5.1) is a steady state, the origin should be a fixed point of the Galerkin system. The first step in the center manifold reduction is to partition the system into slow and fast degrees of freedom, where slow degrees of

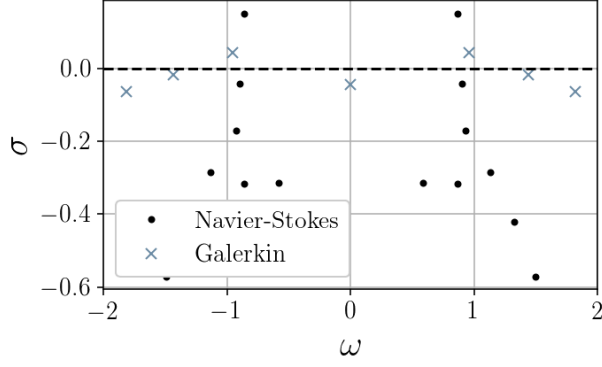


Figure 5.5: **Linear stability of the cylinder flow.** Both the linearized Navier-Stokes and POD-Galerkin systems have one pair of unstable eigenvalues with growth rate $\sigma > 0$. The Galerkin model is based on an expansion around the marginally stable mean flow, so it tends to underestimate the growth rate relative to the true Navier-Stokes equations linearized around the unstable steady state. This can also be seen in the relatively slow growth of perturbations in the Galerkin models shown in Fig. 5.4a.

freedom have eigenvalues with real part $\mathcal{O}(\epsilon)$ for $\epsilon \ll 1$. Assuming the linear operator can be diagonalized with the eigenvector transformation

$$\mathbf{L}\mathbf{V} = \mathbf{V}\mathbf{\Lambda}, \quad (5.2)$$

the coefficient vector can be transformed into generalized eigenvector coordinates with $\mathbf{a} = \mathbf{V}\mathbf{b}$.

This transformation results in a Galerkin model with diagonal linear part, in the form of (2.23). In particular, if \mathbf{V} is invertible then the dynamics of \mathbf{b} are

$$\dot{b}_i = \lambda_i b_i + \sum_{j,k=1}^r Q_{ijk}^\Lambda b_j b_k, \quad (5.3)$$

where $Q_{ijk}^\Lambda = V_{il}^{-1} Q_{lmn} V_{mj} V_{nk}$ is the quadratic tensor transformed into eigenvector coordinates.

In this case there are two weakly unstable eigendirections (active coordinates), as shown in Fig. 5.5. Following the scale separation arguments in Chapter 2, we can separate the state space into active and stable variables \mathbf{b}_a and \mathbf{b}_s , respectively, with

$$\mathbf{b}_a = [b_1 \quad b_2 \quad 0 \quad \cdots \quad 0]^T \quad (5.4)$$

$$\mathbf{b}_s = [0 \quad 0 \quad b_3 \quad b_4 \quad \cdots]^T \approx \boldsymbol{\varphi}(\mathbf{b}_a) \quad (5.5)$$

$$\mathbf{b} \approx \tilde{\mathbf{b}} = \mathbf{b}_a + \boldsymbol{\varphi}(\mathbf{b}_a). \quad (5.6)$$

We first derive the manifold equation $\boldsymbol{\varphi}$ on the fast time scale by neglecting the small real part of the unstable eigenvalues so that $\lambda_{1,2} = \pm i\omega_0$ and \mathbf{b}_a is on a center manifold. The fast variables can then be eliminated from the dynamics on the slow time scale via substitution of the manifold equation.

As introduced in Chapter 2, the dynamic consistency condition for the manifold equa-

tion is that

$$\dot{\mathbf{b}}_s = \frac{d}{dt}\varphi(\mathbf{b}_a) = \mathbf{J}(\mathbf{b}_a)\dot{\mathbf{b}}_a \quad (5.7)$$

$$\mathbf{f}_s(\tilde{\mathbf{b}}) = \mathbf{J}_\varphi(\mathbf{b}_a)\mathbf{f}_a(\tilde{\mathbf{b}}), \quad (5.8)$$

where $\mathbf{J}_\varphi(\mathbf{b}_a)$ is the Jacobian of φ evaluated at \mathbf{b}_a and $\mathbf{f} = [\mathbf{f}_a^T \ \mathbf{f}_s^T]^T$. This consistency condition is a system of PDEs for the function φ that is generally difficult to solve exactly.

However, assuming the coordinates are small, the manifold equation can be expanded in a power series expansion. Requiring that the manifold equation preserves the fixed point and is tangent to the center manifold at the origin implies that $\varphi(0) = 0$ and $\mathbf{J}_\varphi(0) = 0$. Then the first nontrivial terms in the power series are quadratic:

$$\varphi_j(b_1, b_2) = \Phi_{j11}b_1^2 + \Phi_{j12}b_1b_2 + \Phi_{j22}b_2^2 + \dots, \quad j = 3, 4, \dots, r. \quad (5.9)$$

The j^{th} row of the Jacobian (for $j = 3, 4, \dots, r$) is

$$\mathbf{J}_j(b_1, b_2) = [2\Phi_{j11}b_1 + \Phi_{j12}b_2 \quad 2\Phi_{j22}b_2 + \Phi_{j12}b_1] \quad (5.10)$$

and the consistency condition can be expanded as

$$\lambda_j \sum_{k,\ell=1,2} \Phi_{jk\ell}b_kb_\ell + \sum_{k,\ell=1,2} Q_{jk\ell}^\Lambda b_kb_\ell = i\omega_0(2\Phi_{j11}b_1 + \Phi_{j12}b_2)b_1 - i\omega_0(2\Phi_{j22}b_2 + \Phi_{j12}b_1)b_2. \quad (5.11)$$

Comparing powers of \mathbf{b}_a leads to simple closed-form formulae for the expansion coefficients:

$$\Phi_{j11} = \frac{Q_{j11}^\Lambda}{2i\omega_0 - \lambda_j}, \quad \Phi_{j12} = -\frac{Q_{j12}^\Lambda}{\lambda_j}, \quad \Phi_{j22} = \frac{Q_{j22}^\Lambda}{-2i\omega_0 - \lambda_j}. \quad (5.12)$$

Finally, retaining the leading-order correction, the invariant manifold reduction is a generalized Stuart-Landau equation:

$$\dot{b}_j = \lambda_j b_j + \sum_{k,\ell=1,2} Q_{jk\ell}^\Lambda b_k b_\ell + \sum_{k,\ell,m=1,2} \sum_{n=3}^r (Q_{jkn}^\Lambda + Q_{jnk}^\Lambda) \Phi_{nlm} b_k b_\ell b_m. \quad (5.13)$$

This model is two-dimensional, but contains a leading-order approximation to the higher-order modes. As Fig. 5.4 shows, the manifold equation prevents the overshoot of the standard Galerkin model. Although the analytic manifold reduction presented here is an effective technique for stabilizing the cylinder wake, its applicability is somewhat limited in general for two reasons.

First, it relies on the diagonalization and expansion around a fixed point. This not a problem for simple configurations where the Galerkin model preserves fixed points of the fluid flow, or if the steady state is used as the base of the modal expansion. However,

when the flow is far from a bifurcation point it is often better to expand around the mean flow; the resulting Galerkin models may not necessarily have fixed points that coincide with meaningful states of the underlying system.

Second, the invariant manifold technique relies theoretically on smallness of the coefficients (justifying the power series expansion) and the assumption that the higher-order modes can be expressed as instantaneous algebraic functions of the active coefficients. The elimination of constant and linear terms in the parabolic approximation furthermore requires that the real parts of the slow eigenvalues vanish, or at least is small enough to be neglected on the fast time scale. For more complex dynamics far from bifurcation points none of these is necessarily justified. Later chapters propose solutions to these limitations, including data-driven invariant manifold modeling (Chapter 6) and replacing the manifold equation with a stochastic average (Chapter 7).

5.5 SINDy model

Although the invariant manifold model avoids the energy overshoot problem of the standard POD-Galerkin models, it does not solve the problem of overestimating the limit cycle energy (Fig. 5.4a, blue). This is primarily because the power series approximation of the manifold equation is only a first-order approximation to the influence of the higher-order terms. As Fig. 5.6 shows, the parabolic equation roughly approximates the second harmonic, but capturing the third and fourth harmonics would require cubic and quartic Taylor series expansions of the manifold equation.

This is possible in principle, but as the derivation in Sec. 5.4 shows, even the leading-order approximation is relatively involved and the complexity grows rapidly in the polynomial degree. Since we have a sense of the structure of the model, but deriving the exact terms and coefficients analytically is impractical, this problem is a natural candidate for sparse Galerkin regression (SINDy), introduced in Brunton *et al.* (2016a) and reviewed in Sec. 4.2. With this method we look for sparse polynomial representations of both the leading-order dynamics \dot{a}_1, \dot{a}_2 and the manifold equation $a_i = \varphi_i(a_1, a_2)$ for $i = 3, \dots, r$.

Using libraries of polynomials up to third order for the dynamics and fourth order for the manifold, we find

$$\frac{d}{dt} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} = \begin{bmatrix} \lambda_1(1 - \mu_1 r^2) & -\omega(1 + \sigma r^2) \\ \omega(1 + \sigma r^2) & \lambda_2(1 - \mu_2 r^2) \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} \quad (5.14)$$

with $r^2 = a_1^2 + a_2^2$. This is consistent with the normal form for the Hopf bifurcation (2.33) and mean field theory; the linear growth rate is modulated by the perturbation amplitude r^2 until the two are in balance on the limit cycle. We also find that the frequency has a weak amplitude dependence, as can also be observed in the DNS data in Fig. 5.2.

Figure 5.6 shows that the fourth-order SINDy approximation to the manifold equation allows accurate reconstruction of the higher-order modes. The full evolution of the SINDy

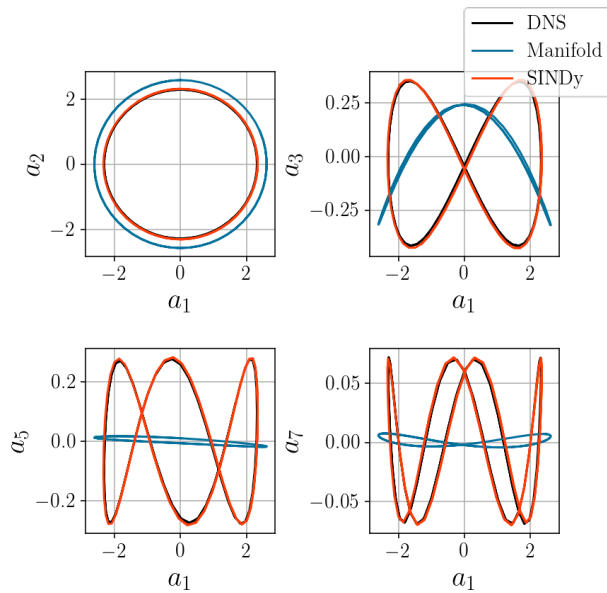


Figure 5.6: Resolving higher harmonics. The analytic invariant manifold reduction computes a quadratic approximation to the manifold, resolving the second harmonic a_3, a_4 . However, a cubic SINDy model with a fourth-order sparse polynomial manifold equation can capture the higher harmonics. The resulting model achieves significantly better accuracy (Fig. 5.4).

model reaches the correct limit cycle energy without overshoot (Fig. 5.4a, orange) by preserving manifold invariance (Fig. 5.4b).

5.6 Summary

The cylinder flow has proven to be a rich test case for the development of reduced-order modeling methodology. It is relatively simple to implement numerically, easy to validate against computational and experimental results, and the flow physics is now well-understood. Still, as illustrated here a full description of the transient dynamics touches a wide range of topics, including linear and nonlinear stability, modal analysis, and dynamical systems theory. So although this section largely reviews existing results (in particular Noack *et al.* (2003); Loiseau *et al.* (2018a)), the example is both demonstrative of much of the conceptual background of Chapters 2-4 and is a good example of fusing physical analysis with targeted applications of simple machine learning methods.

In this case ultimately the SINDy model is both more accurate and significantly simpler than the nearest analytic equivalent (the invariant manifold model). At the same time, it preserves the same deep connection to flow physics and the dynamical systems perspective. Similar methods have recently been used to model open cavity flow (Loiseau & Brunton, 2018; Callaham *et al.*, 2022), chaotic convection (Loiseau, 2020; Guan *et al.*, 2021), magnetohydrodynamics (Kaptanoglu *et al.*, 2021b), and bifurcations in the “fluidic pinball” system (Deng *et al.*, 2020).

Later chapters exploring flows with more complicated dynamics will return to many of the concepts illustrated here on the simpler cylinder wake. In particular, Chapter 6 expands on the idea of empirical manifold modeling, generalizing beyond the case where

the higher-order modes are pure harmonics of the fundamental. Chapter 7 replaces the manifold model with a more general stochastic averaging procedure that results in models with similar cubic structure. Finally, Chapters 8 and 9 introduce a stochastic variant of SINDy and apply it to experimental measurements of a turbulent wake in order to develop a description of mean flow deformation analogous to that presented here for the cylinder wake.

Part III
NONLINEAR DYNAMICS

Chapter 6

EMPIRICAL INVARIANT MANIFOLD REDUCTION

The discussion and examples in Part II of this thesis suggest that many unsteady flows are characterized by emergent low-dimensional structure, proper treatment of which can substantially improve reduced-order models. In particular, when the flow is advection-dominated, the modal coefficients can exhibit strong nonlinear correlations related to the traveling-wave behavior of the PDE solutions. In this chapter we will see that nonlinear dimensionality reduction can identify and enforce this latent structure, as shown schematically in Fig. 6.1. Section 6.1 begins the chapter with an analysis of two model PDEs, demonstrating the connection between advection and nonlinear correlation.

Sections 6.2- 6.5 explore this relationship with an in-depth analysis of quasiperiodic shear-driven cavity flow. After projecting data from a direct numerical solution onto a basis of DMD modes, the recently proposed randomized dependence coefficient (Lopez-Paz *et al.*, 2013) allows us to clearly distinguish the active degrees of freedom from correlated higher harmonics and nonlinear crosstalk. In this minimal representation, the dynamics occur on a 2-torus, while the rest of the modes, which arise as triadic interactions of the active variables in the frequency domain, can be expressed as polynomial functions of the dynamically active variables. The restriction to this manifold stabilizes a standard POD-Galerkin model, avoiding both decoherence and energy imbalance. This representation is also a natural basis for data-driven system identification methods; we apply the sparse identification of nonlinear dynamics (SINDy) algorithm (Brunton *et al.*, 2016a) and show that the flow can be accurately described by two independent Stuart-Landau equations.

The view of the nonlinear correlations as triadic interactions in the frequency domain is only valid for flows with discrete spectral content (i.e. periodic or quasiperiodic dynamics), but the problem of linear modal decompositions overestimating the number of dynamically active degrees of freedom is relevant for advection-dominated flows with more complex dynamics. In this vein, Sec. 6.6 generalizes the polynomial manifold approximation with a deep learning model, extending the validity of the empirical manifold modeling approach to the chaotic regime. A parametric deep manifold model constructed from simulations of the cavity configuration at two Reynolds numbers is also able to reproduce the series of bifurcations in the Ruelle-Takens route to chaos (Fig. 1.2).

6.1 Advection and nonlinear correlation

Many features of projection-based models of advection-dominated flows are demonstrated by simple scalar PDEs. In particular, limitations of the Galerkin representation of hyper-

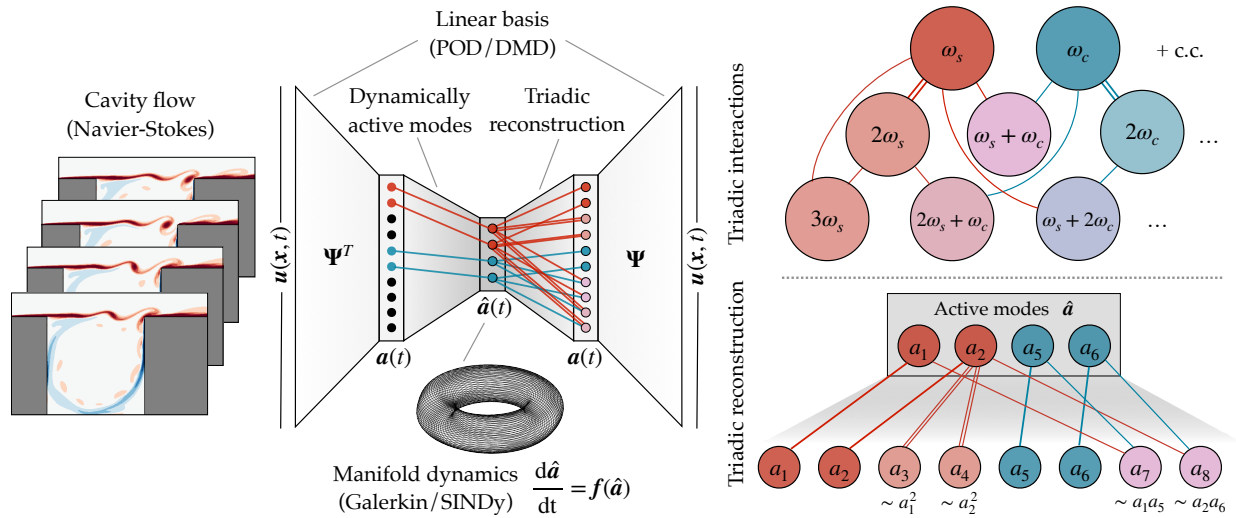


Figure 6.1: **Schematic of the model reduction approach exploiting nonlinear correlations.** The flow fields are first projected onto a linear modal basis Ψ , yielding modal coefficients $a(t)$. The quasiperiodic dynamics can be described by four degrees of freedom; the rest of the modal coefficients can then be reconstructed with polynomial functions consistent with triadic interactions in the frequency domain. The dynamics of the active degrees of freedom can be modeled either by restricting the POD-Galerkin dynamics to the toroidal manifold or by identifying a simple, interpretable dynamical system with the SINDy algorithm.

bolic problems can be seen in the linear constant-coefficient advection equation, while Burgers' equation is a minimal example of the key role of nonlinearity in the full Navier-Stokes equations.

The linear dispersion relation as nonlinear correlation One of the fundamental reasons that Galerkin models of advection-dominated flows tend to be fragile is that they introduce additional variables that do not correspond to physical degrees of freedom. This is perhaps illustrated most clearly by the linear advection equation on a periodic domain:

$$u_t + cu_x = 0, \quad x \in (0, L). \quad (6.1)$$

For any initial condition $u(x, 0) = u_0(x)$, this equation has the simple traveling wave solution $u(x, t) = u_0(x - ct)$. Given the initial condition, the only effective degree of freedom is the phase ct/L . However, the problem could also be solved by means of a Fourier expansion

$$u(x, t) = \sum_{n=-\infty}^{\infty} a_n(t) e^{ik_n x} \quad (6.2)$$

with $k_n = 2\pi n/L$. The Galerkin system in this orthogonal basis (see Chapter. 4) is

$$\dot{a}_n(t) = -i\omega_n a_n(t), \quad \omega_n = k_n c n = 0, \pm 1, \pm 2, \dots \quad (6.3)$$

The relationship between frequency ω and wavenumber k_n is the *dispersion relation*; in this case it implies that all scales are carried at the same speed c .

With this analytic dispersion relation, (6.3) is equivalent to the traveling wave solution, since $a_n(t) = e^{-i\omega_n t} a_n(0)$:

$$u(x, t) = \sum_{n=-\infty}^{\infty} a_n(0) e^{ik_n(x-ct)} = u_0(x - ct). \quad (6.4)$$

However, the Galerkin model has introduced many degrees of freedom in the harmonics a_n by artificially separating space and time. If the projection is approximated numerically or with empirical basis modes, the estimated system may include some error, so that $\omega_n = k_n c + \epsilon_n$. In this case the Galerkin system will be dispersive, i.e. each wavenumber will propagate with a slightly different speed. The traveling wave solution will tend to lose coherence on a timescale $1/\epsilon$, as shown in figure 6.2.

An alternative perspective on the dispersion relation is that it specifies nonlinear correlations between the temporal coefficients a_n , removing the spurious degrees of freedom introduced by Galerkin projection. The linear dispersion relation $\omega_n = nk_1 c$ implies the nonlinear relationship for harmonics

$$a_n(t) = e^{-ink_1 ct} a_n(0) \propto a_1^n, \quad (6.5)$$




Figure 6.2: Linear advection equation with errors $\epsilon_n \sim \mathcal{N}(0, \epsilon^2)$ in the dispersion relation $\omega_n = ck_n$. The Galerkin model (grey) loses coherence with the exact solution (black) over a time scale $1/\epsilon$. If the polynomial correlations implied by the dispersion relation are enforced explicitly, the model is robust to such errors. Nonlinear correlation in the true system, given by (6.5), appears in the Lissajous-type phase portraits of the Fourier coefficients (right). Similar behavior manifests in Galerkin models of nonlinear advection-dominated flows.

with the proportionality determined by the initial condition. Then the only degree of freedom is a_1 , and the traveling wave solution is recovered by the Galerkin model projected onto this mode. In dynamical systems terminology, the solution is restricted to a one-dimensional manifold: a circle representing the phase of the leading Fourier coefficient. In this case the decoherence does not lead to instability because the system is purely linear with purely imaginary eigenvalues, but in nonlinear systems with nonzero linear growth rates the departure from the solution manifold can be catastrophic.

Triadic interactions and the energy cascade For more general linear systems the preceding analysis is complicated by non-normality and physical dispersion, and the concept of a dispersion relation is not well-defined for nonlinear dynamical systems. Nevertheless, analogous concepts are similarly important in models of nonlinear PDEs. For example, Burgers' model is a paradigmatic scalar conservation equation illustrating many features of gas dynamics and nonlinear flows more broadly. Burgers' equation with viscosity ϵ is

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} = \epsilon \frac{\partial^2 u}{\partial x^2}, \quad x \in (0, 2\pi). \quad (6.6)$$

On a periodic domain, we can apply the same Fourier expansion (6.2) with $L = 2\pi$, leading to the Galerkin ODE system

$$\dot{a}_k = -\epsilon k^2 a_k - ik \sum_{\ell=-\infty}^{\infty} a_\ell a_{k-\ell}, \quad k = 0, \pm 1, \pm 2, \dots \quad (6.7)$$

The two right-hand side terms in (6.7), originating from the viscous and nonlinear PDE terms respectively, capture several key features of the full Navier-Stokes equations.

First, the convolution-type sum over wavenumbers ℓ includes only pairs that sum to k ; these are the so-called “triadic” scale interactions. Second, it can be shown that the nonlinear term is energy-preserving in the sense that when the energy $a_k a_k^*$ is summed over all wavenumbers the nonlinear term does not contribute to a net change in energy of the system¹. This suggests that the only role of nonlinearity is to transfer energy between scales. Meanwhile, the dissipation rate of each mode scales quadratically with wavenumber so that the bulk of dissipation occurs at the smallest scales.

The overall picture of the dynamics in the spectral domain is therefore that the nonlinear term transfers energy from the more energetic large scales to the dissipative small scales. Since (6.7) is very similar to the spectral form of the momentum equations for isotropic turbulence (see Chapter 2), this “energy cascade” is an important feature of real viscous flows as well. The energy cascade points to another often-discussed issue with Galerkin models: if the system is truncated at a wavenumber r which is not sufficiently large to capture the net dissipation rate, the energy cascade is interrupted and the system of ODEs will overestimate the energy, potentially even becoming unstable.

This issue is fundamentally different from the decoherence discussed in the context of the linear advection equation. For example, the issue of fine-scale dissipation is also present in the heat equation, given by equations (6.6)-(6.7) without the convective nonlinearity. Whereas the Fourier-Galerkin representation of advection introduces spurious degrees of freedom, this discussion suggests that in the representation of the heat equation *all* coefficients are dynamically important (self-similarity notwithstanding). The Galerkin system is therefore an ideal representation of the parabolic dynamics of the heat equation, where the fundamental assumption of separation of variables is valid.

The inability of Fourier decomposition and POD to produce efficient representations of traveling wave physics has long been recognized. Fundamentally, these decompositions rely on a space-time separation of variables, which is not a valid assumption for traveling waves. Many extensions to POD have been developed for translationally invariant systems and systems with other symmetries (Rowley & Marsden, 2000; Reiss *et al.*, 2018; Rim *et al.*, 2018; Mendible *et al.*, 2020).

For general viscous, nonlinear, advection-driven fluid flows, we might expect advection, triadic interactions, and small-scale dissipation to all be relevant as a result of the joint hyperbolic-parabolic structure of the Navier-Stokes equations. The intrinsic dimensionality of the system and, conversely, the inaptitude of the Galerkin model, may not be *a priori* clear as a result of a complex interplay between these mechanisms.

For example, if the leading degree of freedom a_1 tends to oscillate at a frequency ω_0 , representing either a standing or traveling wave, then the a_2 dynamics include a term of the form $a_1^2 \sim e^{2i\omega_0 t}$. Similarly, $\dot{a}_3 \propto a_1 a_2 \sim e^{3i\omega_0 t}$. In the energy cascade picture, these higher-order modes act as forced, damped oscillators and will tend to respond at the forcing frequencies. In this manner, triadic interactions in the wavenumber domain

¹A similar result holds for inhomogeneous flows (Schlegel & Noack, 2015)

can also give rise to nonlinear correlations in time and triadic structure in the frequency domain.

This effect is not necessarily limited to systems with spatial or temporal periodicity; Rubini *et al.* (2020) investigated the application of system identification methods to a chaotic lid-driven cavity flow and showed that sparse nonlinear coupling, analogous to triadic interactions, were critical for resolving energy transfers across scales of the flow. They distinguish between this empirical, *a posteriori* sparsity appearing in the statistics and data-driven models, and the structural, *a priori* sparsity of systems such as (6.7), which is generally lost in Galerkin models of inhomogeneous flows. Similarly, Schmidt (2020) recently proposed a bispectral modal analysis technique that leverages approximate sparsity in frequency interactions.

In analogy with the dispersion relation, these processes may result in latent structure not immediately obvious in the Galerkin representation. If this structure is ignored, the behavior of the model may depart significantly from that of the underlying system. For example, Majda & Timofeyev (2000) showed that a truncated Galerkin model of the inviscid Burgers equation tends towards equipartition of energy rather than a physical solution as a result of a catastrophic decoherence mechanism. Consequentially, in the following sections we argue that nonlinear correlation and manifold restriction plays an important role in the stability and accuracy of reduced-order models of advection-dominated fluid flows.

6.2 Shear-driven open cavity flow

The flow considered in the present work is the incompressible shear-driven cavity flow visualized in figure 6.3. It is a geometrically-induced separated boundary layer flow having a number of applications in aeronautics (Yu, 1977) or for mixing purposes (Chien *et al.*, 1986). The leading two-dimensional instability of the flow is localized along the shear layer delimiting the outer boundary layer flow and the inner cavity flow (Sipp & Lebedev, 2007; Sipp *et al.*, 2010). This oscillatory instability relies essentially on two mechanisms. First, the convectively unstable nature of the shear layer causes perturbations to grow as they travel downstream. Once the perturbations impinge upon the downstream corner of the cavity, instantaneous pressure feedback re-excites the upstream portion of the shear layer. Coupling of these mechanisms gives rise to a linearly unstable feedback loop at sufficiently high Reynolds numbers ($Re_c \gtrsim 4120$, see Sipp & Lebedev (2007)).

A similar unstable loop exist for compressible shear-driven cavity flows, wherein the instantaneous pressure feedback is replaced by upstream-propagating acoustic waves (Rossiter, 1964; Rowley *et al.*, 2002; Yamouni *et al.*, 2013). At higher Reynolds numbers, the slowly recirculating flow inside the cavity can also perturb the shear layer. This inner cavity mode is similar in spatial structure and oscillation frequency to those observed in two-dimensional lid-driven cavity flows (Arbabi & Mezić, 2017). Since the shear layer instability and inner cavity recirculation occur at incommensurate frequencies, the non-




Figure 6.3: **Open cavity flow configuration** Computational domain and representative instantaneous vorticity field for the shear-driven cavity flow at $Re = 7500$ highlighting the vortical structures developing along the shear layer.

linear coupling between these modes leads to quasiperiodic dynamics, as illustrated in figure 6.4.

Despite its apparent simplicity, this strictly two-dimensional linearly unstable flow configuration has served multiple purposes over the past decade: illustration of optimal control and reduced-order modeling (Barbagallo *et al.*, 2009; Loiseau *et al.*, 2018b; Leclercq *et al.*, 2019), investigation of the nonlinear saturation process of flow oscillators (Sipp & Lebedev, 2007; Meliga, 2017), or as an introduction to dynamic mode decomposition (Schmid, 2010). Recent work has also explored the linear stability of its three-dimensional counterpart, in particular the influence of spanwise end-walls (Liu *et al.*, 2016; Picella *et al.*, 2018).

The dynamics of the flow are governed by the incompressible Navier-Stokes equations (2.1). The Reynolds number is set to $Re = 7500$ based on the free-stream velocity U_∞ and the depth L of the open cavity. The computational domain and boundary conditions considered herein are the same as in Sipp & Lebedev (2007); Sipp *et al.* (2010); Loiseau *et al.* (2018b); Bengana *et al.* (2019); Leclercq *et al.* (2019), shown schematically in figure 6.3.

We perform direct numerical simulation (DNS) of the flow with the Nek5000 spectral element solver (Fischer *et al.*, 2008). The mesh consists of 6100 seventh-order spectral elements, equivalent to roughly 3.8×10^5 grid points, refined towards the walls and shear layer. The domain is therefore somewhat over-resolved compared to similar studies in order to minimize any numerical errors in the Galerkin projection for higher-order modes. Diffusive terms are integrated with third order backwards differentiation, while convective terms are advanced with a third order extrapolation. We retain 30000 snapshots from the DNS at sampling rate $\Delta t = 10^{-2}$, a frequency roughly fifty times larger than the high-frequency oscillation of the shear layer.

Figure 6.3 depicts an instantaneous vorticity field obtained from direct numerical simulation once the flow has reached a statistical steady state. It shows the advection of a vortical structure along the shear layer before it impinges the downstream corner of the cavity. These shear layer oscillations arise as a linear instability mode of the steady base

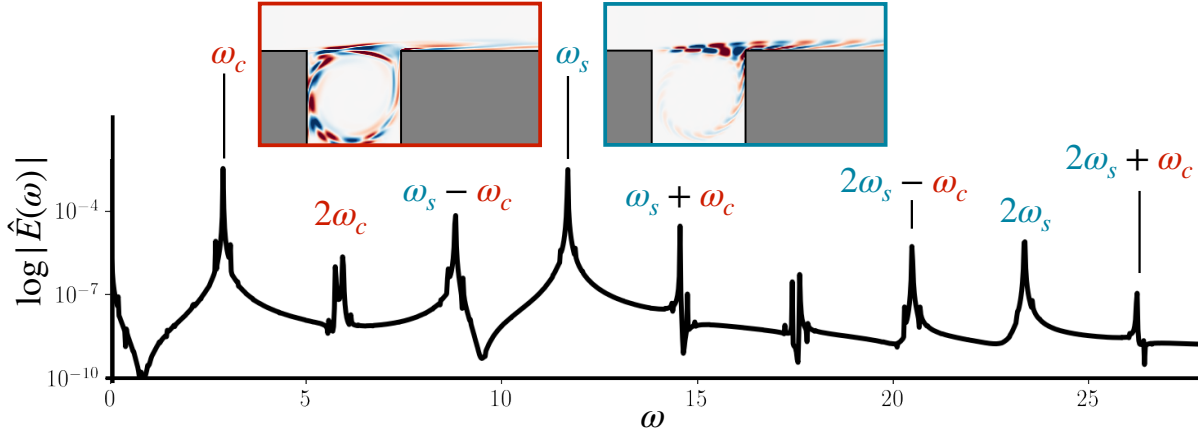


Figure 6.4: **Fourier spectrum** $|\hat{E}(\omega)|$ of the fluctuation's kinetic energy at $Re = 7500$. The high-frequency peak ($\omega_s \simeq 12$) corresponds to the shear layer instability while the low frequency peak ($\omega_c \simeq 3$) are associated to the inner-cavity dynamics. A few other peaks have been labelled based on the quadratic interactions on the two fundamental frequencies for the sake of illustration. Multiple closely-spaced peaks are associated with nearby frequency combinations (e.g. $2\omega_c \approx \omega_s - 2\omega_c \approx 6$). Also shown are the real parts of the DMD modes at ω_s and ω_c .

flow above $Re_c \gtrsim 4120$ (Sipp & Lebedev, 2007). However, the Reynolds number of the present flow ($Re = 7500$) is significantly larger than this critical Reynolds number, so the typical amplitude of fluctuations is not infinitesimal and the associated Reynolds stresses are not negligible.

The physics of this flow are therefore fundamentally nonlinear in at least two respects. First, the growth of the instability modes is checked by the Stuart-Landau nonlinear stability mechanism, in which finite Reynolds stresses deform the steady base flow into the post-transient mean (Sipp & Lebedev, 2007; Meliga, 2017). Second, a stability analysis of the time-averaged mean flow at this Reynolds number reveals a second, weaker instability associated with lower-frequency oscillations inside the cavity; see Sec. 6.3 and Sipp *et al.* (2010). The incommensurate frequencies of these two instabilities give rise to quasiperiodic oscillatory dynamics.

Because the unstable base flow is of limited relevance in the statistically stationary regime, it is more natural to decompose the instantaneous velocity field into a time-averaged mean flow $\bar{\mathbf{u}}(\mathbf{x})$ and zero-mean fluctuations $\mathbf{u}'(\mathbf{x}, t)$. For a detailed analysis of the choice between base- and mean-flow expansions, see Sipp & Lebedev (2007).

Figure 6.4 shows the Fourier spectrum of the kinetic energy of the fluctuating component

$$E(t) = \frac{1}{2} \int_{\Omega} \mathbf{u}'(\mathbf{x}, t) \cdot \mathbf{u}'(\mathbf{x}, t) \, d\Omega \quad (6.8)$$




Figure 6.5: **Eigenspectrum of the Navier-Stokes operator at $Re = 7500$.** The spectrum is estimated in three ways: linearized in the vicinity of the base flow (BF, black circles), mean flow (MF, blue circles), and from dynamic mode decomposition (DMD, red crosses). For the linear stability analyses, only eigenvalues for which a 10^{-8} convergence has been achieved are plotted. These eigenspectra have been computed using a time-stepper Arnoldi algorithm with a sampling period $\Delta T = 0.1$ and a Krylov subspace dimension of 1024 and 512 for the base and mean flows, respectively.

integrated over the domain Ω . Such a spectrum is characteristic of quasiperiodic dynamics, as recently observed for a similar flow by Leclercq *et al.* (2019). As demonstrated below, the two main frequencies correspond either to the dynamics of the vortical structures along the shear layer (ω_s) or to the low-frequency unsteadiness taking place within the cavity (ω_c). The power spectrum consists of approximately discrete peaks, each of which can be accounted for by the sum or difference of these fundamental frequencies and their harmonics. The observation that this spectrum can be generated using only two main frequencies lets us hypothesize that the dynamics of the fluctuation $\mathbf{u}'(\mathbf{x}, t)$ around the mean flow $\bar{\mathbf{u}}(\mathbf{x})$ are amenable to a low-dimensional representation. This flow is therefore of intermediate complexity, between weakly nonlinear flows, which can be accurately described with normal form dynamics, and fully turbulent flows, which have many dynamical degrees of freedom and would likely require careful closure modeling to approximately model the evolution of large-scale coherent structures.

6.3 Reduced-order modeling

Modal analysis Before analyzing the results of the DNS, a global stability analysis offers a useful perspective on the dominant dynamics of the flow. We compute the leading eigenvector-eigenvalue pairs of the Navier-Stokes equations linearized about both the unstable steady state and the mean of the DNS using an in-house Krylov-Schur time-stepping algorithm (Edwards *et al.*, 1994; Stewart, 2001) implemented in the spectral element solver Nek5000 (Fischer *et al.*, 2008).

Figure 6.5 depicts the eigenspectra of the operator linearized about base and mean flows. Four complex conjugate pairs of eigenvalues lie within the upper-half complex

plane, indicating the base flow is strongly unstable. The most unstable eigenvalue is $\sigma \pm i\omega = 0.90 \pm 10.86i$, where σ and ω are the growth rate and circular frequency of the instability mode, respectively.

Although the stability analysis of the base flow provides some insight about the physical origin of the high-frequency shear layer oscillation, there are two main reasons it is insufficient to describe the nonlinear flow. First, there is no trace of the three additional unstable modes predicted by the stability analysis in the DNS. Second, at the Reynolds number considered in this work, linear stability analysis of the base flow is unable to predict the low-frequency inner-cavity oscillation. The higher harmonics are also missing from the stability analysis, though this is to be expected of a linear analysis. To the authors' knowledge, there has not yet been a detailed explanation of the process by which the additional unstable modes are superseded by the low-frequency dynamics. For more details about the shear layer instability and its saturation process at lower Reynolds number, see Meliga (2017).

In this case, the mean flow stability analysis supports the picture suggested by the analysis of the power spectrum; the four active degrees of freedom are related to the two mode pairs corresponding to the shear layer instability and inner cavity oscillation. In the nonlinear DNS, interactions between these modes generate harmonics and frequency crosstalk, although this structure is fully dependent on the active degrees of freedom. Linear stability analysis assumes the perturbations have negligible energy and so it cannot resolve the nonlinear interactions responsible for this behavior.

While the mean flow stability analysis does appear to predict the dominant structures in the DNS, we develop the reduced-order model in an empirical modal basis in order to maximally resolve the fully nonlinear flow. To that end, we first compute the proper orthogonal decomposition (POD) from 4000 fields sampled at $\Delta t = 0.05$, approximately ten times the shear layer frequency, using the method of snapshots (Sirovich, 1987). This is around 13% of the total number of fields retained from the DNS, but is sufficient for statistical convergence of the leading modes.

The singular value spectrum and residual energy are shown in figure 6.6. The singular values converge relatively quickly; the first pair of modes contain 70% of the fluctuation kinetic energy, the first six account for $\sim 90\%$, and by $r = 64$ approximately 99.97% of the energy is recovered. We retain 64 modes for further analysis and note that our modeling results are insensitive to moderate changes in truncation. Still, as we will show in Section 6.4, the intrinsic dimensionality of the system is much smaller than that of the linear subspace required for reconstruction. As with the advection system in Section 6.1, this is partly due to the representation of traveling waves, as shown in Fig. 6.7.

The approximately discrete spectral content of the flow further clarified by a dynamic mode decomposition (DMD) analysis (Sec. 3.3); each DMD eigenvalue (shown in figure 6.8) can be identified with some combination of the fundamental frequencies of modes $k = 1$ and $k = 5$. Based on the results in Sections 6.4 and 6.5, we hypothesize that DMD also filters frequency content and accentuates nonlinear correlations for modes that are




Figure 6.6: **Singular value spectrum of the quasiperiodic cavity flow.** Black dots represent the normalized squared singular values of the snapshot correlation matrix, indicating the fraction of fluctuation kinetic energy resolved by each mode. Red crosses indicate the fraction of residual energy, or normalized cumulative sum of squared singular values. Dashed lines indicate the number of modes retained ($r = 64$).

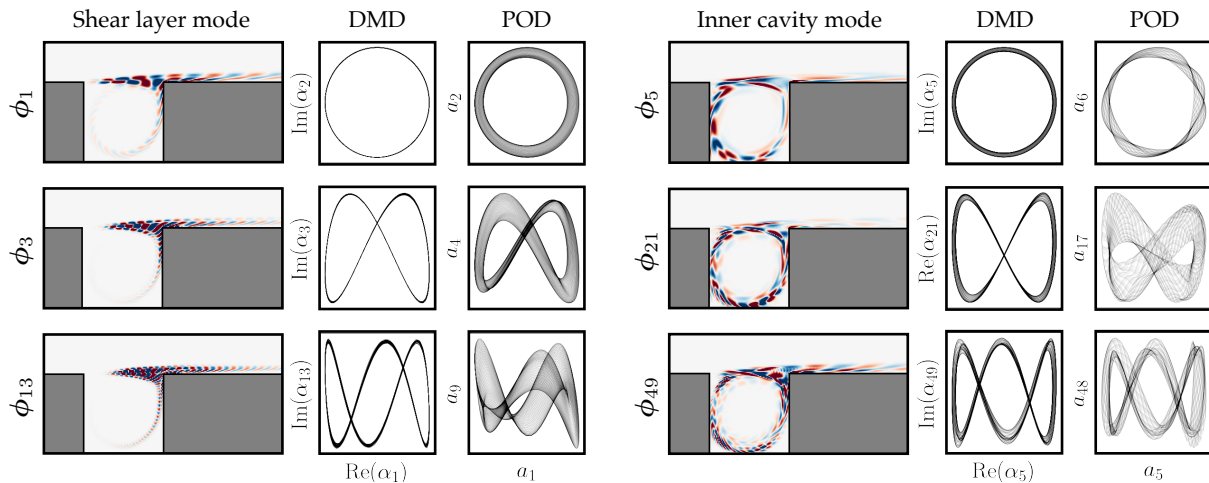


Figure 6.7: **Harmonic modes identified from POD and DMD analysis.** The spatial fields and phase portraits both indicate that certain mode pairs are harmonics arising from the description of wavelike motion in the shear layer and inner cavity. Because DMD is based on both spatial and temporal correlation, this structure is especially pronounced in the DMD coefficients. The vorticity plots are real parts of the DMD modes, but analogous modes exist in the POD basis.

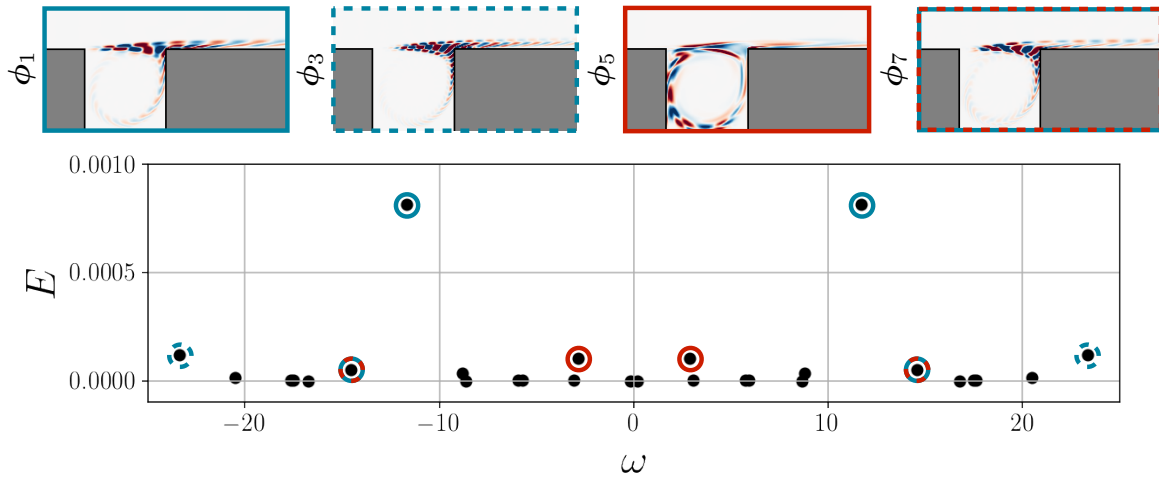


Figure 6.8: DMD frequencies ω_k and average energy $E_k = \langle |\alpha_k|^2 \rangle$ along with vorticity plots for the real part of the most energetic modes. The second mode pair ($k = 3, 4$) is a harmonic of the leading pair ($k = 1, 2$), while the third pair ($k = 5, 6$) represents the low-frequency inner cavity motion. Other modes (e.g. $k = 7, 8$) are either harmonics or indicate nonlinear frequency crosstalk between these leading modes, as in figure 6.4.

Base flow	$0.90 \pm 10.86i$	$0.73 \pm 13.75i$	$0.49 \pm 7.85i$	$0.02 \pm 16.65i$
Mean flow	$0.22 \pm 11.68i$	$0.09 \pm 3.03i$	$0.06 \pm 2.38i$	$0.03 \pm 14.68i$
DMD	$0.00 \pm 11.68i$	$0.00 \pm 23.36i$	$0.00 \pm 2.87i$	$0.02 \pm 3.68i$ height

Table 6.1: Spectrum of the linearized cavity flow. Eigenvalues $\sigma + i\omega$ correspond to the least stable modes in a stability analysis of the base and mean flows, along with DMD eigenvalues for the most energetic modes.

not pure harmonics. This approximate nonlinear algebraic dependence clearly indicates a manifold structure of much lower dimensionality than the linear subspace. Given these implications, the results presented in Secs. 6.4 and 6.5 are based on the DMD expansion (3.22).

The leading eigenvalues of the DMD and global stability analyses are compared in Table 6.1. The dominant frequency predicted by the mean flow stability analysis differs by only 5% to 10% from the dominant peak of the DNS (Figure 6.4) and that given by the DMD analysis. The associated eigenfunction (not shown) also closely resembles the leading DMD mode ϕ_1 shown in Figures 6.7 and 6.8.

POD-Galerkin model Beginning with the POD basis, we compute the linear-quadratic Galerkin system (4.7), as described in Sec. 4.1. Figure 6.9 shows the evolution of the




Figure 6.9: **POD-Galerkin models for open cavity flow.** The time series shows the evolution of the fluctuation kinetic energy predicted by POD-Galerkin reduced-order models of various dimensions along with DNS values. Though all values of r shown here capture sufficient dissipation to remain at finite energy, none resolves the true quasiperiodic dynamics.

fluctuation kinetic energy as predicted by the POD-Galerkin system for various levels of truncation r . Although the estimate does tend to improve with increasing r , none of these models capture the quasiperiodic dynamics of the flow, and most exhibit significant instability. This is true despite (and, we will argue, because of) the fact that these models have many more kinematic degrees of freedom than the true dynamics underlying the post-transient cavity flow.

We also transform the model into the DMD coordinate system $\alpha(t)$, effectively computing a Petrov-Galerkin projection in the DMD basis. The dynamics of the two models are identical at this stage because they are related through a similarity transform. However, for the secondary model reduction based on nonlinear correlations analysis in Sec. 6.4, the DMD-based model is able to more clearly isolate triadic interactions in the frequency domain.

6.4 Nonlinear correlations analysis

With the modal expansion and model reduction described in Sec. 6.3, the formally infinite-dimensional PDE can be reduced to a set of coupled, nonlinear ODEs mimicking the structure of the physical system. However, the POD analysis indicates that dozens of modes are necessary for an approximately complete kinematic representation in a linear basis, while the DMD analysis and power spectrum suggest that the dynamics of the flow are quasiperiodic. A minimal description of the post-transient flow should therefore require only a pair of oscillators evolving on a 2-torus, comprising four degrees of freedom.

This discrepancy can be understood qualitatively in light of the discussion in Sec. 6.1. Advection of nearly periodic fluctuations leads to the appearance of harmonic modes,

as shown in Figs. 6.7 and 6.8. Similarly, crosstalk between the incommensurate dominant frequencies gives rise to modes that are not pure harmonics of either frequency. In the low-dimensional subspace spanned by the leading POD/DMD modes, these can be viewed as triadic interactions in the frequency domain. Again, since we seek structure that is coherent at distinct frequencies, we will focus on the DMD coefficients in the following.

A model quasiperiodic cascade Consider a model system of ODEs including cascading triadic-type interactions of the form of Eq. (6.7) where self-sustaining oscillations drive higher-order degrees of freedom:

$$a_1 = Ae^{i\omega_a t}, \quad b_1 = Be^{i\omega_b t} \quad (6.9a)$$

$$\dot{a}_k = \sum_{|\ell| < k} a_\ell a_{k-\ell}, \quad \dot{b}_k = \sum_{|\ell| < k} b_\ell b_{k-\ell}, \quad \dot{c}_k = \sum_{|\ell| < k} a_\ell b_{k-\ell}, \quad k > 1 \quad (6.9b)$$

For $k = 2$ the only interactions are at $|\ell| = 1$, generating second harmonics for a_2 and b_2 , and oscillations at $\omega_a + \omega_b$ for c_2 . At $k = 3$, the forcing includes the $k = 2$ terms, generating third harmonics for a_3 and b_3 and crosstalk for c_3 . Up to some complex scaling the solutions take the form

$$a_2 = A^2 e^{2i\omega_a t}, \quad b_2 = B^2 e^{2i\omega_b t}, \quad c_2 = AB e^{i(\omega_a + \omega_b)t}, \quad (6.10a)$$

$$a_3 = A^3 e^{3i\omega_a t}, \quad b_3 = B^3 e^{3i\omega_b t}, \quad c_3 = AB^2 e^{i(\omega_a + 2\omega_b)t} + A^2 B e^{i(2\omega_a + \omega_b)t}. \quad (6.10b)$$

By $k = 3$ the response of the c_k crosstalk variable does not represent pure frequency content. If these represented modal coefficients we would expect the dynamics at (for instance) $\omega_a + 2\omega_b$ and $2\omega_a + \omega_b$ correspond to different spatial structures so c_3 and higher-order terms would be separated into distinct coefficients. This serves to emphasize that the temporal coefficients and low-dimensional ODEs are only convenient representations of the full spatiotemporal dynamics and are not fundamental physical quantities.

A global observable such as the energy analogue $\sum_k (|a_k|^2 + |b_k|^2 + |c_k|^2)$ will have frequency content at all integer combinations of ω_a and ω_b . As with the linear advection equation in Section 6.1, these higher-order coefficients can be expressed as algebraic functions of the fundamental oscillators. For instance, the crosstalk variable c_3 can be replaced by $a_1 b_1^2 + a_1^2 b_1$, up to scaling. This example shows that periodic oscillation with cascading triadic interactions can generate quasiperiodic time series, point spectra similar to Fig. 6.4, and nonlinear algebraic dependence without direct interactions between the oscillators, provided the dominant oscillation frequencies are incommensurate.

The randomized dependence coefficient As discussed in Sec. 3.2, the POD coefficients are guaranteed to be linearly uncorrelated. The same is not necessarily true of DMD coefficients, but in practice they tend to be minimally correlated. However, as in the




Figure 6.10: Nonlinear correlations in modal coefficients. Subplots show phase portraits of DMD coefficients along with measures of linear and nonlinear correlation (Pearson’s ρ and the randomized dependence coefficient, respectively). While α_1 and α_3 are linearly uncorrelated ($\rho = 0$), the clear functional relationship between the two is reflected in the RDC value; physically, α_3 is a pure harmonic of α_1 . On the other hand, α_{17} corresponds to a nonlinear crosstalk mode that has no clear correlation with α_1 , either linear or nonlinear. Nevertheless, it can be accurately approximated by a simple polynomial function of the active degrees of freedom, as shown by Table 6.2 and Fig. 6.12.

previous example, a network of triadic interactions forced by a limited number of driving oscillators can exhibit pure algebraic dependence on the active degrees of freedom. In other words, the higher-order variables can have perfect nonlinear correlation, even when uncorrelated in a linear sense.

This is an intuitive result, but challenging to evaluate in a principled way. In a probabilistic setting, mutual information is the most natural metric for generalized correlation, but it requires estimating integrals over conditional probability distributions. This is expensive and difficult for multidimensional signals, and the concept of mutual information itself is not necessarily well-suited for purely deterministic systems. To address this issue, various nonlinear generalizations of the standard (Pearson’s) linear correlation coefficient have been proposed in the statistics community. Recently, Lopez-Paz *et al.* (2013) proposed the *randomized dependence coefficient* (RDC) as an efficient and convenient metric for nonlinear correlation that has the properties defined by R enyi (1959) for generalized measures of dependence between variables.

The RDC combines linear canonical correlations analysis with randomized nonlinear projections to estimate nonlinear dependence; details are presented in Lopez-Paz *et al.* (2013). Figure 6.10 gives examples of both the standard linear correlation coefficient and the RDC for several pairs of DMD coefficients. Coefficient pairs that are pure harmonics tend to score highly on the RDC, while coefficients with nonlinear crosstalk or multiple frequency components do not, even if there is a simple functional dependence on two active coefficients. This limits the potential of the RDC for unraveling the multivariate structure of the manifold, although it is a convenient diagnostic and visualization tool for identifying independent, dynamically active modes.

For example, Fig. 6.11 shows both the phase portraits of leading DMD coefficients (lower triangular portion) and the RDC values (upper triangular). In some cases (colored




Figure 6.11: Identification of the active degrees of freedom with the randomized dependence coefficient (RDC). The lower triangular portion of the figure shows phase portraits of the real (horizontal axes) and imaginary (vertical axes) DMD coefficients, while the upper triangular portion depicts the RDC values scaled linearly in color and radius. Two approximately independent clusters can be identified: the shear layer dynamics (blue) and inner cavity oscillations (red). Each of these is associated with a dominant mode pair (solid borders) and pure harmonics (dashed borders) that are strongly nonlinearly correlated with the dominant modes. The other modes also have simple polynomial relationships with the active degrees of freedom but include cross terms that break the one-to-one nonlinear correlation (see Table 6.2 and Figs. 6.10 and 6.12).

outlines), the modes are clearly pure harmonics of one of the two driving mode pairs. This is reflected in the large values of the RDC for the harmonics, indicating that these coefficients can be directly expressed as algebraic functions of one or the other driving modes. However, based on the previous discussion we expect that coefficients representing frequency crosstalk might be multivariate functions of both driving mode pairs. In this case it is clear based on the energy content and harmonic structure of Fig. 6.11 that the pairs (α_1, α_2) and (α_5, α_6) are the driving degrees of freedom, but chaotic or turbulent flows might have more opaque causal structure.

Manifold reduction via sparse regression Based on the RDC analysis, it is clear that certain modes are pure algebraic functions of one or the other driving mode pairs. In particular, harmonic modes such as those illustrated in Fig. 6.7 are polynomial functions of the fundamental mode. However, as demonstrated by the model quasiperiodic cascade, triadic interactions can also generate multivariate nonlinear correlations for modes representing frequency crosstalk. In addition, from a dynamical systems perspective we expect that a quasiperiodic system with two dominant frequencies should have a post-transient attractor described by four degrees of freedom: two generalized amplitude and phase pairs.

If this is the case, the modes that are not pure harmonics may still be approximately polynomial functions of the shear layer and inner cavity modes. This can be viewed as a data-driven generalization of the analytic center manifold reduction in Sec. 2.3. We explore this hypothesis with the same approach as outlined in Sec. 4.2 for the SINDy algorithm. Denoting these “active” degrees of freedom by $\hat{\alpha}(t) = [\alpha_1 \ \alpha_2 \ \alpha_5 \ \alpha_6]^T$, the library $\Theta(\hat{\alpha})$ is defined as in Eq. (4.12). We assume the full coefficient vector can be approximated as

$$\alpha \approx \Phi \Theta(\hat{\alpha}), \quad (6.11)$$

where the coefficient matrix Φ is relatively sparse, as for Ξ in the SINDy optimization problem. For rows corresponding to active degrees of freedom, Φ are unit vectors that produce an identity map (e.g. $\alpha_1 = \hat{\alpha}_1$)

In this case one motivation for sparse regression is that the library Θ tends to be fairly ill-conditioned, so that approximation with a sparse combination of polynomials may help avoid overfitting. In addition, based on the preceding discussions about DMD and triadic interactions in frequency space, it is reasonable to expect that relatively few combinations of the driving frequencies will correspond to each DMD coefficient. Once the coefficient matrix Φ is identified via a sparse regression algorithm (we use FROLS, as for the SINDy optimization), the functional relationships give a simple nonlinear dimensionality reduction; in this case the “latent variables” are the active degrees of freedom $\hat{\alpha}$ and we can approximate the full coefficient vector with the function $\alpha \approx \varphi(\hat{\alpha}) = \Phi \Theta(\hat{\alpha})$.

The key advantage to this representation of the modal coefficients is that it dramatically restricts the dimensionality of the state space. In the case of the quasiperiodic shear-

driven cavity, we reduce from the $r = 64$ -dimensional subspace spanned by the DMD modes to a 4-dimensional space of $\hat{\alpha}$. Moreover, as $\alpha_2 = \alpha_1^*$ and $\alpha_6 = \alpha_5^*$, these mode pairs represent two generalized amplitude-phase pairs, as expected for dynamics with a toroidal attractor. This eliminates the redundant variables introduced by the linear space-time separation of variables via a nonlinear manifold reduction.

The benefit of this reduced state space is immediately clear for system identification; we can apply SINDy to model the evolution of $\hat{\alpha}$ and reconstruct the full state from this minimal representation. However, the manifold representation can also be used to improve the stability and accuracy of the projection-based Galerkin models. For a general nonlinear embedding of the form $\mathbf{a} = \varphi(\hat{\mathbf{a}})$ and dynamical system $\dot{\mathbf{a}} = \mathbf{f}(\mathbf{a})$, consistency requires $\dot{\hat{\mathbf{a}}} = \mathbf{J}_\varphi(\hat{\mathbf{a}})\dot{\hat{\mathbf{a}}}$, where $\mathbf{J}_\varphi(\hat{\mathbf{a}})$ is the Jacobian of φ evaluated at $\hat{\mathbf{a}}$. Equivalently,

$$\frac{d\hat{\mathbf{a}}}{dt} = \mathbf{J}_\varphi^+(\hat{\mathbf{a}})\dot{\hat{\mathbf{a}}} = \mathbf{J}_\varphi^+(\hat{\mathbf{a}})\mathbf{f}(\varphi(\hat{\mathbf{a}})), \quad (6.12)$$

where \mathbf{J}_φ^+ is the pseudoinverse of \mathbf{J}_φ . This condition defines the reduced dynamics by constraining the velocity of \mathbf{a} to the tangent space of the manifold defined by φ (Guckenheimer & Holmes, 1983; Lee & Carlberg, 2020).

In the case of the linear-quadratic Galerkin dynamics (4.7) with the sparse polynomial manifold Eq. (6.11), the Jacobian consists of rows of the identity matrix by virtue of the fact that the reduced states $\hat{\alpha}$ are also contained in the full coefficient vector α . Then the *manifold Galerkin* dynamics are

$$\dot{\alpha}_k = \tilde{C}_k + \sum_{\ell=1}^r \tilde{L}_{k\ell} \varphi_\ell(\hat{\alpha}) + \sum_{\ell=1}^r \sum_{m=1}^r \tilde{Q}_{k\ell m} \varphi_\ell(\hat{\alpha}) \varphi_m(\hat{\alpha}), \quad k = 1, 2, 5, 6, \quad (6.13)$$

where the tilde denotes the original POD-Galerkin operators rotated to the DMD coordinates (see Chapters 3 and 4). Note that if φ contains polynomials up to order d , the quadratic interactions in (6.13) lead to an effective nonlinearity of order $2d$. This can be viewed as a generalization of Stuart-Landau-type mean field models (Noack *et al.*, 2003) or center manifold expansions (Guckenheimer & Holmes, 1983; Carini *et al.*, 2015), though in these cases the manifold equation φ is usually represented as a Taylor series truncated at second-order. This is sufficient near bifurcations, but the more general form enabled by sparse regression allows improved resolution of the manifold structure.

Despite the near-perfect kinematic resolution of the flow field in the POD basis, there is no level of truncation, up to at least $r = 64$, that leads to a Galerkin system that is both stable and reproduces the quasiperiodic dynamics of the flow, as illustrated in Fig. 6.9. In this section we show that polynomial nonlinear correlations can be used to construct a 4-dimensional model that captures the major structure of the post-transient flow. Based on these results, we argue that accurate kinematic resolution of the advection-dominated flow in the modal basis creates spurious dynamic variables and fragility in the Galerkin systems, as for the linear advection example in Sec. 6.1.

6.5 Polynomial manifold approximation

Nonlinear correlation analysis As discussed in the DMD analysis of Sec. 6.3, several of the modal coefficients are nearly pure harmonics of one of the two dominant frequencies, as a result of the the space-time separation of variables of traveling wave-type structure in the physical field. This signature is clear in the phase portraits in Figs. 6.7 and 6.11, where some coefficient pairs form Lissajous orbits, which are characteristic of harmonic oscillators with frequencies at an integer ratio. This is reflected in the relatively large scores of the RDC metric of dependence between harmonic mode pairs (Fig. 6.11, upper triangular). This measure also confirms that the shear layer dynamics and inner cavity motions are nearly independent.

However, the modal analysis in Sec. 6.3 and the power spectrum in Fig. 6.4 both indicate that the flow cannot be described by purely independent oscillation. Instead, the flow behaves more like the model in Eq. (6.9), where a linear-quadratic system is driven by self-sustaining oscillators at incommensurate frequencies, with higher-order modes connected via cascading nonlinear interactions. If this is indeed the case, the resulting triadic structure would lead to energy content at frequencies that are not pure harmonics. In other words, coefficients that are not significantly correlated with the driving oscillators according to the RDC may instead have multivariate nonlinear correlation, or frequency crosstalk.

Based on this intuition, we apply the sparse manifold regression approach described in Sec. 6.4. By applying FROLS with a residual tolerance of 0.1 we find that DMD coefficients up to α_{52} can be reconstructed with at least 90% accuracy in a library of polynomials in $\hat{\alpha}$ up to seventh order. Of these, 24 coefficients can be approximated with residual $10^{-2} - 10^{-3}$ with only one polynomial function of the active variables, indicating that the mode approximately is a product of a single triadic interaction. None of the coefficients require more than five terms (out of a library of 330).

This analysis also reveals that the higher-order coefficients have three distinct relationships to the active degrees of freedom, as illustrated in Fig. 6.12 and Table 6.2:

1. Pure harmonics

$$\alpha_j \propto \alpha_k^d \alpha_k^{*d'}, \quad k \in (1, 5). \quad (6.14)$$

In this case the dominant frequency will be $\omega_j \sim (d - d')\omega_k$. These coefficients have a high RDC score and have Lissajous-type phase portraits.

2. Nonlinear crosstalk

$$\alpha_j \propto \alpha_1^c \alpha_1^{*c'} \alpha_5^d \alpha_5^{*d'} \quad (6.15)$$

with dominant frequency $\omega_j \sim (c - c')\omega_s + (d - d')\omega_c$. These coefficients have multivariate nonlinear correlation with the active degrees of freedom, so may not have high RDC score. Two-dimensional phase portraits will also not appear meaningful. Still, the coefficients have energy content at a single frequency.

Coefficient	Active terms	Triadic frequencies	DMD eigenvalue	RDC vs. α_1
α_3	$\alpha_1^{*2} \alpha_1 ^2$	$2\omega_s$	$-0.000 + 23.4i$	0.86
α_{17}	$\alpha_1^3\alpha_5$	$3\omega_s + \omega_c$	$-0.007 + 37.9i$	0.08
α_{27}	$\alpha_1\alpha_5^2, \alpha_1\alpha_5^3$	$\omega_s + 2\omega_c, \omega_s + 3\omega_c$	$-0.020 + 17.5i$	0.26

Table 6.2: Representative nonlinear correlations identified by sparse regression. Correlation types include pure harmonic (α_3), nonlinear crosstalk (α_{17}), and mixed frequency content (α_{27}). Polynomial combinations give rise to oscillations at frequencies in terms of the shear layer $\omega_s \approx 11.7$ and inner cavity $\omega_c \approx 2.7$. For modes with nearly pure frequency content (e.g. α_3, α_{17}), the resulting frequencies are close to those predicted by the DMD analysis. The randomized dependence coefficient (RDC) between the coefficient and α_1 is strongest for pure harmonics (α_3), even though mixed-frequency modes can be accurately approximated with a simple polynomial function. See Fig. 6.11 for pairwise RDC values for the leading 24 DMD coefficients.

- Mixed frequency content: these coefficients cannot be expressed as a single polynomial term in $\hat{\alpha}$, but require 2-5 terms for a reasonably accurate approximation. The coefficients may still be an algebraic function of the active variables (i.e. a sum of terms like (6.14) and (6.15)), but will have energy content at various frequencies.

As shown by Fig. 6.12, the polynomial approximations tend to be more accurate for coefficients with pure frequency content, although they do capture the dominant trends for coefficients with mixed content. These sparse polynomial representations of higher-order coefficients determine the manifold equation $\alpha = \varphi(\hat{\alpha})$ based on Eq. (6.11).

Manifold Galerkin model This manifold restriction leads to effective higher-order nonlinearity in the reduced dynamics for $\hat{\alpha}$, given by Eq. (6.13). In particular, since we include up to 7th-order polynomials in the manifold equation, the effective dynamics based on the quadratic Galerkin model involves 14th-order terms. Fortunately, provided the coefficient matrix Φ is sufficiently sparse, the overall cost of evaluating the reduced-order model still only scales with r^3 (from $\mathcal{O}(r)$ evaluations of the quadratic term). The advantage of this additional nonlinearity is that the system is now constrained to the manifold determined by $\varphi(\hat{\alpha})$. This mitigates both the issue of spurious degrees of freedom in the Galerkin representation of hyperbolic dynamics and the effect of truncating the dissipative scales of the energy cascade.

Simulation results for the manifold Galerkin model are shown in Fig. 6.13 along with the SINDy model discussed below. Whereas the standard Galerkin model eventually overestimates the fluctuation energy and becomes aperiodic, the manifold restriction applied to the same operators continues at the correct energy level and with approximately discrete peaks in the frequency spectrum at the correct locations. Of course, there is some phase drift for all models at long times, but the manifold reduction prevents the higher-order coefficients from losing coherence with the dominant oscillations and causing the amplitude drift as in the standard Galerkin model.




Figure 6.12: Example coefficient reconstructions $\alpha \approx \varphi(\hat{\alpha})$ based on the leading DMD coefficients (—). The sparse polynomial approximation (- -) for higher-order modes with pure frequency content (e.g. $\alpha_{17} \approx \varphi_{17}(\hat{\alpha})$) tends to be more accurate than for modes with mixed content (α_{27}).




Figure 6.13: Evolution of the fluctuation kinetic energy for the reduced-order models compared to DNS. By accounting for nonlinear correlations, both the manifold Galerkin and SINDy models remain at the correct energy level at long times, despite having many fewer degrees of freedom than the standard Galerkin model (top). Similarly, both models resolve the nonlinear interactions leading to the discrete peaks in the power spectrum (bottom).

SINDy model The manifold restriction applied to the Galerkin model results in a significant reduction in dimensionality and improvement in stability and accuracy. However, the full evolution equation (6.13) is dense with $\mathcal{O}(r^3)$ entries, where r is the size of the POD/DMD subspace, not the number of active degrees of freedom. This is still a significant improvement over both the DNS and the standard Galerkin model, but the physical picture of coupled nonlinear oscillators giving rise to the quasiperiodic dynamics suggests that a simpler model may capture the dominant features of the flow.

This desire for minimalistic models has been the motivation for several recent applications of SINDy and related system identification techniques to model reduction. However, when these methods are applied to modal coefficients, they also face the fundamental representation issue challenging Galerkin models. That is, models with full kinematic resolution will include spurious dynamical degrees of freedom. The issue is exacerbated in data-driven methods, since both the dimension and the conditioning of the library matrices tend to scale poorly with dimensionality.

For example, it is well known that the flow past a cylinder at Reynolds number 100 can be accurately described by a Stuart-Landau equation with two degrees of freedom. However, fully reconstructing the post-transient vortex street requires on the order of ten POD modes, all of which are harmonics of the leading pair. Loiseau *et al.* (2018a) addressed this by identifying the Stuart-Landau equation with SINDy along with a similar sparse regression approach to Eq. (6.11) to algebraically reconstruct the harmonics.

Here we take a similar approach and assume that we do not need the full order- r dynamics and manifold equation to describe the dynamics of the active degrees of freedom $\hat{\alpha}$. In particular, we anticipate that the minimal description will take the form of coupled Stuart-Landau equations. As described in Sec. 6.4, we construct a library of candidate polynomials including up to cubic terms in $\alpha_1, \alpha_1^*, \alpha_5$, and α_5^* .

We identify symbolic equations for α_1 and α_5 dynamics with the FROLS algorithm, for DMD coefficients $\alpha_2 = \alpha_1^*$ and $\alpha_6 = \alpha_5^*$. FROLS is an iterative, forward greedy algorithm that requires a stopping condition. Often a residual-error criterion is used, as for the sparse polynomial manifold model, but in this case the DMD modes are so close to pure linear oscillation that a single linear term leaves a residual $\sim 10^{-6}$. Instead, we stop the iteration at the second term in each equation, retaining a stabilizing cubic term. The resulting model takes the form

$$\dot{\alpha}_1 = \lambda_1 \alpha_1 - \mu_1 \alpha_1 |\alpha_1|^2 \quad (6.16a)$$

$$\dot{\alpha}_5 = \lambda_5 \alpha_5 - \mu_5 \alpha_5 |\alpha_5|^2, \quad (6.16b)$$

where all λ and μ coefficients are complex.²

The system in Eq. (6.16) is a pair of independent nonlinear Stuart-Landau oscillators. Recalling the toy model in Sec. 6.4, independent oscillators that drive a cascade of triadic

²See <https://github.com/dynamicslab/pysindy/examples> for a simplified presentation of these results using the `pysindy` library

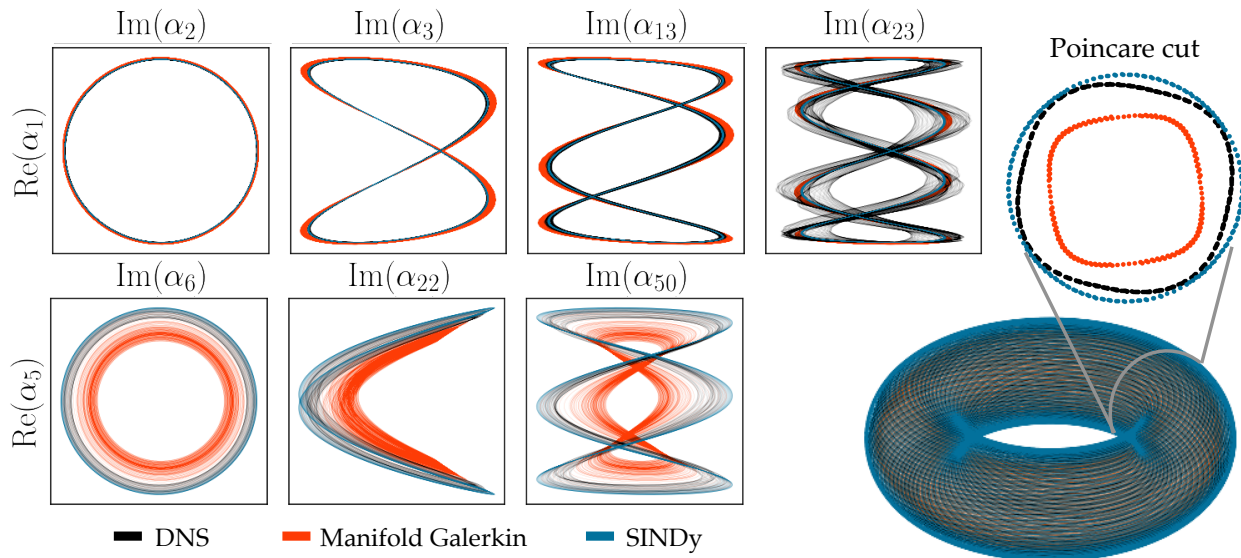


Figure 6.14: Phase portraits for the POD coefficients reconstructed from the manifold Galerkin and SINDy reduced-order models (left). Both models accurately capture the shear layer instability and its harmonics (e.g. $\alpha_1, \alpha_3, \alpha_{13}, \alpha_{23}$), though the manifold Galerkin model tends to underestimate the amplitude of the inner cavity motions (e.g. $\alpha_5, \alpha_{22}, \alpha_{50}$). A Poincarè section of the toroidal attractor confirms this discrepancy, but shows clearly that both models are quasiperiodic and remain on the approximate attractor.

interactions can lead to quasiperiodic dynamics, even without direct dynamical coupling between the oscillators. In contrast to the manifold Galerkin model, it is not necessary to reconstruct the full vector of coefficients to solve this minimal system. Of course, the full vector of coefficients can still be reconstructed after simulating Eq. (6.16) via the manifold function φ .

Figure 6.13 compares the fluctuation kinetic energy based on reconstructions from the SINDy model to the standard and manifold Galerkin models. As for manifold Galerkin, the SINDy model remains at the correct energy level at long times and reproduces the characteristic structure of the power spectrum. A slightly more sensitive evaluation is given in Fig. 6.14, which compares Lissajous figures of reconstructed near-harmonic POD modes. Both models accurately capture the harmonics associated with the shear layer instability, but the Galerkin system somewhat underestimates the amplitude of the inner cavity mode and its harmonics.

Although the reduced state space of the pair of complex coefficients is four-dimensional, the quasiperiodic oscillatory nature of the dynamics also offers a convenient symmetry reduction for the purposes of visualization. With the amplitude-phase representation $\alpha_k = R_k e^{i\phi_k}$, we can approximate the toroidal attractor in the three-dimensional space by

representing α_5 as an expansion about the point in the complex plane defined by α_1 :

$$x = (R_1 + R_2 \cos \phi_2) \cos \phi_1 \quad (6.17a)$$

$$y = (R_1 + R_2 \cos \phi_2) \sin \phi_1 \quad (6.17b)$$

$$z = R_2 \sin \phi_2. \quad (6.17c)$$

Finally, the models can be compared in detail with a Poincarè section of this torus about any convenient plane (we choose $x = 0$); both the three-dimensional phase portrait and Poincarè section are also shown in Fig. 6.14.

As a result of the slight underestimation of energy in the inner cavity motions, the Poincarè section for the manifold Galerkin model shows a somewhat smaller attractor than the DNS and SINDy. In contrast, SINDy slightly overestimates the energy of the inner cavity oscillation, leading to an attractor section with somewhat larger radius. The highly simplified structure of the SINDy model also leads to a circular section, while the Galerkin system captures the rounded-square shape of the true section. This is likely a consequence of the high-order effective nonlinearity in the Galerkin system, which allows it to resolve more complex attractor shapes. Nevertheless, the SINDy system does give an accurate estimate of the typical amplitude in the slice and preserves the coherence of the harmonic modes for both the shear layer and inner cavity oscillations.

In both the manifold Galerkin and SINDy models, the nonlinear correlations play a critical role in the accuracy and stability of the reduced-order dynamics. The space-time separation of variables applied to an advection-dominated flow introduces a significant number of modes that are necessary to reconstruct the field, but do not correspond to independent degrees of freedom in the dynamics. Nonlinear correlations analysis provides a straightforward, principled approach to restricting the Galerkin dynamics to the set of active degrees of freedom, as well as convenient coordinates for system identification.

6.6 Deep manifold models and the route to chaos

As demonstrated in the previous section, the manifold reduction strategy based on nonlinear correlations and sparse polynomial regression is an accurate and efficient approach to deriving compact nonlinear dynamical systems models of flows with approximately discrete frequency content. In the language of Koopman theory, such a flow has a discrete or point spectrum (Mezić, 2013; Arbabi & Mezić, 2017); more sophisticated analysis would be necessary to extend these results to chaotic or turbulent systems with continuous spectra. The reason for the qualifier on frequency content is that the polynomial manifold approximation relies implicitly on the assumption that the higher-order modes are generated by triadic interactions in the frequency domain, as in the model cascade introduced in Sec. 6.4. This includes periodic and quasiperiodic dynamics, for instance, but excludes broadband chaos.

However, in principle any invertible nonlinear dimensionality reduction method could

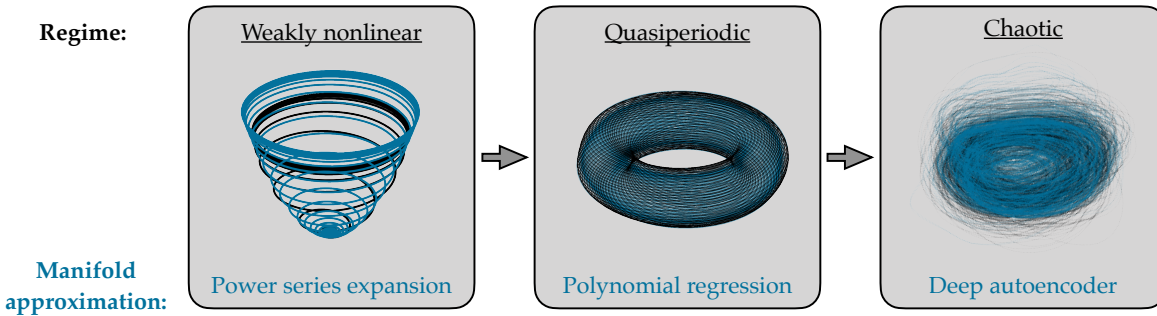


Figure 6.15: **Approximation strategies of the slow manifold.** For small-amplitude perturbations relatively close to a bifurcation, the classical center manifold approximation is a good choice (left, Chapter 5). For phase-locked or quasiperiodic dynamics, a SINDy-type sparse polynomial approximation to the manifold equation is a straightforward, interpretable alternative (middle, Sec. 6.5). When the flow exhibits broadband chaos, these methods will fail and a deep autoencoder offers a more flexible class of function approximation (right, Sec. 6.6).

be used to the same end. For instance, deep autoencoder networks (Bengio *et al.*, 2013) are an unsupervised method that learns a compressed representation of high-dimensional data. Autoencoders have recently been explored for black-box forecasting (Vlachas *et al.*, 2018), system identification (Champion *et al.*, 2019a), and model reduction (Lee & Carlberg, 2020). This section extends the empirical invariant manifold reduction to the chaotic regime of shear-driven cavity flow with a deep autoencoder. See Fig. 6.15 for a schematic of these increasingly general approximations of the slow manifold.

The autoencoder reduction The deep manifold has a similar structure to the encoding diagram shown in Fig. 6.1, but replaces the polynomial-based nonlinear correlations with an autoencoder. As with the polynomial approximation, the starting point for this method is the state vector of modal coefficients. In this case we use the POD coefficients $\mathbf{a}(t) \in \mathbb{R}^r$ rather than the DMD coefficients, since the autoencoder will not rely on approximately pure frequency content. The autoencoder learns a low-dimensional embedding $\hat{\mathbf{a}} \in \mathbb{R}^{r_0}$, with $r_0 < r$ for the coefficient vector using two standard multilayer perceptrons (feedforward neural networks). The first, which we denote $\hat{\mathbf{a}} = \varphi^{-1}(\mathbf{a})$, encodes the state, while the second approximates the manifold equation $\mathbf{a} \approx \varphi(\hat{\mathbf{a}})$. Note that the encoder is not necessarily an exact inverse of the decoder, but it is functionally equivalent.

By learning an approximate identity map $\mathbf{a} \approx \varphi^{-1}(\varphi(\mathbf{a}))$, the autoencoder performs an unsupervised nonlinear dimensionality reduction. Since the embedding is purely kinematic, a “dynamics regularization” can also be applied, which encourages the network to learn an approximation that is consistent with the original dynamics $\dot{\mathbf{a}} = \mathbf{f}(\mathbf{a})$. In partic-

ular, the loss function for the autoencoder is

$$V(\mathbf{a}) = \|\mathbf{a} - \varphi^{-1}\varphi(\mathbf{a})\|^2 + \eta\|\mathbf{f}(\mathbf{a}) - \mathbf{f}(\varphi^{-1}\varphi(\mathbf{a}))\|^2, \quad (6.18)$$

where η is a small weight (hyperparameter). The network is trained with batch stochastic gradient descent as usual: this loss function is evaluated across a batch of training samples and the neural network weights that parameterize φ and φ^{-1} are updated.

The cavity flow models in this section begin from a POD-Galerkin model \mathbf{f} with $r = 16$ and use a latent dimension of $r_0 = 4$. The network is constructed in the `jax` ecosystem (Bradbury *et al.*, 2018; Hennigan *et al.*, 2020; Hessel *et al.*, 2020). The output sizes of the encoder are (16, 32, 64, 64, 32, 16, 8, 4), with the inverse for the decoder, and the dynamics regularization is $\eta = 1e - 3$. These hyperparameters were not chosen by an exhaustive search, but give the network enough flexibility for an accurate reconstruction while still being efficient to evaluate during the online model simulation. The latent space dimension $r_0 = 4$ was simply chosen because that was the dimension of the active POD subspace for the quasiperiodic model in Sec. 6.5. The autoencoder parameters are trained with an ADAM optimizer using learning rate 10^{-4} and batch size 1000. The POD coefficients are divided with a 90% train/test split and the training proceeds until the training and test losses both stop decreasing.

Of course, the true validation of the deep manifold model is an evaluation of its dynamical behavior. Once the neural network is trained, the model can be simulated in the latent space. Chapter 2 derived the consistency condition (6.13) for an arbitrary manifold reduction $\mathbf{a} \approx \varphi(\hat{\mathbf{a}})$ applied to the dynamics $\dot{\mathbf{a}} = \mathbf{f}(\mathbf{a})$, where $\mathbf{a} \in \mathbb{R}^r$ and $\hat{\mathbf{a}} \in \mathbb{R}_0^r$ with $r_0 < r$. The usual dynamical systems view of this condition as a system of nonlinear PDEs for the equation $\varphi(\hat{\mathbf{a}})$, as expressed by Eq. (2.25). This is the starting point of the classical power series approximation of the center manifold (see Secs. 2.3 and 5.4).

However, it can also be viewed as the linear system of equations (2.24) restricting the dynamics to the tangent space of the manifold (Lee & Carlberg, 2020). In the present case this point of view is more convenient because the manifold equation has already been learned by the training procedure. The latent state $\hat{\mathbf{a}}$ is advanced using a standard explicit Dormand-Prince scheme in `scipy`. At each time step the Jacobian $\mathbf{J}_\varphi(\hat{\mathbf{a}}) \in \mathbb{R}^{r \times r_0}$ is evaluated using the automatic differentiation capabilities in `jax` and the linear system is solved as in (6.13). This approach combines the physics-based Galerkin system with the deep manifold reduction for a compact, expressive low-dimensional model.

Modeling the route to chaos The shear-driven cavity flow described in Sec. 6.2 undergoes a supercritical Hopf bifurcation at $Re_c^1 \approx 4126$. Shortly thereafter it passes through a complex series of bifurcations between $4300 < Re < 4600$ before a single stable limit cycle finally emerges. See Sipp & Lebedev (2007); Meliga (2017); Bengana *et al.* (2019) for detailed analyses of this regime. By $Re = 7500$ this limit cycle is no longer stable and the post-transient flow is quasiperiodic, as discussed in Sec. 6.2 and by Sipp *et al.* (2010).

Upon further increase of the Reynolds number, the quasiperiodic 2-torus itself loses stability and the flow breaks down into chaos.

Although this is a fairly coarse description and much remains to be done in characterizing the sequence of bifurcations in this flow, the transitions to periodic, quasiperiodic, and finally chaotic flow appear to follow the canonical Ruelle-Takens route to chaos. The purpose of this section is not a careful study of the sequence of bifurcations in shear-driven cavity flow, but rather (i) to demonstrate that a deep manifold model can accurately model chaotic flow, and (ii) to show that this model can itself reproduce the Ruelle-Takens scenario.

Reduced-order modeling generally becomes much more difficult when parametric dependence is introduced, even in the linear case (Benner *et al.*, 2015). The simplest, and least efficient, approach is to recompute the entire POD-Galerkin model beginning with the high-fidelity simulation at each parameter value (Reynolds number, in this case). Another straightforward method is to perform simulations at a small number of representative parameter values and compute a “global POD” basis from snapshots at different Reynolds numbers. In the latter case, the linear-quadratic reduced-order model (4.7) can then be computed offline and simply evaluated at any Reynolds number, since the parametric dependence is affine in Re^{-1} . Although computationally efficient, this method cannot capture parametric deformation of the modes themselves, which can be significant across large changes in Reynolds number. More sophisticated methods such as Grassmann interpolation (Amsallem & Farhat, 2011; Benner *et al.*, 2015; Loiseau *et al.*, 2018a) could be used to this end, although this is beyond the scope of the present work.

To explore the transition to chaos, we repeat the DNS described in Sec. 6.2, changing only the Reynolds number to 10^4 , at which point the flow is chaotic. We compute both a new “local” POD basis at this Reynolds number and a “global” basis combining the same snapshots from the DNS at $Re = 7500$ and $Re = 10^4$. Figure 6.16 compares simulations of deep manifold models trained at each Reynolds number to the statistics of the true POD coefficients. Although each model misses some detailed features both in the latent space, both the phase portraits and power spectra show close agreement, demonstrating that the deep manifold model can generalize beyond the limitations of polynomial regression to chaotic flows.

Finally, we construct a global POD-Galerkin model by computing separately elements of the constant and linear Galerkin tensors in Eq. 4.8 that vary with Reynolds number (the viscous terms) and those that do not (advective terms). The system may then be immediately assembled and evaluated at any Reynolds number. In addition, we train a new deep manifold model using coefficients from the projected DNS fields at both parameter values without dynamics regularization.

Figure 6.17 shows a bifurcation diagram for this global model. This diagram is constructed by first decoding the 16-dimensional vector of estimated POD coefficients and then taking an approximate Poincaré section, interpolating the coefficients to each signed zero crossing of a_1 as in Fig. 6.14. The coordinates of the estimated POD coefficients cor-




Figure 6.16: Deep manifold models of shear-driven cavity flow. The flow at $Re = 7500$ is quasiperiodic (top), with approximately discrete peaks in the power spectrum, but has transitioned to chaos by $Re = 10^4$ (bottom). Both the phase portraits in the latent space (left) and the power spectra for the reconstructed POD coefficients (right) show qualitative agreement. The quasiperiodic power spectrum is estimated with the FFT, while the chaotic case is estimated with Welch’s method.

responding to the inner cavity oscillations in this global basis (a_4, a_5) are then plotted at each Reynolds number. Isolated points therefore indicate the flow is on a limit cycle, while loops indicate a section of a 2-torus and a cloud of points suggests chaos.

As illustrated by this figure, the deep manifold model does follow the Ruelle-Takens route to chaos parameterized by Reynolds number in this flow. The correspondence with the actual flow is largely qualitative, since the critical Reynolds numbers are not accurately predicted. For instance, the autoencoder model does not become chaotic until $Re \simeq 11000$, while the DNS is chaotic by $Re = 10000$. In addition, the model does not accurately predict the initial Hopf bifurcation, although this might be expected from the relatively high Reynolds numbers at which the POD modes and Galerkin system were computed.

This suggests several possibilities for refinement of this approach. For example, the global POD basis is a very coarse approximation; there is significant mode deformation between the steady flow at $Re \simeq 4000$ and the chaos at $Re \simeq 10^4$ that certainly modifies the dynamics of the Galerkin system. An interpolation approach could significantly improve the fidelity of the Galerkin model. Second, the hyperparameters used in this demonstration are not expected to be optimal; the behavior of the deep manifold model could be improved with a proper choice of dynamics regularization, for instance. The accuracy of the simple SINDy model in Sec.6.5 also raises the possibility of combining the autoen-

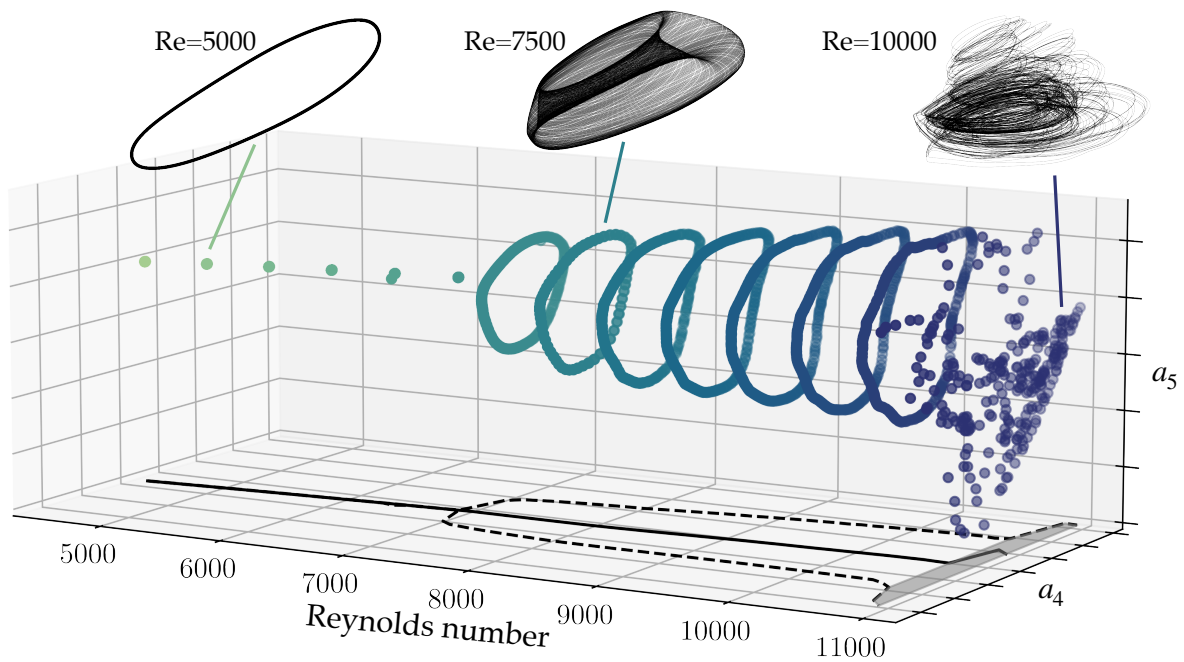


Figure 6.17: **Bifurcation diagram for the parametric deep manifold model.** The model was trained on data from $Re = 7500$ and 10000 and evaluated on $4500 \leq Re \leq 11000$. The points shown here are a Poincaré section of the reconstructed coefficients corresponding to inner cavity oscillations (see text for details). The appearance of a loop near $Re \simeq 7000$ indicates the onset of quasiperiodicity (dashed lines on $x - y$ plane), while the breakdown into a cloud of points shows chaos (shaded section on $x - y$ plane). Selected trajectories in the latent space show a limit cycle, 2-torus, and chaotic motion.

coder framework with SINDy model discovery, as proposed by Champion *et al.* (2019a). Finally, the shear-driven cavity flow at these Reynolds numbers is purely academic, since three-dimensional instabilities would almost surely have appeared by $Re = 10^4$. Future studies of the route to turbulence would benefit from three-dimensional simulations that might give more insight into realistic flow configurations.

Despite the present shortcomings, these results show that a physics-constrained deep learning model can reproduce the Ruelle-Takens route to chaos with a four-dimensional reduced-order model. More broadly, this confirms that even chaotic fluid flow at high Reynolds numbers can have strongly low-dimensional dynamics, provided that a flexible enough dimensionality reduction technique is used. As in the case of polynomial correlations, we have shown here that linear modal bases such as POD and DMD are not necessarily the most compact coordinate systems for low-dimensional modeling. Regardless of the method, nonlinear embedding recognizes the intrinsic dimensionality of the dynamics as distinct from that of the linear subspace required to reconstruct the flow field.

6.7 Discussion

It has been widely recognized for some time that Galerkin-type models of advection-dominated flows are prone to fragility and instability. The majority of work addressing this issue has focused on truncation of the energy cascade, leading to closures in the vein of subgrid-scale large eddy simulation models (Rempfer & Fasel, 1994; Wang *et al.*, 2012; Cordier *et al.*, 2013; Östh *et al.*, 2014; Pan & Duraisamy, 2018).. However, recent work including Carlberg *et al.* (2017); Grimberg *et al.* (2020); Lee & Carlberg (2020) has begun questioning the fundamental suitability of Galerkin projection for hyperbolic problems, pointing out for instance that any notion of optimality associated with the Galerkin system is lost upon time discretization. This perspective is supported by the observation that Galerkin-type reduced-order models often do not significantly improve with increasing rank, as one might expect if the primary issue was under-resolved dissipation, even when the dissipation rate is fully resolved. On the other hand, when the modal basis is selected carefully and the flow is not advection-dominated, heavily truncated Galerkin systems have been shown to accurately reproduce key features of turbulent shear flows (Moehlis *et al.*, 2004; Grimberg *et al.*, 2020; Cavalieri, 2021). Again, this suggests that resolving energy transfer mechanisms is central to constructing accurate reduced-order models.

In this work we have used a nonlinear correlations analysis of a quasiperiodic shear-driven cavity flow to argue that decoherence resulting from the linear modal representation of advecting structures also deserves consideration. This decomposition introduces one temporal coefficient per spatial mode; in many cases this may result in many more coefficients than there are degrees of freedom in the post-transient flow. Galerkin models treat each coefficient as an independent degree of freedom; small errors in the system of differential equations can lead to catastrophic decoherence and instability. Instead, we

show that exploiting statistical structure and algebraic dependence in the temporal coefficients enables the reduction of the dynamical system to the true rank while preserving the kinematic resolution of the modal basis.

The cavity flow is dominated by two key modal structures: a high-frequency shear layer instability and low-frequency inner cavity oscillation. The natural dynamics of the flow are quasiperiodic, as can be seen from the characteristic power spectrum in Figure 6.4. Both the shear layer and inner cavity features can be identified by a stability analysis of the time-averaged mean flow. However, linear modal representations (POD or DMD) approximate the traveling wave structures in the nonlinear flow with not only the fundamental stability modes, but also higher harmonics and nonlinear crosstalk modes, each of which can be associated with one of the approximately discrete peaks in the power spectrum.

The physical coherence (non-dispersion) of the fundamental flow features appears as nonlinear correlation between temporal coefficients associated with harmonics and crosstalk. This can also be conceptualized as triadic interactions in the frequency domain. After using a randomized dependence coefficient (RDC) analysis to identify the dynamically active modes, we use sparse polynomial regression to uncover simple algebraic relationships that account for $\sim 99.5\%$ of the fluctuation kinetic energy.

We give examples of two ways in which these relationships can be used to improve reduced-order models. First, they act as a simple manifold equation constraining the Galerkin dynamics to the post-transient attractor; this may be viewed as a data-driven generalization of analytic invariant manifold reductions (e.g. Noack *et al.* (2003)), which rely on scale-separation arguments. Alternatively, the driving coefficients offer a convenient basis for nonlinear system identification; we use the SINDy framework (Brunton *et al.*, 2016a) to identify a simple model of the flow as a pair of independent Stuart-Landau oscillators. The full flow field may then be reconstructed with the manifold equation.

The manifold Galerkin system connects the reduced-order system to the governing equations and may allow for more natural parametric variation, but requires accurate estimates of the gradients of the POD modes and/or intrusive access to the full-order solver. On the other hand, the SINDy model is compact, non-intrusive, and more amenable to analytical treatment, though it cannot be directly connected to the underlying physical equations and it is more difficult to capture parametric variation. In general, the most appropriate choice is likely to be application-dependent. Regardless of the chosen model reduction technique, we conclude that exploiting nonlinear structure in the modal coefficients is a natural and efficient approach to improving the stability and accuracy of low-order models of this flow.

In a broader sense, approaching this analysis (with the RDC and sparse polynomial regression) can be seen as simple, interpretable manifold learning. This is sufficient for quasiperiodic dynamics, since the form of the nonlinear dependence can be readily deduced by reasoning about the triadic interactions. For more complex dynamics, the polynomial approximation can be extended with a deep manifold model based on an autoen-

coder embedding. This generalization can reproduce the dynamics in the chaotic regime of the shear-driven cavity flow and capture the sequence of bifurcations characteristic of the Ruelle-Takens route to chaos.

There are also many opportunities for further work in reduced-order model development. For instance, accounting for nonlinear correlations does not address the issue of truncating the energy cascade. This does not pose a problem for the present laminar, two-dimensional flow, but severe dimensionality reduction of any multiscale or turbulent dynamics will necessarily act as a spatiotemporal filter. In other words, the dissipative scales will generally not be correlated in any way with the large-scale dynamics and a closure strategy is likely to be necessary in order to accurately capture the dissipation rate. This is typically less of an issue in the system identification framework, since the natural dynamics are estimated at once, including any effective closure models within the span of the candidate functions. However, for more complex flows it may be necessary to employ a more sophisticated optimization (Champion *et al.*, 2019b), physics-based constraints (Loiseau *et al.*, 2018b), or enforcement of long-term stability (Schlegel & Noack, 2015; Kaptanoglu *et al.*, 2021a). Despite prospective challenges in scaling this approach to chaotic and turbulent flows, we expect that there are significant stability and robustness benefits to be realized by exploiting nonlinear correlations in reduced-order models of coherent structures in advection-dominated flows.

Chapter 7

MULTISCALE CLOSURE MODELS VIA STOCHASTIC AVERAGING

Despite the complex and often chaotic dynamics exhibited by many unsteady fluid flows, they are often dominated by energetic coherent structures evolving on relatively long length and time scales (Holmes *et al.*, 1996). This realization made it possible to study fluid flows as high-dimensional dynamical systems evolving on a low-dimensional manifold, providing answers to longstanding questions on topics such as the route to turbulence (Landau, 1944; Hopf, 1948; Ruelle & Takens, 1971; Swinney & Gollub, 1981) and the role of nonlinear interactions (Stuart, 1958; Landau & Lifshitz, 1959). The persistent nature of these coherent structures eventually also raised the possibility of using low-dimensional surrogate models for optimization and control objectives (Noack *et al.*, 2011; Brunton & Noack, 2015a; Rowley & Dawson, 2017).

The fundamental challenges in constructing such reduced-order models may be broken into two categories: approximating the structures themselves, and approximating their evolution. We refer to these as the kinematic and dynamic approximations, respectively. A modal separation of variables assumption is often employed for the former task (Taira *et al.*, 2017, 2020), but although this may be suitable for many closed or diffusion-dominated flows, it is not a natural representation of traveling waves or advection-dominated flows (Rowley & Marsden, 2000; Reiss *et al.*, 2018; Rim *et al.*, 2018; Grimberg *et al.*, 2020; Mendible *et al.*, 2020). This issue is inextricably linked to the problem of modeling the coherent structure dynamics; as is well known in many domains, a proper choice of coordinates can greatly simplify the modeling task (Champion *et al.*, 2019a).

The multiscale nature of fluid flows further complicates both the kinematic and dynamic aspects of low-dimensional modeling. For instance, the effective dimensionality of a chaotic or turbulent flow may be orders of magnitude greater than that of a laminar flow with one or two instability modes. Meanwhile, the “triadic” structure of the nonlinear interactions in the wavenumber or frequency domain ensures that the dynamics of all scales across the flow are linked. Thus, even if the large, energetic coherent structures can be approximated with a low-dimensional basis, models of their dynamics that do not account for the role played by the unresolved degrees of freedom are often unstable or physically inconsistent (Noack *et al.*, 2011; Callahan *et al.*, 2022). Similar considerations impact the development of numerical methods (Bazilevs *et al.*, 2007), self-consistent mean flow modeling (Meliga, 2017), and resolvent analysis (Padovan *et al.*, 2020; Rigas *et al.*, 2021; Barthel *et al.*, 2021, 2022).

The need for subscale modeling was apparent even in early work combining empir-

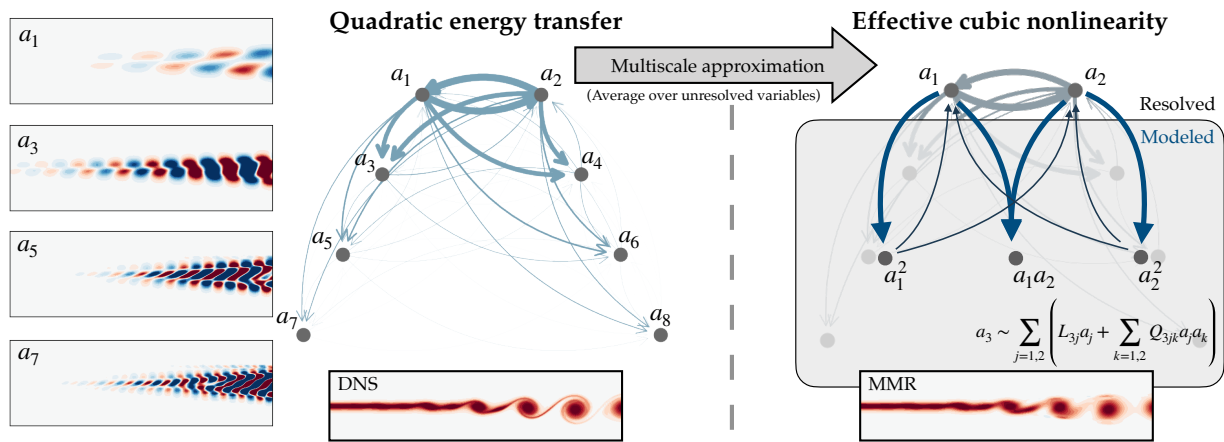


Figure 7.1: **Illustration of multiscale closure model applied to a mixing layer.** The visualization of the network of average energy transfer between the leading harmonic modes shows the cascade of energy to higher-order modes. The multiscale model reduction (MMR) method approximates the effects of unresolved higher-order modes via stochastic averaging, which leads to a generalized Stuart-Landau-type equation with cubic nonlinear interactions. Note that the while the network visualization of energy transfers (left) is computed from the modal coefficients and Galerkin model for the leading harmonics, that of of the multiscale approximation (right) is a notional illustration of the origin of the cubic terms in Eq. (7.24b). See Sec. 7.4 for details on the construction of this figure and the low-dimensional model of the mixing layer.

ical modal approximations, such as the proper orthogonal decomposition (POD), with physics-based model reduction, such as Galerkin projection. For example, low-dimensional models of vortex shedding in the globally unstable cylinder wake could accurately predict the dynamics over short times, but were subject to structural instability over a longer time horizon (Deane *et al.*, 1991; Ma & Karniadakis, 2002). Eventually Noack *et al.* (2003) showed that this was a result of the failure of the standard post-transient POD basis to resolve the Stuart-Landau mechanism of mean flow deformation associated with nonlinear interactions between the fluctuations (Landau, 1944; Stuart, 1958), an insight that enabled low-dimensional modeling of natural and actuated flows with increasingly more complex dynamics (Luchtenburg *et al.*, 2009; Deng *et al.*, 2020; Sieber *et al.*, 2020; Deng *et al.*, 2020; Callahan *et al.*, 2021c).

Contemporary work on POD-Galerkin modeling of turbulent shear flows also recognized the need for closure models that could approximate the effect of unresolved scales (Aubry *et al.*, 1988; Rempfer & Fasel, 1994; Ukeiley *et al.*, 2001). These efforts targeted two distinct physical mechanisms: mean flow deformation and subscale dissipation. For flows that are either parallel (Aubry *et al.*, 1988) or weakly nonparallel under the assumption of Taylor’s frozen turbulence hypothesis (Ukeiley *et al.*, 2001), these authors developed a Boussinesq eddy viscosity relationship between the resolved scales and a slowly varying parallel mean flow, leading to stabilizing cubic terms consistent with the Stuart-Landau description. To capture dissipation due to the unresolved scales, these models also adopted a linear mixing length approximation.

Although these studies remain landmark explorations of low-dimensional coherent structure modeling, there are several opportunities to improve the proposed closure strategies. First, it is difficult to generalize the Reynolds stress models applied to parallel shear flows by Aubry *et al.* (1988) and Ukeiley *et al.* (2001) to fully inhomogeneous flows. While the “shift mode” approach to mean flow modeling via an augmented POD basis introduced by Noack *et al.* (2003) is agnostic to the geometry of the flow, it requires computation of the unstable steady state of the Navier-Stokes equations, which is not generally experimentally accessible. Second, in the Richardson-Kolmogorov energy cascade description, energy is transferred from large to small scales through *nonlinear* interactions, where it is finally dissipated. This is not consistent with a linear mixing length model, a fact exploited by later work investigating nonlinear models of subscale dissipation (Wang *et al.*, 2012; Cordier *et al.*, 2013; Östh *et al.*, 2014), particularly the finite time thermodynamics approach, which is centered on modeling unresolved nonlinear energy transfers (Noack *et al.*, 2008).

Recent years have seen a surge in interest in data-driven and machine learning-based methods (Brenner *et al.*, 2019; Duraisamy *et al.*, 2019; Brunton *et al.*, 2020), including a number of proposed closure and stabilization schemes for reduced-order models, either through regression to additional linear-quadratic terms (Xie *et al.*, 2018; Mohebujjaman *et al.*, 2017, 2018) or by adding a deep learning model to approximate the residual (San & Maulik, 2018b,a; Menier *et al.*, 2022). Alternative work has explored interpretable system

identification methods that forego the projection-based model altogether (Brunton *et al.*, 2016a; Loiseau & Brunton, 2018; Peherstorfer & Willcox, 2016; Qian *et al.*, 2020; Callahan *et al.*, 2021c), but may incorporate physical constraints derived from the Galerkin system (Loiseau *et al.*, 2018b; Deng *et al.*, 2020; Kaptanoglu *et al.*, 2021a). If an accurate, non-intrusive model is more important than interpretability or satisfying physical constraints, then the traditional projection-based framework can be eliminated altogether with black-box neural network forecasting methods (Hesthaven & Ubbiali, 2018; Wan *et al.*, 2018).

These empirical closure models are constructed on the assumption that the influence of the unresolved variables can be approximated based on information from the resolved variables alone. The success of the Reynolds stress-mean flow models (Aubry *et al.*, 1988; Ukeiley *et al.*, 2001; Mantič-Lugo *et al.*, 2014), invariant or center manifold reductions (Guckenheimer & Holmes, 1983; Couillet & Spiegel, 1983; Noack *et al.*, 2003; Carini *et al.*, 2015), and weakly nonlinear analysis (Stuart, 1958; Sipp & Lebedev, 2007; Meliga *et al.*, 2009; Meliga & Chomaz, 2011) suggests that there may be circumstances where this relationship may be derived analytically via traditional analysis. The unifying thread between these methods is the assumption of a scale separation between the resolved and unresolved variables that can be exploited to develop an asymptotically correct closure model.

Beyond the field of fluid dynamics, the method of adiabatic elimination (Haken, 1983; Risken, 1996) has long been used to discard the fast variables in systems with emergent large-scale coherence when there is a separation in time scales, while heterogeneous multiscale methods have played an important role in simulating physical systems with widely separated scales (E & Engquist, 2003; E *et al.*, 2007; E, 2011). A similar stochastic averaging approach has been successful in climate modeling (Majda *et al.*, 2001), where the primitive equations have the same quadratic nonlinearity as the usual Navier-Stokes equations without rotation, buoyancy, topography, etc. This method, also called homogenization in the multiscale modeling literature, has its roots in the theory of singular perturbations of Markov processes (Kurtz, 1973; Papanicolaou, 1976) and is rigorously supported for stochastic systems with asymptotic scale separation; see Majda *et al.* (2001); Givon *et al.* (2004); E (2011), or Pavliotis & Stuart (2012) for in-depth presentations. With some assumptions on ergodicity, a similar approach can also be taken with deterministic systems, even in a regime where the scale separation is not in the asymptotic limit (Majda *et al.*, 2006).

This chapter explores the application of multiscale stochastic averaging methods developed by Majda *et al.* (2001), Givon *et al.* (2004); E (2011); Pradas *et al.* (2012), Pavliotis & Stuart (2012), and others to the closure problem in reduced-order models of incompressible flow or other systems reducible to linear-quadratic dynamics (Rowley *et al.*, 2004; Qian *et al.*, 2020; Kaptanoglu *et al.*, 2021b), including introducing an approximation to the form of the fast dynamics that allows for computation of the averaged dynamics in closed form. The procedure is summarized in Fig. 7.1 and shown schematically in Fig. 7.2. We refer to this secondary dimensionality reduction of the Galerkin system as multiscale model

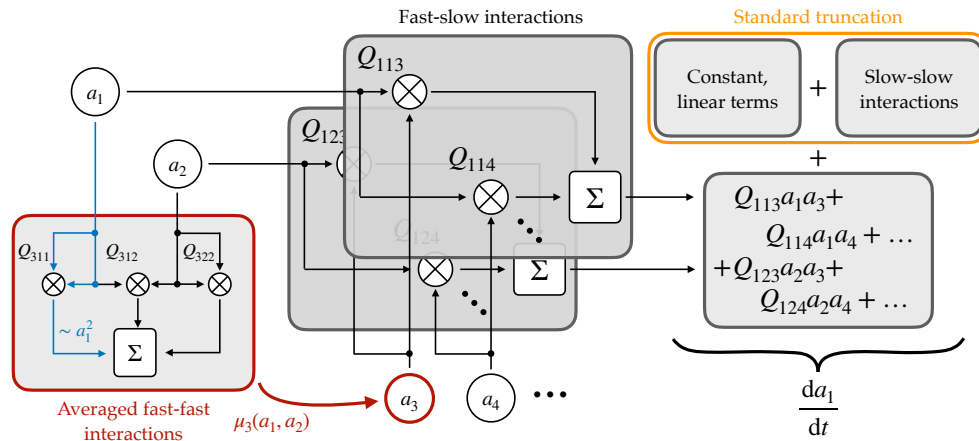


Figure 7.2: **Schematic of nonlinear interactions in the multiscale closure scheme.** The dynamics of one variable (a_1) in a system with two slow variables involve quadratic interactions between fast and slow variables that would be neglected in a standard truncation. Instead, the proposed method averages over the fast scales, ultimately generating effective cubic nonlinearities in the closed equations.

reduction (MMR).

Since fluid flows generally do not have a true scale separation away from the threshold of instability, we demonstrate via numerical simulations that this method is a robust and systematic approach to stabilizing low-dimensional models. This extends the work of Majda *et al.* (2006) exploring the application of this class of methods to systems beyond the parameter regimes where their validity can be rigorously proven. Multiscale model reduction is a unified framework for understanding the origin and importance of cubic terms in reduced-order models of the linear-quadratic Navier-Stokes equations, capable of capturing both mean flow deformation and subscale dissipation. Throughout this work we also highlight connections to other modeling methods, including Koopman theory and weakly nonlinear analysis.

7.1 Multiscale closure modeling

One of the primary difficulties of reduced-order modeling for fluid flows is that the nonlinear interactions in the Navier-Stokes equations transfer energy between all scales of the flow. This means that the restriction to a low-dimensional subspace can lead to a poor dynamic approximation, as discussed at length in Chapter 6. Generally speaking, the goal of closure modeling is to augment this truncated model with additional terms that account for the effect of the unresolved scales.

The multiscale approach to closure modeling accomplishes this by first partitioning the dynamics into fast and slow variables, and then approximating the solution to the

associated backwards Kolmogorov equation with a perturbation series expansion. The result can be interpreted as a Koopman generator of the form of Eq. (2.39) corresponding to the coarse-grained dynamics for the slow variables. As we will show, when applied to the linear-quadratic Galerkin system (4.7), this leads to a generalized Stuart-Landau equation including cubic terms.

An interesting feature of cubic Stuart-Landau-type models of fluid flow is that they are often more accurate than linear-quadratic models, even though the nonlinearity in the underlying governing equations is quadratic (Noack *et al.*, 2003; Loiseau & Brunton, 2018). One reason for this is the artificial space-time separation of variables introduced by the modal representation (3.1), which can introduce spurious degrees of freedom that represent phase-locked harmonics of traveling waves, for instance (Callahan *et al.*, 2022). As a result, POD-Galerkin models are vulnerable to decoherence and can fail to improve with increasing rank, even when the kinematic approximation of the flow field becomes nearly perfect. As illustrated by the examples in Secs. 7.3 and 7.4, the introduction of cubic terms can mitigate this, either by modifying the dynamics to resemble phase-locked nonlinear oscillators or by eliminating spurious degrees of freedom altogether.

Although the more general method of multiscale averaging can be rigorously justified when there is a strict separation of timescales (Majda *et al.*, 2001; Pavliotis & Stuart, 2012), there is no such strict separation of timescales in most fluid flows of practical interest. The proposed method for model reduction should therefore properly be viewed as an approximate closure model motivated by the asymptotic limit in the spirit of the variational multiscale approach to turbulence modeling (Bazilevs *et al.*, 2007). Nevertheless, it provides a systematic method for stabilizing low-dimensional models and also highlights a generic mechanism by which higher-order nonlinearities can arise from the quadratic term in the governing equations.

Averaging over unresolved variables Beginning with the Galerkin dynamics (4.7) for the state consisting of POD coefficients $\mathbf{a}(t) \in \mathbb{R}^r$, where we assume that r is large enough for an accurate kinematic reconstruction of a typical flow field, we partition the system into slow variables $\mathbf{x}(t) \in \mathcal{X} = \mathbb{R}^{r_0}$ and fast variables $\mathbf{y}(t) \in \mathcal{Y} = \mathbb{R}^{r-r_0}$, so that $\mathbf{a} = [\mathbf{x}^T \ \mathbf{y}^T]^T$.

For concise notation we will use the Einstein convention that repeated indices imply summation and we will omit explicit summation unless not doing so would lead to ambiguity. We will index the slow variables with Roman subscripts i, j, k, ℓ ranging from 1 to r_0 and the fast variables with Greek subscripts α, β, γ ranging from 1 to $r - r_0$. Without loss of generality we will also assume that the quadratic term has been symmetrized in the last two indices, so that $Q_{ijk} = Q_{ikj}$. Then the partitioned Galerkin system is

$$\dot{x}_i = f_i^x(\mathbf{x}, \mathbf{y}) = F_i^1 + L_{ij}^{11}x_j + L_{i\beta}^{12}y_\beta + Q_{ijk}^{111}x_jx_k + 2Q_{ij\beta}^{112}x_jy_\beta + Q_{i\beta\gamma}^{122}y_\beta y_\gamma \quad (7.1a)$$

$$\dot{y}_\alpha = f_\alpha^y(\mathbf{x}, \mathbf{y}) = F_\alpha^2 + L_{\alpha j}^{21}x_j + L_{\alpha\beta}^{22}y_\beta + Q_{\alpha jk}^{211}x_jx_k + 2Q_{\alpha j\beta}^{212}x_jy_\beta + Q_{\alpha\beta\gamma}^{222}y_\beta y_\gamma. \quad (7.1b)$$

Standard truncation of this system is equivalent to retaining only the terms \mathbf{F}^1 , \mathbf{L}^{11} , and \mathbf{Q}^{111} . Here we will attempt to approximate the terms involving the fast variables in $\mathbf{f}^x(\mathbf{x}, \mathbf{y})$ in an average sense in order to derive a closed system $\dot{\mathbf{x}} = \tilde{\mathbf{f}}^x(\mathbf{x})$, largely following the approach of Pavliotis & Stuart (2012).

We assume that the time scales of the fast and slow dynamics are separated by a parameter $\epsilon \ll 1$, so that we may define $\mathbf{f}^x(\mathbf{x}, \mathbf{y}) \equiv \epsilon \tilde{\mathbf{f}}^x(\mathbf{x}, \mathbf{y})$. Furthermore, we note that the role of the fast self-interaction term $\mathbf{Q}^{222}(\mathbf{y}, \mathbf{y})$ is to transfer energy between the unresolved scales. Since this mechanism is of secondary importance to the transfers between slow and fast scales, we apply a version of the “working assumption of stochastic modeling” (Majda *et al.*, 2001):

$$Q_{\alpha\beta\gamma}^{222} y_\beta y_\gamma \approx \sigma_\alpha w_\alpha(t), \quad (7.2)$$

where $w(t)$ is a Wiener process and σ is an as-yet-undefined constant forcing amplitude. Finally, we coarse-grain the dynamics on the time scale $\tau = \epsilon t$, leading to the slow/fast system

$$\frac{\partial x_i}{\partial \tau} = \tilde{f}_i^x(\mathbf{x}, \mathbf{y}) \quad (7.3a)$$

$$\frac{\partial y_\alpha}{\partial \tau} = \frac{1}{\epsilon} f_\alpha^y(\mathbf{x}, \mathbf{y}) + \frac{1}{\sqrt{\epsilon}} \sigma_{\alpha\beta} w_\beta(\tau), \quad (7.3b)$$

where we have also used the scaling property of Wiener processes that $w(\epsilon t) = \sqrt{\epsilon} w(t)$.

Defining an arbitrary scalar-valued observable $g^\epsilon(\mathbf{x}, \mathbf{y})$, the backwards Kolmogorov equation (2.43) associated with Eq. (7.3) is

$$\frac{\partial g^\epsilon}{\partial \tau} = \epsilon^{-1} \underbrace{\left[f_\alpha^y(\mathbf{x}, \mathbf{y}) \frac{\partial}{\partial y_\alpha} + \frac{\sigma_\alpha^2}{2} \frac{\partial^2}{\partial y_\alpha^2} \right]}_{\mathcal{L}_0} g^\epsilon + \underbrace{\tilde{f}_i^x(\mathbf{x}, \mathbf{y}) \frac{\partial}{\partial x_i}}_{\mathcal{L}_1} g^\epsilon. \quad (7.4)$$

Note that \mathcal{L}_0 can be viewed as the generator of a stochastic process in \mathbf{y} with \mathbf{x} as a fixed parameter. Consequently, we make the following assumptions related to the ergodicity of \mathbf{y} :

1. The operator \mathcal{L}_0 has a one-dimensional nullspace spanned by constants in \mathbf{y} :

$$\mathcal{L}_0 g(\mathbf{x}) = 0 \quad (7.5)$$

2. The Fokker-Planck operator \mathcal{L}_0^\dagger has a one-dimensional nullspace corresponding to the stationary distribution ρ_x^∞ , where again \mathbf{x} is treated as a fixed parameter:

$$\mathcal{L}_0^\dagger \rho_x^\infty(\mathbf{y}) = 0, \quad (7.6)$$

along with the usual normalization condition $\int_{\mathbf{y}} \rho_x^\infty d\mathbf{y} = 1$.

These assumptions obviate the need to express the fast scale \mathbf{y} as an instantaneous function of \mathbf{x} , as in an invariant manifold model (Guckenheimer & Holmes, 1983; Pavliotis & Stuart, 2012). Instead, the fast variable is modeled with a simplified distribution that can be averaged over as follows. We assume that the backwards Kolmogorov equation (7.4) can be approximated by means of an asymptotic expansion

$$g^\epsilon(\mathbf{x}, \mathbf{y}, \tau) = g_0 + \epsilon g_1 + \mathcal{O}(\epsilon^2). \quad (7.7)$$

Substituting into Eq. (7.4) and equating powers of ϵ gives the consistency conditions

$$0 = \mathcal{L}_0 g_0, \quad (7.8a)$$

$$\frac{\partial g_0}{\partial \tau} = \mathcal{L}_0 g_1 + \mathcal{L}_1 g_0. \quad (7.8b)$$

By virtue of the assumption of a one-dimensional nullspace, Eq. (7.8a) is satisfied if g_0 is not a function of \mathbf{y} , or $g_0 = g_0(\mathbf{x}, \tau)$. As expected, the leading-order solution does not depend on the fast variable.

An effective evolution equation for g_0 can be derived by considering Eq. (7.8b) as a linear equation for $g_1(\mathbf{x}, \mathbf{y}, \tau)$. The Fredholm alternative introduced in Sec. 2.2 in the context of weakly nonlinear analysis specifies that for any equation of the form $\mathcal{L}_0 g_1 = b$ to have a unique solution, all functions ρ in the nullspace of the adjoint operator \mathcal{L}_0^\dagger must be orthogonal to b . Since we have assumed that the nullspace of \mathcal{L}_0^\dagger is one-dimensional and spanned by the stationary distribution $\rho_{\mathbf{x}}^\infty(\mathbf{y})$, this implies that $\langle \rho_{\mathbf{x}}^\infty, b \rangle = 0$. In other words, the solvability condition for

$$\mathcal{L}_0 g_1 = -\frac{\partial g_0}{\partial \tau} + \mathcal{L}_1 g_0 \quad (7.9)$$

is that the right-hand side of Eq. (7.9) has zero mean with respect to the stochastic process generated by \mathcal{L}_0 :

$$\int_{\mathbf{y}} \rho_{\mathbf{x}}^\infty(\mathbf{y}) \left[-\frac{\partial}{\partial \tau} + \mathcal{L}_1 \right] g_0(\mathbf{x}, \tau) d\mathbf{y} = 0. \quad (7.10)$$

Using the normalization condition for $\rho_{\mathbf{x}}^\infty$, this simplifies to a closed evolution equation for $g_0(\mathbf{x}, \tau)$:

$$\frac{\partial g_0}{\partial \tau} = \left[\int_{\mathbf{y}} \rho_{\mathbf{x}}^\infty(\mathbf{y}) \tilde{f}_i^x(\mathbf{x}, \mathbf{y}) d\mathbf{y} \right] \frac{\partial g_0}{\partial x_i}. \quad (7.11)$$

Again, the observable g_0 itself is arbitrary, but Eq. (7.11) has the form of the generator of a Koopman operator corresponding to the coarse-grained dynamics in \mathbf{x} alone. Undoing the ϵ scaling in τ and \tilde{f}^x , Eq. (7.11) can be written as

$$\frac{\partial g_0}{\partial t} = \hat{\mathbf{f}}(\mathbf{x}) \cdot \nabla g_0, \quad (7.12)$$

corresponding to the averaged dynamics

$$\dot{\mathbf{x}} = \hat{\mathbf{f}}(\mathbf{x}), \quad (7.13a)$$

$$\hat{\mathbf{f}}(\mathbf{x}) = \int_{\mathcal{Y}} \rho_{\mathbf{x}}^{\infty}(\mathbf{y}) f_i^x(\mathbf{x}, \mathbf{y}) d\mathbf{y}. \quad (7.13b)$$

The distribution $\rho_{\mathbf{x}}^{\infty}(\mathbf{y})$ specifies the probability distribution of the fast variables \mathbf{y} and is stationary on the fast timescale, but implicitly time-varying since it is parameterized by the state \mathbf{x} of the slow variables. In the simplest case $\rho_{\mathbf{x}}^{\infty}$ might be proportional to a delta function in \mathbf{x} , indicating that the fast variables are a direct function of the slow variables. Physically this might correspond to the case where \mathbf{y} represents a slow amplitude-dependent deformation of the base flow or phase-locked higher harmonics, for instance.

More generally the functional form of this distribution might be complicated, but it is not necessary to specify it in closed form provided the integral in Eq. 7.13b can be evaluated. For instance, when the dynamics are linear-quadratic then Eq. 7.13b simplifies to first and second moments of the distribution. Then the key to the MMR closure for the POD-Galerkin is deriving an approximation of the fast dynamics that is consistent with the original system but allows for computation of these moments in closed form.

As an aside, although the disappearance of the fictitious small parameter ϵ is necessary for consistency with the original Galerkin system, it does call into question the validity of the perturbation series approximation (7.7). In this work we do not attempt to make this more rigorous, but instead demonstrate by example that it is a useful heuristic capable of resolving important features of the flow physics. We justify this based on the observation that the underlying linear-quadratic Galerkin systems often poorly approximate the true dynamics, so it would not be useful to perfectly match the original model even if it was possible to do so.

Stochastic averaging for a model system Before applying the formal solution (7.13) to the linear-quadratic Galerkin dynamics, we first illustrate the stochastic averaging procedure on a simple two-dimensional fast-slow system:

$$\frac{dx}{d\tau} = \lambda x - xy \quad (7.14a)$$

$$\epsilon \frac{dy}{d\tau} = -y + \mu x^2 + \sqrt{\epsilon} \sigma w(t), \quad (7.14b)$$

This SDE system is associated with the following backwards Kolmogorov equation for an arbitrary scalar observable $g^{\epsilon}(x, y, \tau)$:

$$\frac{\partial g^{\epsilon}}{\partial \tau} = \epsilon^{-1} \underbrace{\left[(-y + \mu x^2) \frac{\partial}{\partial y} + \frac{\sigma^2}{2} \frac{\partial^2}{\partial y^2} \right]}_{\mathcal{L}_0} g^{\epsilon} + \underbrace{\left[(\lambda x - xy) \frac{\partial}{\partial x} \right]}_{\mathcal{L}_1} g^{\epsilon} \quad (7.15)$$

Here \mathcal{L}_0 is the backwards Kolmogorov operator associated with an Ornstein-Uhlenbeck process in y , which we assume satisfies the ergodicity requirements (7.5) and (7.6). By comparison of \mathcal{L}_0^\dagger with the steady-state Fokker-Planck equation for an Ornstein-Uhlenbeck process (2.47), the stationary distribution ρ_x^∞ is a Gaussian with mean and variance

$$\int_{-\infty}^{\infty} y \rho_x^\infty(y) dy = \mu x^2, \quad \int_{-\infty}^{\infty} y^2 \rho_x^\infty(y) dy = \frac{\sigma^2}{2}. \quad (7.16)$$

In this case, the solvability condition (7.10) reduces to

$$\frac{\partial g_0}{\partial \tau} = \left[\int_{-\infty}^{\infty} \rho_x^\infty(y) (\lambda x - xy) dy \right] \frac{\partial g_0}{\partial x} = 0. \quad (7.17)$$

Comparing with Eq. (7.11), this condition corresponds to a Koopman operator for the averaged dynamics. Using the normalization condition on the probability distribution ρ_x^∞ and the mean given by Eq. (7.16), the fast variable can be integrated analytically, giving the coarse-grained dynamics

$$\frac{dx}{d\tau} = \int_{-\infty}^{\infty} \rho_x^\infty(y) (\lambda x - xy) dy \quad (7.18)$$

$$= \lambda x - \mu x^3. \quad (7.19)$$

Finally, we note that in this case the same solution could have been reached by neglecting the diffusion term in the y dynamics from the beginning, as in an invariant manifold-type model (Guckenheimer & Holmes, 1983; Noack *et al.*, 2003; Pavliotis & Stuart, 2012), since due to the form of the slow dynamics the effect of the noise disappears on averaging. This would give the approximation $y = \mu x^2$, which could then be substituted into the x -dynamics to recover Eq. (7.18). The difference between that approach and the one described here lies in Eq. (7.17); rather than assuming the fast variables can be expressed as an instantaneous function of the slow variables, they are assumed to fluctuate and their effect is accounted for by averaging. Later we will discuss some general conditions under which the two are equivalent, but in this example if there was a term proportional to y^2 in the x dynamics, the solvability condition would include the second moment of y and the diffusion σ would appear in the coarse-grained dynamics.

Application to the Galerkin system As illustrated by the previous example, in order to make practical use of the averaging procedure, we must be able to perform the integral over the distribution $\rho_x^\infty(\mathbf{y})$ of the fast variables. This distribution is the solution for fixed x to the steady-state Fokker-Planck equation

$$\mathcal{L}_0^\dagger \rho = -\frac{\partial(\rho f_\alpha^y)}{\partial y_\alpha} + \frac{\sigma_\alpha^2}{2} \frac{\partial^2 \rho}{\partial y_\alpha^2} = 0, \quad (7.20)$$

corresponding to the linear stochastic process

$$\dot{y}_\alpha = [F_\alpha^2 + L_{\alpha j}^{21}x_j + Q_{\alpha j k}^{211}x_jx_k] + [L_{\alpha\beta}^{22} + 2Q_{\alpha j\beta}^{212}x_j]y_\beta + \sigma_\alpha w_\alpha(t). \quad (7.21)$$

The stationary distribution is a multivariate Gaussian with mean and covariance that can be determined analytically from the solution of a Lyapunov equation (Risken, 1996), but this would need to be done at each value of \mathbf{x} . Instead, we propose the diagonal drift approximation

$$\dot{y}_\alpha \approx \nu_\alpha(\mathbf{x}) (\mu_\alpha(\mathbf{x}) - y_\alpha) + \sigma_\alpha w_\alpha(t), \quad (7.22)$$

for which the stationary distribution is a product of univariate Gaussians solving the 1-dimensional Fokker-Planck equation (2.47), i.e. $y_\alpha \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha^2/2\nu_\alpha)$. By comparison with Eq. (7.21), the conditional mean is naturally defined as

$$\mu_\alpha(\mathbf{x}) = \sum_{j,k=1}^r \nu_\alpha^{-1} [L_{\alpha j}^{21}x_j + Q_{\alpha j k}^{211}x_jx_k]. \quad (7.23)$$

Here we have omitted the contribution of the constant forcing term \mathbf{F}^2 so that the approximate fast process \mathbf{y} preserves the zero-mean property of the POD coefficients for $\mathbf{x} = 0$. In this work we will make the simplifying assumption that the effective damping coefficients ν_α are constant, although more generally they could be functions of \mathbf{x} . Appropriate values for ν_α can be determined from energy balance, as described below.

With these approximation the average in Eq. (7.13) is a straightforward calculation resulting in the generalized Stuart-Landau model

$$\hat{f}_i(\mathbf{x}) = F_i^1 + L_{ij}^{11}x_j + L_{i\alpha}^{12}\mu_\alpha(\mathbf{x}) + Q_{ijk}^{111}x_jx_k + 2Q_{ij\alpha}^{112}x_j\mu_\alpha(\mathbf{x}) + \nu_\alpha^{-1}Q_{i\alpha\alpha}^{122}\sigma_\alpha\sigma_\alpha/2 \quad (7.24a)$$

$$\equiv \hat{F}_i + \hat{L}_{ij}x_j + \hat{Q}_{ijk}x_jx_k + \hat{C}_{ijkl}x_jx_kx_l, \quad (7.24b)$$

with the following closed quantities denoted by a hat:

$$\hat{F}_i = F_i^1 + \nu_\alpha^{-1}Q_{i\alpha\alpha}^{122}\sigma_\alpha\sigma_\alpha/2 \quad (7.25a)$$

$$\hat{L}_{ij} = L_{ij}^{11} + \nu_\alpha^{-1}L_{i\alpha}^{12}L_{\alpha j}^{21} \quad (7.25b)$$

$$\hat{Q}_{ijk} = Q_{ijk}^{111} + \nu_\alpha^{-1}(L_{i\alpha}^{12}Q_{\alpha j k}^{211} + 2Q_{ij\alpha}^{112}L_{\alpha k}^{21}) \quad (7.25c)$$

$$\hat{C}_{ijkl} = 2\nu_\alpha^{-1}Q_{ij\alpha}^{112}Q_{\alpha kl}^{211}. \quad (7.25d)$$

We will refer to the secondary reduction of a standard rank- r POD-Galerkin model to the cubic model (7.24b) as the multiscale model reduction (MMR) approach to closure modeling.

The terms in the closure model are computed by summing the Galerkin tensors over the fast variables. As a mnemonic for these averaged quantities, the superscripts for the partitioned system could be thought of as tensor contractions with the modification of the

damping ν_α^{-1} . For example, the cubic term $Q_{ij\alpha}^{112}\nu_\alpha^{-1}Q_{\alpha kl}^{211}$ is suggestive of the slow variables “filtering through” the fast dynamics via the quadratic interactions \mathbf{Q}^{211} and \mathbf{Q}^{112} .

Although the constant, linear, and quadratic terms are all modified as a result of the stochastic averaging procedure, the appearance of this cubic term is perhaps the most noteworthy. It represents the leading-order contribution of the fast-slow nonlinear interaction in the slow dynamics due to the slow-slow interaction in the fast dynamics. We will explore this in more detail in the following section, but it is a generalization of the weakly nonlinear Stuart-Landau mechanism (Stuart, 1958; Landau & Lifshitz, 1959) that is capable of resolving the stabilizing influences of both mean flow deformation and the energy cascade.

The final ingredient in the multiscale closure model (7.24b) is the determination of the damping and diffusion coefficients ν_α and σ_α . A reasonable approximation for the diffusion could be the mean of the neglected term $\sigma_\alpha = \sum_\beta Q_{\alpha\beta\beta}^{222}\overline{y_\beta^2}$, although in our numerical examples we find little difference with neglecting the diffusion altogether.

For the damping term we apply an energy balance condition on the closed model, derived from the assumption that the system is statistically stationary so that the average variation of kinetic energy vanishes (Noack *et al.*, 2011):

$$\frac{d}{dt}\overline{x_i^2} = \overline{\hat{f}_i(\mathbf{x})x_i} = 0. \quad (7.26)$$

When the closed tensors (7.25) are substituted for $\hat{f}(\mathbf{x})$, this simplifies to a linear system of equations for ν^{-1} . The energy balance approximation does not require statistics of the fast variables, but due to the cubic term it does require well-converged fourth moments of the slow variables. In particular, the vector ν_{inv} which is the element-wise inverse of ν is the solution to the linear system

$$\mathbf{A}\nu_{\text{inv}} + \mathbf{b} = 0, \quad (7.27a)$$

$$\begin{aligned} A_{i\alpha} = & \sum_{j=1}^{r_0} L_{i\alpha}^{12} L_{\alpha j}^{21} \overline{a_i a_j} + \sum_{j,k=1}^{r_0} L_{i\alpha}^{12} Q_{\alpha j k}^{211} \overline{a_i a_j a_k} \\ & + 2 \sum_{j,k=1}^{r_0} Q_{ij\alpha}^{112} L_{\alpha k}^{21} \overline{a_i a_j a_k} + 2 \sum_{j,k,\ell=1}^{r_0} Q_{ij\alpha}^{112} Q_{\alpha k \ell}^{211} \overline{a_i a_j a_k a_\ell} \end{aligned} \quad (7.27b)$$

$$b_i = L_{ii}^{11} \overline{a_i^2} + \sum_{j,k=1}^{r_0} Q_{ijk}^{111} \overline{a_i a_j a_k} \quad (7.27c)$$

In general the damping should be positive since the linear stochastic approximation (7.21) and its diagonal simplification is only plausible if the mean $\mu(\mathbf{x})$ is linearly stable for all \mathbf{x} .

This is physically consistent with the energy cascade picture, in which we expect that energy will primarily flow “downhill” from the slow variables, representing large coherent structures and global instabilities, to the fast variables representing the smaller, dissipative scales of the flow. On this basis negative values of the damping coefficients can also be set to zero to avoid introducing unphysical instabilities in the cubic term, for example.

We conclude the discussion of the application to Galerkin-type systems with a note on the computational scaling of the closure model. While the simulation of the original linear-quadratic model is dominated by the quadratic term, which requires $\mathcal{O}(r^3)$ operations to evaluate, the cubic term in the closure model (7.24b) requires $\mathcal{O}(r_0^4)$ operations to evaluate. In some cases this may mean that the cubic model is actually more expensive to simulate than the original Galerkin system, although whether or not this is the case will depend on the specific values of r , the dimension of the linear-quadratic system, and r_0 , the dimension of the closed model. For the examples given below either $r_0^4 < r^3$, as in Secs. 7.2 and 7.3, or the two are the same order of magnitude, as in Sec. 7.4. As we will show, the primary advantage of this approach is not necessarily that it is faster to simulate than the usual POD-Galerkin system, but that it is more stable and physically faithful with many fewer modes.

Relationship to invariant manifold modeling The invariant manifold reduction, introduced in Sec. 2.3 and applied to the cylinder flow in Sec. 5.4, is a method in which the fast variables are assumed to depend algebraically on the slow variables. It is often applied near the threshold of bifurcation, when one may assume that there are only a small number of weakly unstable directions in the modal subspace and that the modal coefficients are typically small.

Under these assumptions, the fast variables can be approximated with a power series expansion with coefficients determined by a dynamical consistency condition (Guckenheimer & Holmes, 1983; Pavliotis & Stuart, 2012). For example, if the dynamics are close to a supercritical Hopf bifurcation then the leading-order approximation is a parabolic manifold that depends quadratically on the slow variables; replacing the fast variables with this approximation results in a cubic evolution equation for the slow variables similar to Eq. (7.24b). This approach is closely related to the technique of center manifold reduction, which is typically applied to the large-scale system at the bifurcation point, yielding normal form-type amplitude equations for the evolution of the instability modes (Coullet & Spiegel, 1983; Carini *et al.*, 2015). Here we recapitulate the parabolic invariant manifold approximation in order to give a comparison with the stochastic averaging method, but we will conclude by showing that a different manifold model gives an equivalent result to Eq. (7.25) for $\sigma = 0$.

The central assumption of the invariant manifold model is that the fast variables can be expressed as an instantaneous function $\mathbf{y} \approx \varphi(\mathbf{x})$ of the slow variables. Dynamical

consistency for the fast variables then requires that, to leading order,

$$\dot{\mathbf{y}} = \mathbf{f}^y(\mathbf{x}, \varphi(\mathbf{x})) = \frac{d}{dt}\varphi(\mathbf{x}) \quad (7.28)$$

$$= \mathbf{J}_\varphi(\mathbf{x})\mathbf{f}^x(\mathbf{x}, \varphi(\mathbf{x})) \quad (7.29)$$

where $\mathbf{J}_\varphi(\mathbf{x})$ is the Jacobian of φ evaluated at \mathbf{x} (Guckenheimer & Holmes, 1983; Pavliotis & Stuart, 2012).

Assuming all of the state variables in \mathbf{b} are typically small, the manifold equation can be expanded as a power series in \mathbf{x} . Preservation of the fixed point at the origin requires $\mathbf{h}(0) = 0$, and in the case of a center manifold reduction where the real parts of the slow eigenvalues vanish, the condition that the $\varphi(\mathbf{x})$ be tangent to the center manifold at the origin also implies that $\mathbf{J}_\varphi(0) = 0$. Under these conditions the first nontrivial terms in the expansion are quadratic and the expansion coefficients can be determined by matching powers of \mathbf{x} in the consistency condition (7.29).

The connection between this methods and stochastic averaging can be made clear by viewing $\mathbf{y}(t) = \varphi(\mathbf{x}, t)$ as an explicit function of time and rewriting the dynamical consistency condition as a Koopman generator acting on the α^{th} component of the manifold equation:

$$\frac{\partial \varphi_\alpha}{\partial t} = \mathbf{f}^x(\mathbf{x}, \varphi(\mathbf{x})) \cdot \nabla \varphi_\alpha. \quad (7.30)$$

Since the observable in such an equation is arbitrary, this is equivalent to the averaged backwards Kolmogorov equation (7.12) if

$$f_i^x(\mathbf{x}, \varphi(\mathbf{x})) = \int_{\mathbf{y}} \rho_{\mathbf{x}}^\infty(\mathbf{y}) f_i^x(\mathbf{x}, \mathbf{y}) d\mathbf{y}. \quad (7.31)$$

In the limit of zero diffusion in the fast process, the stationary distribution approaches a Dirac delta function $\rho_{\mathbf{x}}^\infty = \delta^{N-r}(\mathbf{y} - \boldsymbol{\mu}(\mathbf{x}))$ and the two are equivalent with $\varphi(\mathbf{x}) = \boldsymbol{\mu}(\mathbf{x})$. With the diagonal drift approximation (7.22) and constant damping ν , the zero-diffusion limit of the proposed Galerkin closure in Sec. 7.1 can be seen as an invariant manifold model with $\varphi(\mathbf{x})$ defined as the conditional mean (7.23). Because we derived the stochastic averaging model with the assumption of real-valued coefficients and the parabolic model uses the complex-valued eigenvector coordinates, it is less obvious that the parabolic invariant manifold corresponds to a particular choice of ν and $\boldsymbol{\mu}$ in Eq. (7.22).

Although the invariant manifold reduction is more theoretically straightforward than the stochastic averaging closure, it generally requires stronger assumptions on the system. In particular, the development of the parabolic manifold model requires that the origin be a fixed point and that the real part of the slow eigenvalues is zero. This does not raise any difficulties for simple weakly unstable configurations where the Galerkin model preserves fixed points of the fluid flow, or if the steady state is used as the base of the modal expansion. However, when the flow is far from a bifurcation point the mean flow is often

	MMR	FTT
ρ	$\rho(\mathbf{y}; \mathbf{x})$	$\rho(\mathbf{y}; t)$
$\mathbb{E}[y_i]$	$\mu_i(\mathbf{x})$	0
$\mathbb{E}[y_i^2]$	σ_i^2	$2E_i(t)$

Table 7.1: Statistics of unresolved variables for the multiscale and finite-time thermodynamics models. MMR assumes the distribution is parameterized by the slow variables \mathbf{x} , while FTT assumes it is independent and zero-mean, although with possibly fluctuating variance. We have found the variance σ_i^2 to be negligible in MMR if it is assumed constant, but more generally it could be also modeled as a function of \mathbf{x} . Note that the unresolved variables \mathbf{y} are also defined differently between the two methods (see Sec. 7.1).

a more relevant base of the expansion; in this case the resulting Galerkin models may not necessarily have fixed points that coincide with meaningful states of the underlying system. In addition, the invariant manifold reduction relies on smallness of the coefficients (justifying the power series expansion) and the assumption that the higher-order modes can be expressed as instantaneous algebraic functions of the active coefficients. For chaotic or turbulent dynamics neither of these is necessarily justified.

Finite-time thermodynamics Before concluding the discussion of the multiscale closure model, we give a brief comparison to the finite-time thermodynamics (FTT) framework introduced by Noack *et al.* (2008). This approach begins from the linear-quadratic POD-Galerkin system, Eq. (4.7), but rather than assuming that the unresolved variables are the higher-order modal coefficients, the FTT model decomposes the modal coefficients $\mathbf{a}(t) \in \mathbb{R}^r$ into an “ensemble-averaged” value $\mathbf{x}(t) \in \mathbb{R}^r$ and a “fluctuating” value $\mathbf{y}(t) \in \mathbb{R}^r$ in the spirit of the Reynolds decomposition. As a result, both the resolved and unresolved variables correspond to the *same* spatial modes, and the modal expansion (3.1) becomes

$$\mathbf{u}(\mathbf{x}, t) \approx \boldsymbol{\psi}_0(\mathbf{x}) + \sum_{i=1}^r (x_i(t) + y_i(t)) \boldsymbol{\psi}_i(\mathbf{x}). \quad (7.32)$$

The second major difference is the assumption of the behavior of the unresolved variables \mathbf{y} on ensemble averaging. The ensemble average is with respect to the probability distribution ρ of the fluctuations \mathbf{y} . The MMR method assumes that this distribution is parameterized explicitly by the slow variables \mathbf{x} , so that $\rho = \rho(\mathbf{y}; \mathbf{x}(t))$ (written as $\rho_{\mathbf{x}}^{\infty}(\mathbf{y})$ in the derivation of MMR above). On the other hand, the Reynolds-averaging analogy in FTT leads to the assumption that the distribution ρ may be time-varying but does not depend on \mathbf{x} , so that $\rho = \rho(\mathbf{y}; t)$.

By definition, the fluctuating amplitudes \mathbf{y} have zero mean with respect to this distribution, so that

$$\mathbb{E}[y_i] \equiv \int_{\mathbf{y}} y_i \rho(\mathbf{y}; t) d\mathbf{y} = 0, \quad i = 1, 2, \dots, r. \quad (7.33)$$

Here we use the notation $\mathbb{E}[\cdot]$ to denote the ensemble average with respect to the unknown distribution $\rho(\mathbf{y}, t)$ to distinguish it from the usual time average of observed values used in the rest of this work.

As a result of Eq. (7.33), the mean and fluctuating amplitudes are linearly uncorrelated with respect to this average:

$$\mathbb{E}[x_i y_j] = 0, \quad i, j = 1, 2, \dots, r. \quad (7.34)$$

The normalization of the probability distribution also implies that $\mathbb{E}[x_i] = x_i$. Finally, FTT assumes that the fluctuations themselves are uncorrelated, with the usual variance-kinetic energy relationship for POD coefficients. Combining these assumptions, the second-order statistics of the POD coefficients with respect to the ensemble average defined by $\rho(\mathbf{y}; t)$ are

$$\mathbb{E}[a_i a_j] = x_i x_j + 2E_i(t) \delta_{ij}, \quad (7.35)$$

where $E_i(t) = \mathbb{E}[y_i^2]/2$. A comparison of the assumptions on the distributions is shown in Table 7.1. Note that since the variance E_i must be non-negative, Eq. (7.35) implies that the statistics of \mathbf{a} and \mathbf{x} are not identical unless $\mathbf{E}(t) = 0$ for all time.

Expanding the POD coefficients into mean and fluctuating component and taking the ensemble average gives a set of $2r$ evolution equations for the mean values \mathbf{x} and the fluctuation kinetic energy \mathbf{E} :

$$\dot{x}_i = \sum_{j=1}^r L_{ij} x_j + \sum_{j,k=1}^r Q_{ijk} x_j x_k + 2 \sum_{j=1}^r Q_{ijj} E_j \quad (7.36a)$$

$$\dot{E}_i = 2 \sum_{j=1}^r (L_{ii} + 2Q_{iij} x_j) E_j + \sum_{j,k=1}^r Q_{ijk} \mathbb{E}[y_i y_j y_k]. \quad (7.36b)$$

The remainder of the FTT method centers on deriving a nonlinear closure model for the third-order statistics $\mathbb{E}[y_i y_j y_k]$; see Noack *et al.* (2008) for details.

The main similarity between the FTT and multiscale methods is that both seek to approximate the effect of unresolved fluctuations on the resolved variables, resulting in nonlinear closure models. However, the two are philosophically different in that MMR assumes the unresolved variables correspond to different spatial structures and evolve on faster time scales, while FTT models them as linearly uncorrelated fluctuations associated with the same spatial modes as the resolved variables.

Noack *et al.* (2008) showed that the FTT model is quite successful at predicting post-transient energy levels in several flow configurations, including the cylinder wake and a turbulent homogeneous shear flow. This method also has the advantage that it does not assume any separation of scales in amplitude or time. On the other hand, there are some theoretical difficulties in working with the FTT framework, including that Eq. (7.35) muddies the relationship between the second-order statistics of the modal coefficients \mathbf{a}

and x , which are foundational to the POD method. FTT also increases the state dimension by a factor of two and cannot be used to derive the generalized Stuart-Landau equations that directly account for effects such as mean flow deformation and the energy cascade. However, Cordier *et al.* (2013) showed that a simplified nonlinear closure model based on Eq. (7.36) has been shown to stabilize POD-Galerkin models of a mixing layer similar to that examined in Sec. 7.4 without increasing the state dimension. Although this FTT closure model is energetically stable, we will show that it does not reproduce the statistics or phase relationships of the DNS in the present case.

7.2 Prototypical fast effects

There are two principal mechanisms by which nonlinear interactions with higher-order modes effectively stabilize the flow: Reynolds stress-induced mean flow deformation and the energy cascade. Before demonstrating the MMR closure introduced in Sec. 7.1 on DNS of unsteady fluid flow, we first confirm that it can resolve both of these effects by examining simple representative analytic models.

Nonlinear stability in a mean-field model Nonlinear stability analysis was born out of the observation that the exponential energy growth predicted by the linear theory when the flow has at least one unstable eigenvalue is untenable at long times, and furthermore that the key assumption of infinitesimal perturbation amplitudes is violated once the instability grows to finite amplitude. It is conceivable that the energy growth eventually leads to a breakdown into turbulence, but it was observed experimentally (e.g. Taylor (1923)) that this is not always the case.

Instead, if the perturbations are assumed to be of small but finite amplitude, the linear stability problem can be considered the first term in an asymptotic expansion. The next-order terms suggest that the Reynolds stresses from the nonlinear self-interaction of the perturbation tends to deform the slowly-varying background flow from the unstable steady-state towards a neutrally stable mean flow Stuart (1958). This type of weakly nonlinear analysis has since been applied to a wide range of flows and bifurcation scenarios (Sipp & Lebedev, 2007; Meliga *et al.*, 2009; Meliga & Chomaz, 2011; Rigas *et al.*, 2017). Although this has proven to be a highly successful approach, it is only strictly applicable near the threshold of bifurcation and requires invasive access to various linear operators. The self-consistent mean flow modeling framework (Mantič-Lugo *et al.*, 2014; Mantič-Lugo & Gallaire, 2016; Meliga, 2017; Rigas *et al.*, 2021) is more general in terms of parametric dependence, but it does not lead to a low-dimensional dynamical system.

Resolving mean flow deformation was shown to be a key ingredient in stabilizing reduced-order models. While models that only included POD modes derived from expansion around the mean gave accurate approximations on short time horizons but eventually tended towards structural instability (Deane *et al.*, 1991; Ma & Karniadakis, 2002), explicitly incorporating this stabilizing feedback mechanism via Reynolds stress model-

ing (Aubry *et al.*, 1988; Holmes *et al.*, 1996), an augmented POD basis (Noack *et al.*, 2003; Luchtenburg *et al.*, 2009; Noack *et al.*, 2011), or a normal form ansatz (Deng *et al.*, 2020) in models of globally unstable flows significantly improved their robustness. Here we show that if the POD basis spans the mean flow deformation, the proposed multiscale model reduction method can resolve its stabilizing effect without introducing additional degrees of freedom. This is consistent with the invariant manifold reduction approach of Noack *et al.* (2003), but does not require treating the mean flow deformation differently than any of the other unresolved fast variables.

A classic illustrative model of the nonlinear stability mechanism is the three-dimensional “mean-field” system proposed by Noack *et al.* (2003), identical to the model problem in Sec. 2.3 without nonlinear frequency modulation. The model consists of a pair of unstable slow variables a_1 and a_2 along with the fast variable a_Δ :

$$\dot{a}_1 = \lambda a_1 - \omega a_2 - a_1 a_\Delta \quad (7.37a)$$

$$\dot{a}_2 = \omega a_1 - \lambda a_2 - a_2 a_\Delta \quad (7.37b)$$

$$\dot{a}_\Delta = -a_\Delta + \mu(a_1^2 + a_2^2). \quad (7.37c)$$

The slow variables represent amplitudes of a pair of global instability modes, while a_Δ models a background deformation that tends to reduce the linear growth rate. This system is a simplified model with similar behavior to the cylinder wake discussed in Sec. 7.3.

In many cases the growth rate λ is relatively small, so that the standard approach is to use the approximation $\dot{a}_\Delta = 0$. The fast variable is then given by an algebraic relationship that pins the system to the parabolic invariant manifold $a_\Delta \approx \mu(a_1^2 + a_2^2)$. Replacing a_Δ in Eq. (7.37) to a single cubic Stuart-Landau equation for the complex amplitude $A(t) = a_1 + ia_2$:

$$\dot{A} = (\lambda + i\omega)A - \mu A|A|^2. \quad (7.38)$$

We restate these results in order to show that this mechanism is also resolved by the multiscale closure (7.24b). Here the fast dynamics are already linear in a_Δ with damping constant $\nu = 1$. The fast-slow linear blocks \mathbf{L}^{12} and \mathbf{L}^{21} are also zero, so the only modification due to the closure model is the cubic term $\hat{\mathbf{C}}$. The only nonzero entries in the quadratic tensor are $Q_{11\Delta} = Q_{22\Delta} = -1$ and $Q_{\Delta 11} = Q_{\Delta 22} = \mu$, leading to cubic terms $\hat{C}_{1111} = \hat{C}_{1122} = \hat{C}_{2211} = \hat{C}_{2222} = -\mu$. Replacing the fast-slow nonlinear interactions in Eq. (7.37) with the cubic closure, the three-dimensional model reduces to

$$\dot{a}_1 = \lambda a_1 - \omega a_2 - \mu a_1(a_1^2 + a_2^2) \quad (7.39a)$$

$$\dot{a}_2 = \omega a_1 - \lambda a_2 - \mu a_2(a_1^2 + a_2^2), \quad (7.39b)$$

identical to the real and imaginary parts of the invariant manifold reduction (7.38).

Energy transfer in the Burgers model Along with mean flow interactions, one of the primary challenges in applying model reduction to the Navier-Stokes equations is the under-resolution of the fine scales of the flow, which typically results in underestimating the energy dissipation. In the classic spectral dynamics picture of isotropic turbulence, although the bulk of the energy is contained in the larger scales of the flow, the viscous energy dissipation scales with the square of the wavenumber, so that even very small scales play a key role in the energy balance of the flow. The problem of truncating the cascade of energy from large to small scales thus applies equally well to “model reduction” in the sense of discretizing the governing equations on a mesh that cannot resolve the finest scales of the flow.

Since POD modes are ordered hierarchically by average energy content, this problem often manifests in reduced-order modeling when attempting to severely truncate a POD representation; the higher-order modes are often much more dissipative, even though they may not be of practical interest (Noack *et al.*, 2011). Further complicating this is the issue that advection-dominated flows may require many more modes for even a good energetic approximation than for a good dynamic approximation, and these additional degrees of freedom can lead to decoherence and instability in the reduced-order model, as shown in Chapter 6. The approach of partitioning the system and averaging over the “fast” variables can address both of these issues by retaining only the most energetic degrees of freedom while accounting for the average effect of unresolved dissipation at fine scales.

To illustrate the role of nonlinear advection Burgers (1948) and Hopf (1948) introduced simple one-dimensional toy models of turbulence, now generally studied as Burgers’ equation. Here we apply the multiscale closure approach to the Fourier-Galerkin representation of a forced, viscous Burgers’ equation on a periodic domain (Noack *et al.*, 2008):

$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} = \eta \frac{\partial^2 u}{\partial x^2} + \zeta g(x, t), \quad x \in (0, 2\pi), \quad (7.40)$$

where η and ζ are the viscous and forcing amplitudes, respectively, and the forcing $g(x, t)$ is constructed to act on only the leading wavenumbers. Majda *et al.* (2006) used a similar model to explore a similar stochastic model reduction procedure on large deterministic systems without asymptotic scale separation.

Since the domain is periodic, we forego the empirical POD basis and expand the solution in a Fourier basis

$$u(x, t) = \sum_{k=-\infty}^{\infty} a_k(t) e^{ikx}. \quad (7.41)$$

The standard Galerkin projection gives the coupled system of ODEs

$$\dot{a}_k = -\eta k^2 a_k - i\ell a_\ell a_{k-\ell} + \zeta \delta_{1|k|} a_k, \quad (7.42)$$

where the forcing reduces to the Kronecker delta function $\delta_{1|k|}$ by construction. The

quadratic nonlinearity would typically result in a double sum over wavenumbers, but the orthogonality property of the Fourier basis restricts the interaction to the “triads” $(k, \ell, k - \ell)$. As for the full Navier-Stokes equations, this nonlinearity satisfies the conservative property discussed in Sec. 4.1. The system is forced at large scales $|k| = 1$, dissipative at smaller scales $|k| \gg 1$, and the triadic interactions mediate energy transfers between these scales. Although the closure method described in Sec. 7.1 was derived for real-valued systems, Eq. (7.42) could easily be transformed to be purely real by separating the real and imaginary parts of the complex-conjugate pairs, as was done for instance by Noack *et al.* (2008). In this case the system is simple enough that it does not matter, so we work with the complex representation for compact notation.

We can apply the results of Sec. 7.1 directly by rewriting the quadratic term with the sparse tensor \mathbf{Q} defined by

$$Q_{klm} = \begin{cases} -i\ell, & \ell + m = k \\ 0, & \text{otherwise.} \end{cases} \quad (7.43)$$

The closure modeling is simplified by the diagonal structure of the linear term, so that the only modification to Eq. (7.42) is the cubic term $\hat{\mathbf{C}}$ given by Eq. (7.25d). Based on the sparse structure of \mathbf{Q} , the sum over fast scales disappears so that

$$\hat{C}_{klmn} = \frac{1}{\nu_{k-\ell}} \begin{cases} -km, & \ell + m + n = k \\ 0, & \text{otherwise.} \end{cases} \quad (7.44)$$

Note that these entries are also nonzero only for k, ℓ, m , and n in the slow scales and $k - \ell$ in the fast scales. For example, if we model only the forced scales $|k| = 1$ as slow variables, the only nonvanishing contributions to $\hat{\mathbf{C}}$ are from the wavenumber sets $(k, \ell, m, n) = (1, -1, 1, 1)$ and $(-1, 1, -1, -1)$. Since the slow quadratic tensor $\hat{\mathbf{Q}}$ is also zero for $|k| = 1$, the closed ODE system simplifies to the one-dimensional Stuart-Landau equation

$$\dot{a}_1 = (\zeta - \eta)a_1 - \nu_2^{-1}a_1|a_1|^2. \quad (7.45)$$

The cubic term in Eq. (7.45) captures the leading-order influence of the energy transfers to smaller scales. This effect is naturally stabilizing; as in the physical mechanism, if the slow variables are excited, the amplitude term $|a_1|^2$ representing the unresolved scales increases and the cubic damping term removes energy from the resolved scales.

In this case the multiscale closure directly accounts for only the next set of wavenumbers $|k| = 2$ because the purely triadic structure of the quadratic interactions does not admit contributions from any smaller scales. The damping constant ν_2 must therefore be selected so that the explicitly modeled transfers from $|k| = 1$ to $|k| = 2$ are representative of the entire cascade. This can be accomplished with the energy balance method described in Sec. 7.1, for example.

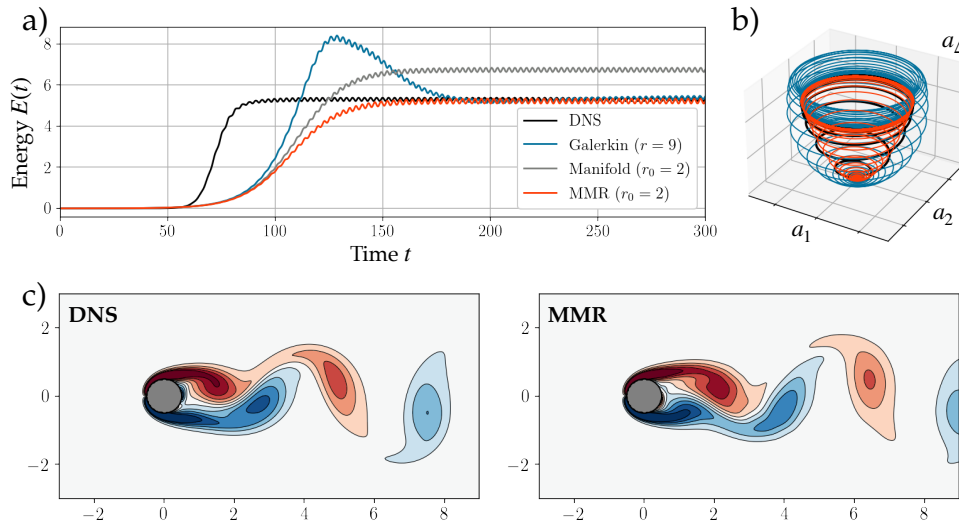


Figure 7.3: **Reduced-order models of the cylinder wake.** The standard 9-mode Galerkin model accurately estimates the stable limit cycle (a), but the transient dynamics deviate from the slow manifold, leading to an energy overshoot (b). Both the two-dimensional invariant manifold reduction in Sec. 5.4 and MMR closure models prevent this by eliminating the rapidly-equilibrating variable associated with mean flow deformation, but MMR more accurately estimates the limit cycle amplitude. The full flow field can be reconstructed by approximating the harmonics with polynomial regression (c).

While traditional closure models for Galerkin systems approximate the dissipation at higher orders with a linear linear eddy-viscosity-type closure (Aubry *et al.*, 1988; Rempfer & Fasel, 1994), the mechanism of energy transfer to smaller scales is fundamentally non-linear in nature. As clearly illustrated by this model problem, the multiscale closure model preserves the nonlinearity of the energy transfer in its approximation of the effects of the unresolved small scales. This is consistent with the findings of (Östh *et al.*, 2014) that a nonlinear term modeling subscale turbulence can outperform the linear eddy viscosity closure in POD-Galerkin models.

7.3 Vortex shedding in the cylinder wake

In this section we revisit the canonical two-dimensional flow past a circular cylinder examined at length in Chapter 5 to show that the MMR closure can naturally resolve the stabilizing effects both of mean flow deformation and of energy transfer to higher-order modes without assuming proximity to a bifurcation.

In keeping with previous work demonstrating that the cylinder wake can be modeled accurately with as few as two degrees of freedom, we reduce the 9-mode POD-

Galerkin system to a 2-mode generalized Stuart-Landau model (7.24b) with the closure model (7.25) and damping derived from the energy-balance relation (7.26). In this case, the elimination of the higher-order modes does not require sacrificing kinematic resolution, since the unresolved coefficients can be approximated with sparse polynomial functions of the resolved variables, as detailed in Chapter 6.

Figure 7.3a compares the energy of the first mode pair $E(t) = a_1^2(t) + a_2^2(t)$ predicted by the POD-Galerkin model and the 2-mode closure model to DNS of the transient flow initialized from the unstable steady state perturbed by the least-stable eigenmode. Also included for comparison in the invariant manifold model introduced by Noack *et al.* (2003) and described in Sec. 5.4. Although the Galerkin model is stable and accurately predicts the amplitude of the vortex shedding, it exhibits a transient energy overshoot before settling onto the limit cycle. Figure 7.3b shows the reason for this; whereas the DNS is restricted to an approximately parabolic invariant manifold, the finite relaxation time in the shift mode dynamics means that the energy of the leading modes continues its exponential growth for longer before the energy can be absorbed into the mean flow deformation.

In contrast, the cubic term in the MMR and invariant manifold models approximates the relaxation of the shift mode dynamics as instantaneous, pinning its amplitude to that of the primary mode pair. While the structure and interpretation of the two models is the same, there are quantitative differences in their accuracy. In particular, the manifold model slightly overestimates the energy of the limit cycle, with mode amplitudes larger than the DNS by $\sim 13\%$, while the multiscale closure accurately predicts the limit cycle amplitude. This discrepancy can most likely be attributed to the energy balance procedure used to estimate the subscale damping ν in the MMR model, which is not used in the invariant manifold model.

Finally, the full flow field can also be reconstructed from the two degree-of-freedom multiscale model by means of the sparse polynomial approximation to the nonlinear correlations in the higher-order modes (Fig. 7.3c). The flow field predicted by the model gives a good approximation of the DNS solution, although over long times the phase of the vortex shedding tends to drift from that of the high-dimensional solution. By systematically averaging over the fast variables, the multiscale closure method is thus able to reduce the dimensionality of the model while also improving its fidelity.

7.4 Convective instability and vortex pairing in the mixing layer

As a more difficult test of the proposed MMR closure method, we also examine models of an incompressible mixing layer, pictured in Fig. 7.4. The mixing layer is a canonical example of a free shear flow exhibiting convective instability (Monkewitz & Huerre, 1982; Huerre & Monkewitz, 1985, 1990). Inlet forcing excited Kelvin-Helmholtz instability waves, which roll up into approximately discrete vortices. These vortices are subject to a secondary instability (Brancher & Chomaz, 1997) and undergo successive helical pairing at the subharmonic of the primary Kelvin-Helmholtz frequency, a process that drives




Figure 7.4: An incompressible mixing layer at $Re = 500$. The hyperbolic tangent base flow is convectively unstable, amplifying small perturbations as they are carried downstream. The primary shear layer instability forms Kelvin-Helmholtz waves that roll up into vortices, which in turn undergo successive vortex pairing. We force the flow at the inlet with eigenfunctions of inviscid flow equations linearized around the base flow (left). The central part of the domain is used for further modeling to avoid boundary effects, while the downstream extent $x \in (250, 300)$ is strongly damped to prevent numerical instability at the outlet.

the linear growth of the mixing layer (Winant & Browand, 1974). At higher Reynolds numbers and in three dimensions, the turbulent mixing layer is dominated by the linear growth of coherent structures similar to those seen in 2D simulation (Brown & Roshko, 1974; Stanley & Sarkar, 1997), which play a major role in mixing, transport, and entrainment of the turbulent flow (Hussain, 1981). From a global perspective, the flow behaves as an amplifier, where strong non-normality in the linear operator results in transient algebraic energy growth of small perturbations (Chomaz, 2005).

By analogy to numerical methods, modeling an advection-dominated flow with a POD-Galerkin approach is generally ill-advised in much the same way that pseudospectral methods are not ideal for solving hyperbolic problems. Although the vortices shed in the flow past a cylinder are also traveling waves in a sense, the Kelvin-Helmholtz instability in the mixing layer is convective in nature as opposed to the absolute instability in the cylinder wake. For these reasons, Galerkin projection onto a global modal basis such as POD is not a natural way to approximate shear flows like the mixing layer. Nevertheless, both free and wall-bounded shear layers are present in a wide variety of flows, so in order to employ a low-dimensional dynamical systems model it is important that it be able to capture these features. The challenge of modeling advection-driven phenomena in a global reduced-order modeling framework motivates the use of similar configurations in several studies of stabilization methods for Galerkin-type systems (Ukeiley *et al.*, 2001; Noack *et al.*, 2005; Cordier *et al.*, 2013; Balajewicz *et al.*, 2013).

Since the flow becomes increasingly complex downstream, the difficulty of modeling the flow can be tuned by varying the streamwise extent of the domain. We construct models on both short ($L_x = 150$) and long domains ($L_x = 250$) as shown in Fig. 7.5. Over the extent of the short domain, the flow is mainly characterized by the initial instability and vortex rollup, with the earliest signs of vortex pairing at $x \gtrsim 130$. On the full domain, the vortex pairing accounts for a much larger fraction of the fluctuation kinetic energy, as shown by the POD analysis below.

Flow configuration The inlet profile is a standard hyperbolic tangent:

$$U(y) = \bar{U} + (\Delta U/2) \tanh 2y, \quad (7.46)$$

where $\bar{U} = (U_1 + U_2)/2 \equiv 1$ and $\Delta U = U_1 - U_2$ in terms of the free-stream velocities U_1 and U_2 above and below the layer, respectively, and the length scale is nondimensionalized by the initial vorticity thickness $\delta_\omega = \frac{\Delta U}{U'(0)}$. We define a Reynolds number $Re = \delta_\omega(0)\Delta U/\nu$ based on initial vorticity thickness, velocity difference ΔU across the layer, and kinematic viscosity ν and set $Re = 500$ in the simulation. The domain consists of 5200 eleventh-order spectral elements on $x, y \in (-10, 300) \times (-50, 50)$, but the full streamwise extent is masked to $x, y \in (0, 250) \times (-20, 20)$ for modeling in order to discount any boundary effects. The domain is shown in Fig. 7.4.

This configuration roughly follows that described by Balajewicz *et al.* (2013) and Cordier *et al.* (2013), with two key differences. First, we simulate incompressible flow, while these authors simulate isothermal compressible flow at low Mach numbers ($Ma \sim 0.1$). Consequently, we replace the “sponge” boundary conditions for far-field acoustic waves with a region of increased viscous damping corresponding to $Re = 50$ for $x > 250$ to avoid numerical instability at the outflow. Second, rather than disturbing the inlet with random solenoidal perturbations, we use eigenfunction forcing similar to that employed by Colonius *et al.* (1997). The flow acts an amplifier (Huerre & Monkewitz, 1985, 1990; Chomaz, 2005), so the downstream flow is highly sensitive to the inlet conditions; eigenfunction forcing accentuates the natural dynamics of the flow.

Since the dominant Kelvin-Helmholtz instability is inviscid (Michalke, 1964), we derive the forcing from a temporal stability analysis of the inviscid equations linearized about the parallel hyperbolic tangent inlet profile $\mathbf{u}(\mathbf{x}) = (U(y), 0)$. Defining a perturbation streamfunction $\psi(\mathbf{x}, t)$ with real wavenumber α and complex velocity $c = c_r + ic_i$:

$$\psi(\mathbf{x}, t) = \text{Re} \left\{ \hat{\psi}(y) e^{i\alpha(x-ct)} \right\}, \quad (7.47)$$

the evolution of the perturbation is given by the Rayleigh equation (Schmid & Henningson, 2001)

$$(U - c) \left(\frac{d^2}{dy^2} - \alpha^2 \right) \hat{\psi} - U'' \hat{\psi} = 0. \quad (7.48)$$

For each wavenumber α , this can be written as a generalized eigenvalue problem

$$\mathcal{A}(\alpha)\hat{\psi}_j(y; \alpha) = c^{(j)}\mathcal{B}(\alpha)\hat{\psi}_j(y; \alpha) \quad (7.49)$$

for the j^{th} velocity c and perturbation streamfunction $\hat{\psi}$, where we use the convention of sorting by decreasing $\text{Im}\{c^{(j)}\}$.

The base profile is odd, so the real phase velocity is the midline value $c_r = \bar{U}$ (Michalke, 1964). With $\bar{U} = 1$, perturbations oscillate with frequency $\omega = \alpha$ and growth rate $\sigma(\alpha) = \alpha c_i(\alpha)$. For further details on linear stability analysis of shear flows, see for example Huerre & Monkewitz (1985); Schmid & Henningson (2001); Chomaz (2005).

We choose a spectral discretization of Eq. (7.48) using a Hermite collocation scheme from a Python implementation of the `dmsuite` library (Weidemann & Reddy, 2000). We find the least stable wavenumber α^* by maximizing the growth rate σ (i.e. minimizing $-\alpha c_i^{(1)}$) using an implementation of the BFGS algorithm available in `scipy`. Since this departs from the standard shooting method, we validate this method against Michalke (1964), who uses the profile $U(y) = 0.5(1 + \tanh y)$ based on nondimensionalizing with the shear layer thickness $\delta_\omega/2$ rather than vorticity thickness. With that profile we find a least-stable wavenumber $\alpha^* = 0.4449$ and growth rate $\sigma = 0.0949$ compared to the reported results of $\alpha^* = 0.4446$ and $\sigma = 0.0949$. For the present base profile (7.46) we find $\alpha^* = 0.8912$ and $\sigma = 0.2129$.

Once the wavenumber α^* and eigenfunction $\hat{\psi}_1$ corresponding to the maximum growth rate is identified, the inlet velocity perturbation can be computed from

$$\hat{u}_1(y; \alpha^*) = \hat{\psi}'_1(y; \alpha^*), \quad \hat{v}_1(y; \alpha^*) = i\alpha^*\hat{\psi}_1(y; \alpha^*). \quad (7.50)$$

The optimal perturbations are shown in Fig. 7.4. Following Colonius *et al.* (1997), we also apply forcing at the first three subharmonics of the least-stable mode computed from solving (7.48) at $\alpha^*/2$, $\alpha^*/4$, and $\alpha^*/8$, rescaling each so so that the maximum value of \hat{u}_1 is $10^{-3}\Delta U$. We perturb the inlet profile by $\text{Re}\{\hat{u}_1(y, \alpha_n)e^{i\omega_n t}\}$ for $n = 1, 2, 4, 8$ and $\alpha_n = \omega_n = \alpha^*/n$ and use a simulation time step of $\Delta t = 0.00705 = 10^{-3} \times 2\pi/\alpha^*$, so that the sampling rate of the simulation is commensurate with the forcing frequency.

The DNS is run until a final time of $t = 2820$, corresponding to fifty periods of the lowest-frequency component of the eigenfunction forcing, saving at every tenth time step. The POD modes are then computed from the first 10% of these snapshots (4000 fields), with the remainder reserved for statistical comparison with the models. The domain truncation can be accomplished by simply setting elements of the weight matrix \mathbf{W} in Eq. (3.4) to zero for mesh locations with $x > L_x$ so that these locations do not contribute to the statistics in the correlation matrix. For the mixing layer we compute modes with both velocity and pressure fields as described in Chapter 3, since the pressure gradient term is not necessarily negligible in Galerkin models of free shear flows (Noack *et al.*, 2005).

Figure 7.5 shows the leading POD eigenvalues and modes for both domains. In both

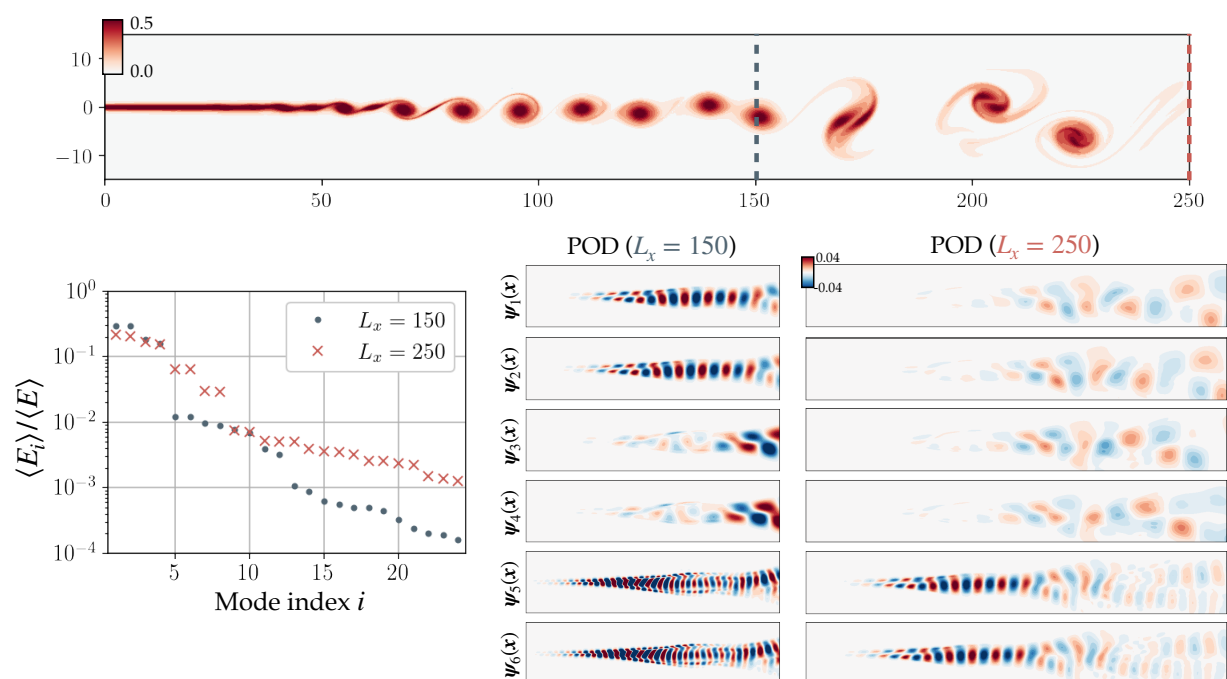


Figure 7.5: **Proper orthogonal decomposition (POD) applied to the mixing layer.** Modes computed on both the short and long domains reveal modes related to the dominant flow features: the shear layer instability and the downstream vortex pairing. Although the vortex pairing is secondary to the upstream instability, on the longer domain it accounts for a larger proportion of the fluctuation kinetic energy. Higher-order modes not pictured here include harmonics, nonlinear crosstalk, and modes related to irregularity in the location of the vortex pairing events.

cases the bulk of the fluctuation kinetic energy is contained in the first four modes, with the remainder of the spectrum decaying relatively slowly. As expected, the dominant mode pair for the shorter domain corresponds to the Kelvin-Helmholtz waves, while the second pair represents the onset of vortex shedding at roughly half the spatial wavenumber. Higher-order modes are either harmonics (such as the third mode pair) or less regular downstream structures related to slight variations in the vortex pairing.

The directed graph of energy transfers in Fig. 7.1 is constructed from the modes computed on the short domain. Specifically, the leading mode pair and its first three harmonics are shown for a straightforward visualization of the energy cascade; the nodes labeled $a_1 - a_8$ thus actually correspond to mode pairs (φ_1, φ_2) , (φ_5, φ_6) , $(\varphi_{17}, \varphi_{18})$, and $(\varphi_{24}, \varphi_{25})$. The thickness of the edge from node a_i to a_j is proportional to $\sum_k Q_{ijk} \overline{a_i a_j a_k}$ for quadratic Galerkin tensor \mathbf{Q} , which is roughly related to average nonlinear energy transfer from one mode to the other, although it should not necessarily be interpreted in any rigorous way beyond the purposes of conceptual visualization of the energy cascade.

On the longer domain the vortex pairing makes up a larger proportion of the fluctuation kinetic energy. Consequently, the leading four modes are spatially localized downstream, representing vortex pairing, while the upstream linear instability is not evident until the third mode pair, even though the vortex pairing is secondary to the shear layer instability from a physical perspective.

The slowly decaying POD spectrum and wavelike spatial structure of the modes are a consequence of the advection-driven nature of this flow. It is well known that the space-time separation of variables assumed by POD is not a natural representation of traveling waves. This is one reason that free shear flows have long posed a challenge for POD-Galerkin modeling.

Reduced-order model of the mixing layer In one sense, the dynamics of the mixing layer are relatively simple. The flow close to the inlet is dominated by the shear layer instability and subsequent vortex roll-up. These vortices are advected downstream at approximately the midline flow velocity, until the secondary vortex pairing begins at the subharmonic of the primary instability frequency. However, since this behavior unfolds as the disturbances are carried downstream, a principal challenge for low-dimensional models based on global POD modes is to preserve phase coherence between the harmonics and subharmonics, so that the reconstructed flow fields consist of approximately discrete vortices undergoing the vortex pairing.

As discussed above, both the POD modes and the difficulty of modeling this flow therefore depends on the streamwise extent of the domain. Since the subharmonic vortex pairings are energetically dominant downstream, the leading POD modes computed from a longer domain will be related to vortex pairing, even though the primary shear layer instability drives the dynamics from a causal perspective. To illustrate this effect we consider two domain lengths: one that is dominated the primary instability and vortex roll-up, and the full modeling domain with vortex pairing.




Figure 7.6: **Standard reduced-order models of the mixing layer.** The plots show the fluctuating kinetic energy in the first 16 modes of POD-Galerkin (top) and FTT closure models (bottom) for the mixing layer on the longer domain. Although the FTT closure model does stabilize the energy at a much lower level than the unmodified Galerkin system, it does not reproduce the natural dynamics of the flow over long time horizons.

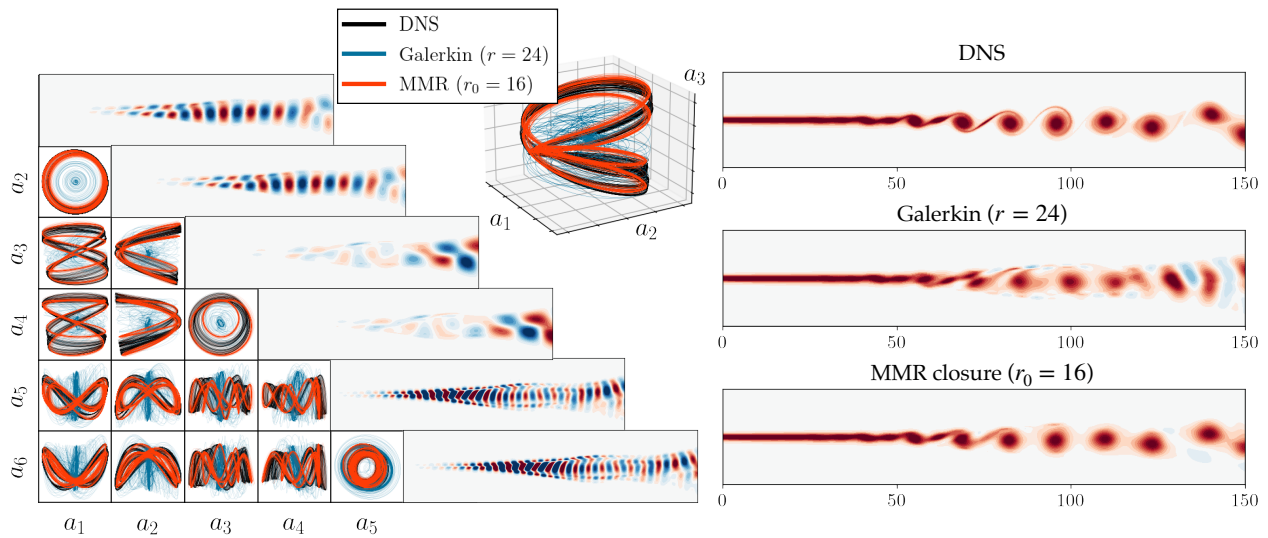


Figure 7.7: **Models of the mixing layer on a short domain.** The phase portraits of the mode pairs (left) show phase-locking between the fundamental mode pair (a_1, a_2) , the subharmonic vortex pairing (a_3, a_4) , and the first harmonic (a_5, a_6) . These phase relationships are preserved by the MMR closure, resulting in physically consistent flow field estimates (right).

In this case the POD-Galerkin models do not necessarily improve with rank truncation; for both domains the best dynamical approximation appears to be models constructed from $r = 24$ modes (Fig. 7.6, including comparison with FTT closure model). On the short domain ($L_x = 150$), this model is energetically stable but eventually approaches an unphysical fixed point (Fig. 7.7), while for the longer domain ($L_x = 250$) it initially follows the DNS trajectory before eventually becoming unstable (Fig. 7.8). In either case, increasing the dimension of the POD-Galerkin model resulted in worse models. We then reduce the 24-dimensional Galerkin systems to $r_0 = 16$ -dimensional generalized Stuart-Landau models with the MMR closure.

Figure 7.7 shows the comparison of the closed model to the DNS and Galerkin systems for the short domain. The MMR model closely matches the phase portraits of the DNS, indicating that it preserves the phase relationships between the various modes, including the early subharmonic vortex pairing (a_3, a_4) and higher harmonics, e.g. (a_5, a_6) . This remains true even when integrated to very long times (the final integration time is $t = 2820$ in this case, as with the DNS). As a result, the approximated fields are highly coherent, with the same pattern of vortex roll-up, advection, and pairing as seen in the DNS (Fig. 7.7, right).

On the longer domain, the multiscale closure also stabilizes the Galerkin model Figure 7.8 compares the true energy content of the POD coefficients to that predicted by the model both in the time and frequency domain. Although the evolution of the en-

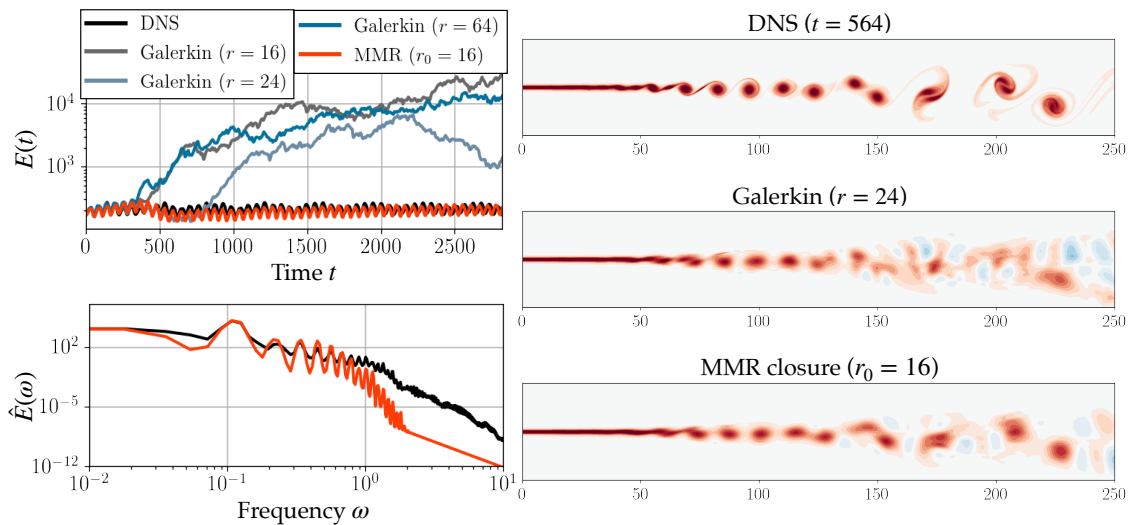


Figure 7.8: Models of the mixing layer on the long domain. While standard POD-Galerkin models are energetically unstable until at least $r = 64$, the multiscale closure stabilizes the model with only 16 variables (left). The model also remains coherent on long integration times (see also Fig. 7.9), producing flow field predictions that are consistent with the large-scale structure of the DNS (right).

ergy matches at the dominant frequencies, the MMR model does not capture the weakly chaotic dynamics exhibited by the DNS as a result of slight irregularity in the vortex pairing. Instead, after a brief transient the solution becomes approximately periodic with phase locking between the modes (Fig. 7.9). As with the short domain, this synchronization ensures that the predicted flow fields remain coherent, even when integrated for long times (Fig. 7.8, right).

As might be anticipated based on the more complex dynamics on the longer domain, the fields predicted by the MMR model do not match the DNS quite as closely as on the shorter domain, particularly closer to the inlet. This is likely because fine-scale resolution of the upstream region of the flow requires higher harmonics of the primary instability that do not appear in the POD modes of the longer domain until much higher-order modes. Higher-resolution predictions might therefore be constructed by applying nonlinear correlations analysis and polynomial regression to selected higher-order modes, as in the cylinder reconstruction in Chapter 6. Still, the dominant vortices are still evident in the prediction of the low-dimensional model, even in the downstream pairing region.

On both domains, the synchronization and long-time phase coherence between the modes reinforces the picture of the closure model as consisting of coupled nonlinear oscillators. Although the description of this flow in terms of spatially fixed global modes is not an ideal representation of the advection-driven dynamics, the proposed MMR closure method is able to successfully stabilize the POD-Galerkin models and accurately capture

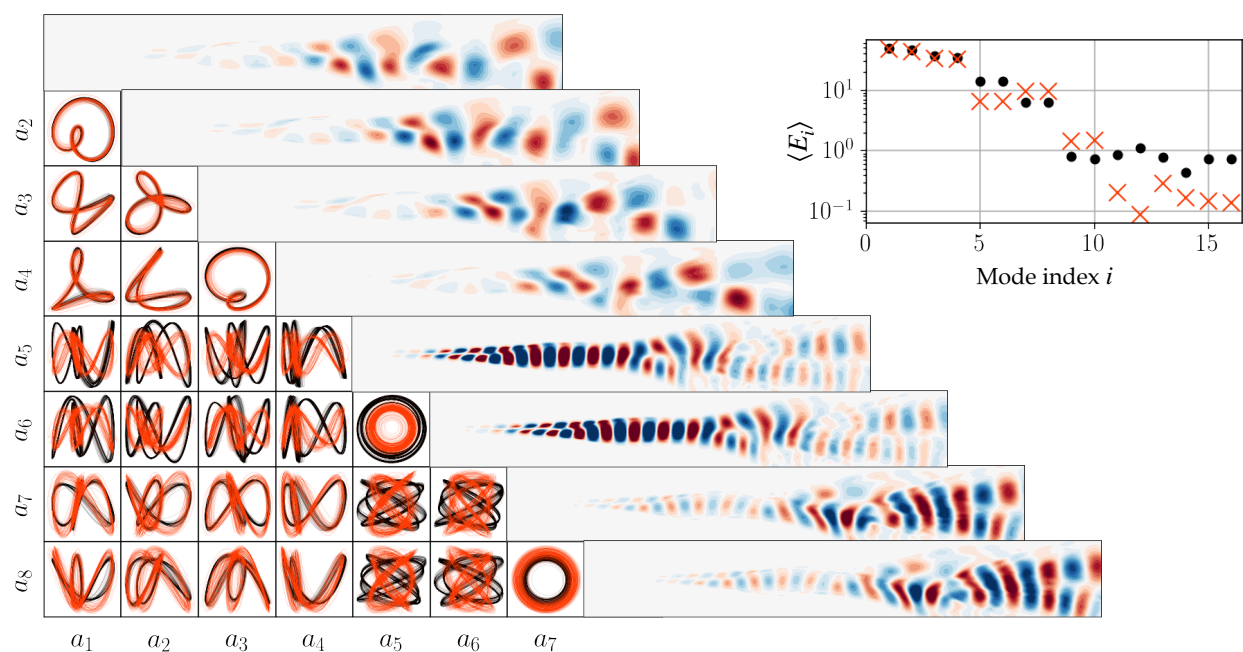


Figure 7.9: Phase portraits of the mixing layer model on the long domain. The MMR closure (orange) closely matches the DNS trajectories (black) for the most energetic modes, even though the structure is more complicated than the Lissajous-type harmonics on the short domain (Fig. 7.7). This figure does not show a Galerkin model for comparison because they are energetically unstable (Fig. 7.8).

the phase relationships necessary for accurate predictions of the full flow fields.

7.5 Discussion

In this chapter, we have developed a multiscale model reduction (MMR) approach to improve linear-quadratic dynamical systems, derived from POD-Galerkin projection of the Navier-Stokes equations, with the addition of systematically computed cubic closure terms. These cubic closure terms are derived through an adaptation and application of a multiscale stochastic averaging method that originated in singular perturbation theories of Markov processes (Kurtz, 1973; Papanicolaou, 1976; Majda *et al.*, 2001; E, 2011; Pavliotis & Stuart, 2012). Whereas the standard truncation of the Galerkin system disregards the influence of unresolved variables, the proposed multiscale model reduction method, accounts for their effect in an average sense through averaging via the stochastic Koopman operator. In particular, this approach is able to model nonlinear interactions between resolved and unresolved variables, capturing key mechanisms such as mean flow deformation and the energy cascade. The closed model includes cubic terms, taking the form of generalized Stuart-Landau equations that often act as coupled nonlinear oscillators.

Since the derivation of the MMR closure method relies on an asymptotic timescale separation between resolved and subscale variables that is generally absent in fluid flows, we do not attempt to prove the general validity of this approach when the flow is not close to a bifurcation. Instead, we have demonstrated by example that the method dramatically improves the stability and accuracy of low-dimensional models of unsteady fluid flow compared to the standard POD-Galerkin models. After a comparison to the well-known benchmark problem of vortex shedding in a cylinder wake, we have shown that the proposed method can reproduce chaotic dynamics in lid-driven cavity flow. Finally, we developed models of an incompressible mixing layer and showed that the coupled-oscillator MMR closures not only stabilize the models, but preserve phase relationships between the modes, which is critical for physically consistent predictions of advection-driven flow.

These results raise several possibilities for interesting future work. For example, the mixing layer is an “amplifier” flow characterized by convective instability and transient energy growth. It is therefore highly sensitive to the nature of the upstream flow; models developed for a specific inlet forcing may not generalize. Further consideration of this issue could be the foundation of a multiscale approach to input/output models in the resolvent framework (McKeon & Sharma, 2010; Sharma & McKeon, 2013; Zare *et al.*, 2017; Pickering *et al.*, 2021).

As discussed for the mixing layer in Sec. 7.4, we have observed that the MMR models often tend towards phased-locked periodic or quasiperiodic solutions rather than chaos, consistent with the idea of a system of weakly coupled oscillators. However, we expect that in the development of models of chaotic or turbulent flows it may be helpful to reintroduce some degree of stochasticity to better match the flow statistics. Fortunately,

this is straightforward in the multiscale modeling framework (E, 2011; Pavliotis & Stuart, 2012). For instance, the homogenization method (Majda *et al.*, 2001) begins from slightly different scaling assumptions and results in a closed model in which the fast-scale stochastic forcing appears as a new diffusion term. A careful investigation of different scale separation assumptions and stochastic approximations could be valuable in extending this framework to turbulent flows.

In this work we have only considered linear-quadratic dynamics derived from POD-Galerkin projection. This has the advantage of being based on the governing equations, but recent work has shown that data-driven model discovery can be a powerful alternative (Brunton *et al.*, 2016a; Loiseau & Brunton, 2018; Loiseau *et al.*, 2018a,b; Rubini *et al.*, 2020; Deng *et al.*, 2020; Callaham *et al.*, 2022). Data-driven modeling is especially useful when the governing equations are unknown or projection-based modeling is not as straightforward as for incompressible fluid flow (Loiseau, 2020; Guan *et al.*, 2021). In a sense, data-driven modeling might circumvent the need for closure modeling by fitting directly to the time series, including leveraging the intuition that cubic terms will appear in the effective dynamics (Loiseau & Brunton, 2018). However, the linear-quadratic system has certain symmetries and conservation properties that can be enforced in constrained regression to a quadratic model (Schlegel & Noack, 2015; Loiseau & Brunton, 2018; Kaptanoglu *et al.*, 2021a,b), while extending these properties to cubic nonlinearity is not straightforward. Another interesting future direction could be exploring a two-step process, first identifying a quadratic model that satisfies the constraints, and then applying the MMR closure to reduce it to a lower-dimensional cubic model. Since the empirical models are often sparse, this could potentially result in models with improved stability and scaling compared to the projection-based approach. Alternatively, a more flexible deep-learning model could be used to replace the proposed diagonal drift approximation, for instance with a modified variational autoencoder (Kingma & Welling, 2013) that parameterizes the distribution of subscale variables conditioned on the slow variables.

It will also be interesting to extend this analysis to systems where the leading-order cubic terms tend to *destabilize* the system, such as the subcritical bifurcation in plane Poiseuille flow (Stewartson & Stuart, 1971). This is an important scenario in the transition to turbulence, particularly for shear flows in which non-normal energy amplification can activate the nonlinearity even when the flow is globally stable. Higher-order effective nonlinearity could potentially be incorporated into the framework via state-dependent diffusion or an alternative fast drift model, for instance.

Finally, the analytic elimination of small scales in the Fourier-Galerkin model of Burgers' equation raises the possibility of a subscale LES closure model based on multiscale analysis. This might take the form of a multiple-scale expansion in both space and time or of a filter designed for the isotropic scales of turbulence, for instance. In either case, the Stuart-Landau form of the MMR closure model suggests that a spatiotemporal analysis would be akin to a Ginzburg-Landau theory of variational multiscale turbulence modeling.

Beyond the context of POD-Galerkin reduced-order modeling, the multiscale model reduction framework represents a general theory of effective cubic nonlinearity arising in amplitude equations for global modes in quadratically nonlinear fluid flow. The system of generalized Stuart-Landau equations produced by the closure model presents a picture of the flow as a set of coupled nonlinear oscillators describing the evolution of mode pairs. Close to the threshold of instability the MMR closure is consistent with a weakly nonlinear analysis, but it does not rely on proximity to a bifurcation. The multiscale modeling methodology is therefore an alternative to standard asymptotic expansions that may have implications for a variety of theoretical and numerical approaches to modeling the large, slow scales of unsteady fluid flows.

Part IV

STOCHASTIC SYSTEM IDENTIFICATION

Chapter 8

STOCHASTIC MODELING WITH LANGEVIN REGRESSION

It is widely accepted in physics that nominally deterministic systems with many degrees of freedom can often be modeled more effectively from a statistical point of view. In many complex multiscale systems, a variety of processes lead to emergent large-scale structures whose dynamics are described by a relatively small set of macroscopic variables (Haken, 1983; Cross & Hohenberg, 1993). The influence of the unresolved degrees of freedom can be approximated with random forcing in the spirit of statistical mechanics (Risken, 1996; Friedrich *et al.*, 2011; Tabar, 2019). This stochastic treatment of unresolved variables has become commonplace in fields including climate science (Majda & Wang, 2006), ecology (Levin, 1992), epidemiology (Allen, 2008), protein folding (Prinz *et al.*, 2011), neuroscience (Harrison *et al.*, 2005), and turbulence (Kraichnan & Chen, 1989).

The stochastic evolution of a state x is often represented with Langevin dynamics

$$\dot{x} = f(x) + \sigma(x)w(t). \quad (8.1)$$

The deterministic “drift” dynamics $f(x)$ describe the evolution of the slow macroscopic variables, while the fluctuations are parameterized by the diffusion term $\sigma(x)w(t)$, where $w(t)$ is typically assumed to be a Gaussian white noise process. If this model cannot be derived from first principles, a model can sometimes be inferred from observations of the natural dynamics of the system, as illustrated in Fig. 8.1

A significant challenge in constructing approximate stochastic models is that many of these systems are far enough from thermal equilibrium that even when the microscopic governing equations are known, fundamental principles such as detailed balance and the fluctuation-dissipation theorem cannot be readily applied. For example, widely separated time scales for forcing and dissipation prevents viscous turbulence from approaching a state of equipartition (Kraichnan & Chen, 1989; Noack *et al.*, 2008). This scale separation is captured by the Reynolds number, which can be interpreted as a ratio of the time scales characterizing the energetic large-scale dynamics and the small-scale viscous motions (Tennekes & Lumley, 1972).

As with equilibrium statistical physics and quantum mechanics, there are several ways to represent stochastic dynamics, as exemplified by the differing treatments of Brownian motion by Einstein (1905) and Langevin (1908). Einstein’s theory is constructed around a diffusion equation governing the evolution of the distribution of particles, while Langevin’s describes an individual trajectory of a particle subject to friction and a random fluctuating force. This duality persists in the modern theory; the same stochastic

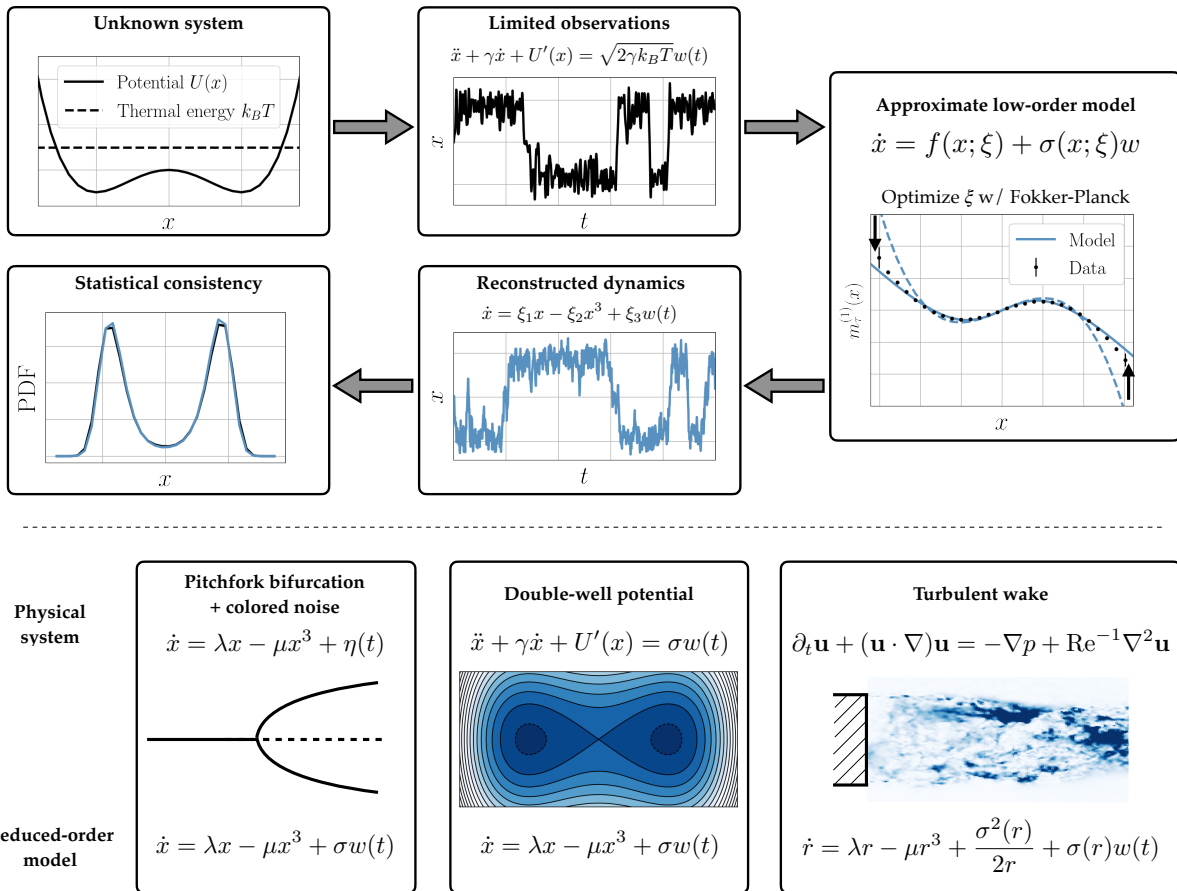


Figure 8.1: Schematic of Langevin regression (top) with example applications (bottom). Given a long time series of a macroscopic variable describing a complex system, we seek to identify an approximate stochastic model. The variable x might represent a reaction coordinate capturing metastable protein configurations or the temporal coefficient of a dominant global hydrodynamic mode, for instance. Langevin regression uses both the forward and adjoint Fokker-Planck operators to optimize free parameters ξ of the model, ensuring consistency with observed statistics such as the finite-time Kramers-Moyal coefficients $m_\tau^{(n)}(x)$ (see Sec. 8.1).

process can be represented with a generalized Langevin-type differential equation governing trajectories or a Fokker-Planck equation for the evolution of the probability distribution (Risken, 1996).

Linear dynamics can be analyzed from either perspective, but in the nonlinear case the most convenient representation often depends on the application. Nonlinear Langevin-type stochastic differential equations are difficult to treat analytically, but fit more naturally with low-dimensional modeling and control objectives. On the other hand, Fokker-Planck equations replace nonlinear trajectory dynamics with a linear partial differential equation for the probability distribution. The ensemble perspective also facilitates comparison with long time series measurements of ergodic systems.

In this work we seek to exploit these equivalent representations to identify Langevin dynamics by using both the forward and adjoint Fokker-Planck equations to ensure consistency with observations. We propose a framework for identifying nonlinear stochastic models from noisy experimental data, building on previous work in sparse regression (Brunton *et al.*, 2016a; Boninsegna *et al.*, 2018) and adjoint-based parameter estimation (Honisch & Friedrich, 2011). After presenting background material on stochastic dynamics in Sec. 8.1, we describe the method in Sec. 8.2. Finally, we demonstrate its application to three example systems in Sec. 8.3: a cubic model driven by colored noise, a particle in a double-well potential, and a symmetry-breaking instability in a turbulent wake.

Related work Stochastic systems can be broadly categorized according to the type of dynamics, noise, and stochasticity. Dynamics may be linear or nonlinear, the noise process may be white (uncorrelated in time) or colored (time-correlated), and the strength of the fluctuations may be constant (additive diffusion) or state-dependent (multiplicative diffusion). Similarly, stochastic model identification methods can be similarly categorized by the type of models they are able to construct.

For example, realization algorithms are a mainstay of engineering disciplines, although these methods are restricted to linear input/output systems with additive white noise (Juang & Pappa, 1985; Juang *et al.*, 1991; Van Overschie & De Moor, 1994). Perhaps the most general and successful nonlinear approach is the NARMAX framework, which can construct nonlinear models driven by state-dependent colored noise (Billings, 2013). However, NARMAX models typically cannot be transformed to continuous time, which is often the most natural setting for physical problems, making them difficult to interpret. More recent work has explored a variety of strategies for modeling nonlinear stochastic systems, including operator theoretic methods (Klus *et al.*, 2020), optimal transport (Arbabi & Sapsis, 2019), deep learning (Want *et al.*, 2019; Noé *et al.*, 2019, 2020), and identifying distribution evolution equations (Bakarji & Tartakovsky, 2020), although none of these pursues a representation in terms of nonlinear state-space dynamics. On the other hand, recent work has demonstrated that a stable linear system driven by colored noise can accurately reproduce second-order turbulent statistics (Zare *et al.*, 2017).

Recent advances have made significant inroads towards continuous-time model discovery in deterministic nonlinear systems (Kantz & Schreiber, 2003; Schmidt & Lipson, 2009). For example, sparse identification of nonlinear dynamics (SINDy) approximates time derivatives with a sparse linear combination of candidate functions (Brunton *et al.*, 2016a). However, even without the difficulties of estimating time derivatives from noisy data, a major challenge for extending deterministic methods to stochastic modeling lies in disambiguating the macroscopic dynamics from the unresolved degrees of freedom. In cases where the dynamics can be closely approximated by one-dimensional dynamics forced by additive white noise, parameters may be identified by regression to analytic solutions of the PDF (Rigas *et al.*, 2015; Brackston *et al.*, 2016). For more general systems, recent work has approached the parameter estimation problem with inference methods based on ensemble Kalman filtering Schneider *et al.* (2020) and information theory (Frishman & Ronceray, 2020; Brückner *et al.*, 2020). Alternatively, Boninsegna *et al.* introduced a major contribution to stochastic modeling by demonstrating that SINDy could be extended to stochastic systems without Monte Carlo approximation via the conditional moments used in the Kramers-Moyal expansion (Boninsegna *et al.*, 2018). This stochastic SINDy method was capable of recovering the correct model structure and parameters from large libraries of candidate functions.

Approximating stochastic dynamics from data with the Kramers-Moyal average has a long history of successful modeling in a wide range of fields (Friedrich *et al.*, 2011). However, as with many theoretical results, it is predicated on the assumption that the dynamics are driven by Gaussian white noise. As recognized by Einstein, even the molecular forcing involved in Brownian motion has some finite decorrelation time since it is a continuous physical system (Einstein, 1905). Moreover, omitting degrees of freedom from an otherwise Markovian¹ system generally leads to explicit memory effects in the dynamics (Zwanzig, 2001). We therefore expect that all systems will have some characteristic “Einstein-Markov” time scale over which the time evolution of macroscopic variables may depart significantly from the standard assumptions of stochastic modeling (Friedrich *et al.*, 2011). For Brownian motion this time scale is on the order of picoseconds, while for complex, multiscale, far-from-equilibrium systems it may even be longer than experimental sampling rates.

These considerations make the sampling rate used to construct stochastic models from experimental time series an important choice. Theoretical difficulties introduced by time-correlated forcing and non-Markovian effects can be avoided to some extent by deliberately subsampling. This has been established in finance (Zhang *et al.*, 2005; Aït-Sahalia *et al.*, 2005) and shown for artificial dynamics with two widely separated time scales (Pavliotis & Stuart, 2007). Qualitatively, coarse sampling allows the unresolved degrees of freedom to decorrelate, while ideally still resolving the coherent macroscopic scales (see Fig. 8.2). If the fluctuations appear uncorrelated in time, standard theoretical tools, such

¹In this context, meaning that the evolution of the system only depends on its current state and not its time history




Figure 8.2: Dual scale separation for stochastic modeling. Even the fastest scales of continuous physical systems are characterized by some finite decorrelation rate α (e.g. Eq. (8.7)). However, if the macroscopic dynamics have a much slower characteristic timescale $\omega \ll \alpha$, we may be able to choose a sampling rate τ^{-1} which can simultaneously resolve the dominant dynamics and allow the unresolved scales to decorrelate. For example, the power spectrum of the radial center of pressure of the turbulent wake is shown at bottom along with the subsampling rate used in Sec. 8.3.

as the Kramers-Moyal average, may once again be applied. However, coarse sampling leads to distorted estimates of the conditional moments used in the Kramers-Moyal expansion (Ragwitz & Kantz, 2001), although these finite-time effects can be accounted for using the adjoint Fokker-Planck equation (Lade, 2009).

The relevance of this result for parameter estimation in stochastic models was realized by Honisch & Friedrich, who proposed an optimization framework designed to correct for finite sampling-rate effects (Honisch & Friedrich, 2011). This technique has recently been refined and applied to parameter estimation for amplitude equations describing several different physical systems by Boujo & Noiray (2017); Boujo & Cadot (2019); Boujo *et al.* (2020). As with the majority of work in nonlinear time series analysis (Takens, 1981; Kantz & Schreiber, 2003; Bradley & Kantz, 2015), existing studies have focused on modeling scalar observables, although the theory readily generalizes to complex- or vector-valued systems.

Here we show that these finite-time corrections generalize the stochastic SINDy method (Boninsegna *et al.*, 2018) to the broad class of systems for which the forcing cannot be treated as white noise. Specifically, we explore systems for which the fast scales have nontrivial dynamics, the exclusion of which formally breaks the Markovian properties of the full physical system. This includes colored noise, latent variables, and “microscopic” degrees of freedom with significantly nonzero time correlations. Furthermore, we extend the adjoint Fokker-Planck optimization problem with the forward steady-state solution to enforce consistency between the model and the empirical probability distribution.

Contributions The proposed modeling framework, which we refer to as Langevin regression, is designed to identify nonlinear Langevin-type equations directly from noisy

experimental data. This method combines the advantages of three previously distinct approaches: adjoint-based parameter estimation with the Kramers-Moyal average (Lade, 2009; Honisch & Friedrich, 2011), learning unknown model structure with sparse regression (Brunton *et al.*, 2016a; Boninsegna *et al.*, 2018), and steady-state PDF fitting (Rigas *et al.*, 2015).

We explore the proposed nonlinear stochastic model identification method on several example systems. First, we illustrate the importance of judicious subsampling and finite-time corrections for correlated forcing by recovering a nonlinear Langevin equation driven by colored noise. We then show that the Langevin regression can construct a reduced-order model approximating the second-order dynamics of a particle in a double-well potential with a first-order bifurcation normal form. Both of these illustrative examples avoid the latent variable problem for the unresolved degrees of freedom by learning stochastic closure models. An implementation of Langevin regression along with code to reproduce the results from the simulated system is available on GitHub².

Finally, we apply Langevin regression to experimental measurements of a turbulent bluff-body wake; the sparse model selection procedure identifies a model similar to that proposed by Rigas *et al.* (2015), but with an additional nonlinear noise term that improves the correspondence with both the empirical probability distribution and power spectral density. This wake model will be expanded in Chapter 9, with a particular focus on exploring the spatial fields and modeling the mean flow deformation mechanism. Langevin regression draws from both the long legacy of stochastic modeling and recent advances in data-driven methods to form a flexible and general framework for approximating complex nonlinear dynamics with statistically consistent stochastic models.

8.1 Background on stochastic modeling

This section briefly reviews relevant theoretical concepts in stochastic modeling, including the Fokker-Planck equation, the Kramers-Moyal conditional average, and adjoint corrections for finite-time sampling effects. For more comprehensive background on the topics of nonequilibrium statistical mechanics and the physical applications of stochastic differential equations we refer the reader to reference texts such as Risken (1996); Zwanzig (2001); Tabar (2019).

In this work we seek to model the macroscopic variables $x \in \mathbb{R}^d$ with Langevin dynamics of the form given by Eq. (8.1). The influence of unresolved degrees of freedom on the standard deterministic dynamics $\dot{x} = f(x)$ is modeled with the diffusion term $\sigma(x)w(t)$, where $w(t)$ is a white noise process. The diffusion is called additive if $\sigma(x)$ is a constant, or multiplicative if it is state-dependent. The majority of stochastic models assume additive noise, since it is easier to treat analytically and is often a reasonable approximation for systems without strong coupling across scales. For instance, additive

²<https://github.com/dynamicslab/langevin-regression>

process noise is the standard assumption for linear state-space models of electronic circuits subject to thermal fluctuations. However, for systems such as fluid dynamics where the quadratic nonlinearity leads to bidirectional coupling between the coherent and turbulent degrees of freedom, state-dependent noise may improve the model (Majda *et al.*, 2001).

Perhaps the most restrictive assumption is that placed on the noise process $w(t)$. In order to treat Langevin dynamics analytically, the forcing is typically taken to be Gaussian-distributed and delta-correlated in time, i.e. $\langle w(t)w(t') \rangle = \delta(t - t')$. If this is not the case, the following discussion becomes more complicated (Risken, 1996; Zwanzig, 2001). For a macroscopic or integral quantity, the assumption of Gaussian statistics can be argued by appealing to the central limit theorem, but the time-correlation requirement is generally more difficult to justify.

Fokker-Planck equation The Langevin equation (8.1) describes individual trajectories, but due to the variability inherent in stochastic dynamics, it is often more natural to approach stochastic dynamics from the ensemble perspective. For the Langevin dynamics given by Eq. (8.1), conservation of probability requires that the probability density function (PDF) $p(x, t)$ evolves in time according to the *Fokker-Planck equation*³:

$$\frac{\partial p(x, t)}{\partial t} = -\frac{\partial}{\partial x_i} [f_i(x)p(x, t)] + \frac{\partial^2}{\partial x_i \partial x_j} [a_{ij}(x)p(x, t)] \equiv \mathcal{L}p, \quad (8.2)$$

where the diffusion tensor $a(x)$ is given by $a_{ij}(x) = \sigma_i(x)\sigma_j(x)/2$ and the subscripts indicate Einstein summation notation. See also Sec. 2.4 for a more abstract operator-theoretic discussion of the Fokker-Planck and related equations. As an alternative approach to Langevin equations, the multiscale interactions in complex spatiotemporal systems can be modeled explicitly with Fokker-Planck equations (Peinke *et al.*, 2019).

Often we are interested in the case where the system is statistically stationary, so that $\mathcal{L}p = 0$ and $p = p(x)$. For example, if x is a scalar and the diffusion σ is a constant, the steady-state solution can be determined analytically:

$$p(x) = C \exp \left[\frac{2}{\sigma^2} \int f(x) dx \right], \quad (8.3)$$

where the constant C is determined by the normalization condition $\int p(x) dx = 1$. However, solving the Fokker-Planck equation for general nonlinear dynamics is challenging and typically must be approached approximately or numerically (Risken, 1996).

Kramers-Moyal average The Fokker-Planck equation may also be derived by expressing the time evolution of a general PDF as a Taylor series of conditional finite-time mo-

³Here and throughout we use the Itô interpretation of stochastic integrals.




Figure 8.3: **Schematic of Kramers-Moyal coefficient estimation for the first moment (drift).** The drift estimate is determined by the conditional mean of the state evolution over the interval, while the diffusion is given by the conditional variance. The conditional moments can be approximated by dividing a long time series into histogram bins and taking the mean and variance within each bin. For example, the Kramers-Moyal drift estimate gives an approximate discretized vector field for the deterministic component of the dynamics (right).

ments $m_\tau^{(n)}(x)$, where the conditional finite-time moments $m_\tau^{(n)}(x)$ are

$$m_\tau^{(n)}(x) = \langle (x'(t + \tau) - x'(t))^n \rangle_{x'(t)=x}. \quad (8.4)$$

This leads to the *Kramers-Moyal expansion*. For scalar x ,

$$\frac{\partial p}{\partial t} = \lim_{\tau \rightarrow 0} \frac{1}{\tau} \sum_{n=1}^{\infty} \left(-\frac{\partial}{\partial x} \right)^n \frac{m_\tau^{(n)}(x)}{n!} p(x, t). \quad (8.5)$$

Viewing the conditional mean in Eq. (8.4) as a finite-difference formula, the first moment $m_\tau^{(1)}(x)$ gives the average displacement over a time interval τ if the system is in state x . Likewise, the second moment $m_\tau^{(2)}(x)$ gives a conditional variance of the short-time evolution.

However, according to Pawula's theorem, if the system is driven by Gaussian white noise all moments $n \geq 3$ vanish (Risken, 1996), leading to the Fokker-Planck equation (8.2). The leading moments are related to the drift and diffusion functions of the corresponding Langevin equation in the limit of vanishing time interval, i.e.

$$f_i(x) = \lim_{\tau \rightarrow 0} \frac{1}{\tau} \langle x'_i(t + \tau) - x'_i \rangle_{x'(t)=x} \quad (8.6a)$$

$$a_{ij}(x) = \lim_{\tau \rightarrow 0} \frac{1}{2\tau} \langle (x_i(t + \tau) - x_i)(x_j(t + \tau) - x_j) \rangle_{x'(t)=x}. \quad (8.6b)$$

We refer to these relationships as the *Kramers-Moyal average*.

In principle, these averages could offer a way to approximate unknown drift and diffusion functions from data by binning the time series into histograms and computing (8.6)

with the sampling rate $\tau = 1/f_s$, as illustrated in Fig. 8.3. This is a classic approach to constructing approximate Langevin equations from data (Friedrich & Peinke, 1997; Friedrich *et al.*, 2011). Boninsegna *et al.* also demonstrated that the Kramers-Moyal average can be combined with SINDy sparse regression to discover analytic drift and diffusion equations from data (Boninsegna *et al.*, 2018).

Finite-time effects The Kramers-Moyal average has chiefly been a theoretical tool; its application to time series analysis depends strongly on the assumption of Gaussian white noise and fast enough sampling rates to approximate $\tau \rightarrow 0$. In practice, these two requirements tend to be in tension; if the forcing originates with unresolved scales, then increasing the sampling rate leads to stronger correlations in the “noise”. The assumption of uncorrelated forcing is never strictly satisfied for continuous physical systems. Einstein recognized this in his work on Brownian motion (Einstein, 1905), although in that case the separation of scales is pronounced enough that experimental sampling rates run little risk of capturing correlation effects in the molecular forcing. In complex systems of modern interest, the scale separation is typically much less obvious.

As a simple illustrative model, consider a system driven by a colored noise process $\eta(t)$:

$$\dot{x} = f(x) + \sigma_x \eta \quad (8.7a)$$

$$\dot{\eta} = -\alpha \eta + \sigma_\eta w(t), \quad (8.7b)$$

where $w(t)$ is a true Gaussian white noise process. In this case, the forcing $\eta(t)$ is characterized by a decorrelation time α^{-1} . When $f(x)$ is a stable Navier-Stokes operator linearized about a turbulent mean profile, colored noise forcing has been shown to accurately reproduce turbulent statistics (Zare *et al.*, 2017). A generalized Langevin equation driven by time-correlated forcing can also be derived from the Euler equations using the direct-interaction approximation (Kraichnan & Chen, 1989). This perspective is also popular in climate modeling, where models of this form can be derived from the governing equations using perturbation arguments (Majda *et al.*, 2001).

Despite the formally non-Markovian nature of the resolved variables, it is advantageous to approximate the forcing as white-in-time, since most foundational theoretical tools are based on this assumption. One strategy to avoid dealing with the latent variable $\eta(t)$ is deliberate subsampling, as illustrated in Fig. 8.2. If the macroscopic dynamics have a characteristic timescale ω , we may be able to choose a sampling rate $f_s = \tau^{-1}$ such that $\omega \ll f_s \ll \alpha$ and the forcing appears decorrelated. That is, $\langle \eta(t + k\tau)\eta(t) \rangle_t \approx \delta_{k0}$. In this case the forcing appears to be (band-limited) white noise.

The minimum τ for which this is true is often called the Einstein-Markov scale (Friedrich *et al.*, 2011). This time scale is not necessarily related to the autocorrelation time of the macroscopic dynamics x , but may be identified from data by checking the Markov property at different sampling rates (Friedrich *et al.*, 2011; Callahan *et al.*, 2021b). This strat-




Figure 8.4: PDF evolution from the Fokker-Planck operator \mathcal{L} . The distribution used to evaluate the conditional finite-time moments (8.4) can be interpreted as the evolution of a delta function initial condition over the sampling interval τ , where the state is known to be x' at time t .

egy of subsampling at the approximate Einstein-Markov scale was proposed to model spatial fluctuations in turbulence (Friedrich *et al.*, 1998) and the cosmic microwave background (Ghasemi *et al.*, 2006).

In order to sample the resolved dynamics coarsely enough that the forcing is decorrelated, finite-time effects in the Kramers-Moyal estimates of drift and diffusion must be accounted for (Ragwitz & Kantz, 2001). However, these effects can be determined exactly in terms of the adjoint Fokker-Planck operator (Lade, 2009). For notational clarity, we give the result for scalar x , although the result generalizes naturally to higher dimensions.

The conditional moments defined in Eq. (8.4) can be written equivalently as

$$m_\tau^{(n)}(x) = \int_{-\infty}^{\infty} (x' - x)^n p(x', t + \tau | x, t) dx', \quad (8.8)$$

where $p(x', t + \tau | x, t)$ indicates the conditional joint probability that the system is in state x' at time $t + \tau$ given that it was in state x at time t . If the drift and diffusion are not time-dependent, then the conditional probability $p(x', t + \tau | x, t)$ can be interpreted as the propagation of uncertainty over an interval τ if the state x is known at time t , as shown in Fig. 8.4. According to Eq. (8.2), this evolution is given by the Fokker-Planck equation acting on a Dirac delta function:

$$p(x', t + \tau | x, t) = e^{\tau \mathcal{L}(x')} \delta(x' - x). \quad (8.9)$$

Using the definition of the adjoint operator and evaluating the integral with the delta function,

$$m_\tau^{(n)}(x) = \left[e^{\tau \mathcal{L}^\dagger(x')} (x' - x)^n \right]_{x'=x}, \quad (8.10)$$

where the adjoint Fokker-Planck operator is given in tensor summation notation by

$$\mathcal{L}^\dagger(x) = f_i(x) \frac{\partial}{\partial x_i} + a_{ij}(x) \frac{\partial^2}{\partial x_i \partial x_j}. \quad (8.11)$$

Thus, Lade (2009) showed that the effect of coarse sampling rates can be understood as the adjoint Fokker-Planck operator evolving the moments of the distribution in time.

Honisch and Friedrich demonstrated the use of this relationship to optimize free parameters in a Langevin model (Honisch & Friedrich, 2011); the proposed method in Sec. 8.2 builds on this result.

The close correspondence between the Fokker-Planck and Liouville operators also suggests an interpretation of this result in terms of Koopman theory (Mezić, 2005, 2013). The adjoint Fokker-Planck operator is a *linear* generator of the time evolution of observable functions, even when the underlying dynamics are nonlinear (Zwanzig, 2001; Klus *et al.*, 2020). In this case, the observables are the conditional moments $(x' - x)^n$.

These corrections are difficult to compute analytically for nonlinear dynamics, so Eq. (8.10) is typically approximated on a discretized domain. The matrix exponential $e^{\tau \mathcal{L}^\dagger}$ is relatively inexpensive in one dimension, but more generally it may be helpful to interpret the correction as the solution of a PDE. That is, if $w^{(n)}(x, t)$ is the solution to the adjoint Fokker-Planck equation

$$\frac{\partial w^{(n)}}{\partial t} = \mathcal{L}^\dagger(x)w^{(n)}(x, t), \quad w^{(n)}(x, 0) = x^n, \quad (8.12)$$

then by linearity of \mathcal{L}^\dagger the finite-time conditional moments $m_\tau^{(1)}(x)$, $m_\tau^{(2)}(x)$ are given by

$$m_\tau^{(1)}(x) = w^{(1)}(x, \tau) - x \quad (8.13a)$$

$$m_\tau^{(2)}(x) = w^{(2)}(x, \tau) - 2xw^{(1)}(x, \tau) - x^2. \quad (8.13b)$$

8.2 Proposed Langevin regression method

The decomposition of a multiscale system into dominant deterministic dynamics and stochastic forcing is conceptually simple. However, the gulf between detailed first-principles descriptions and a simple Langevin model is large enough that developing accurate stochastic models is difficult. This is especially true when the dynamics are nonlinear and the unresolved scales have internal dynamics and cannot be treated as true Gaussian white noise.

Due to the intrinsic volatility of individual trajectories, model identification methods typically rely on ergodicity and exploit the connection to ensemble properties via the Fokker-Planck equation. For example, a scalar model with additive noise can be determined by fitting to the analytic steady-state probability distribution given by Eq. (8.3). However, this fitting procedure cannot independently estimate drift and diffusion; an additional quantity, such as mean-square displacement, must also be used (Rigas *et al.*, 2015; Brackston *et al.*, 2016). This presents a challenge for modeling multiplicative noise, which some studies have suggested is important for capturing interactions between the resolved and unresolved degrees of freedom (Friedrich & Peinke, 1997; Majda *et al.*, 2001). Furthermore, the use of a time average to estimate steady-state statistics destroys temporal information, so that oscillatory dynamics cannot be resolved.

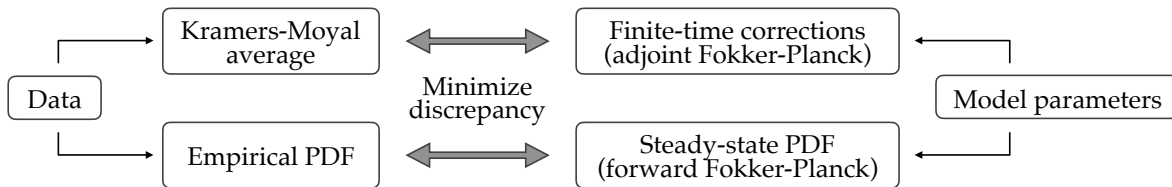


Figure 8.5: **Schematic of the Langevin regression optimization problem.** The Fokker-Planck equation can be used to compare the conditional moments and distribution for a proposed model to those observed empirically. The model parameters are chosen to minimize the discrepancy between the model and observations. The form of the model may be simultaneously identified with the model selection procedure outlined in Sec. 8.2

Seeking to address a number of these challenges, several recent studies have proposed methods whereby stochastic models may be derived from the Kramers-Moyal average (Honisch & Friedrich, 2011; Boujo & Noiray, 2017; Boninsegna *et al.*, 2018). Here we synthesize and extend these methods into a single optimization framework for identifying sparse, interpretable Langevin-type models from experimental data by constraining the model to both the stationary PDF and the Kramers-Moyal average. In particular, we generalize the stochastic SINDy method proposed by Boninsegna *et al.* (2018) with finite-time corrections that enable modeling systems whose forcing is not approximately given by Gaussian white noise.

Known model structure: Fokker-Planck optimization Estimating the Kramers-Moyal coefficients (8.6a) and (8.6b) by binning a long time series is an attractive option for estimating drift and diffusion functions directly from data. However, as discussed in Sec. 8.1, the sampling rate must be chosen carefully for unresolved dynamics to decorrelate. This subsampling can introduce significant finite-time distortion to the empirical Kramers-Moyal coefficients. This distortion is given in terms of the adjoint Fokker-Planck equation by Eq. (8.10), although constructing the Fokker-Planck operator itself requires the drift and diffusion.

Based on these considerations, Honisch & Friedrich (2011) suggested an iterative procedure to estimate free parameters of drift and diffusion functions, which may be summarized as follows. If the Langevin model is given in terms of a set of parameters ξ , so that

$$\dot{x} = f(x; \xi) + \sigma(x; \xi)w(t), \quad (8.14)$$

then the problem is to choose the parameters such that the finite-time conditional moments are consistent with the empirical Kramers-Moyal estimates. Leaving aside numerical details, the optimization consists of the following:

Parameter optimization

1. Select an appropriate sampling rate τ (see appendix of Callaham *et al.* (2021b) for some suggestions);
2. Estimate empirical finite-time conditional moments $\hat{m}_\tau^{(n)}(x)$ for $n = 1, 2$ using Eq. (8.4);
3. For a set of parameters ξ , construct the adjoint Fokker-Planck operator \mathcal{L}^\dagger in Eq. (8.11);
4. Compute the exact moments $m_\tau^{(n)}(x, \xi)$ using the adjoint correction given by Eq. (8.10);
5. Choose ξ to minimize the discrepancy between the empirical moments $\hat{m}_\tau^{(n)}(x)$ and the exact moments $m_\tau^{(n)}(x, \xi)$.

More concretely, the optimal ξ solves the following problem on a discrete domain of N points x_i :

$$\min_{\xi} \sum_{n=1}^2 \sum_{i=1}^N w_i^{(n)} [m_\tau^{(n)}(x_i, \xi) - \hat{m}_\tau^{(n)}(x_i)]^2. \quad (8.15)$$

Here the weights $w_i^{(n)}$ reflect pointwise uncertainty in the empirical estimate of the moments.

Due to the diffusive nature of the Fokker-Planck equation, it is not clear that this problem is necessarily well-posed. That is, when the system is sampled coarsely there may be a range of parameters that are consistent with the observed conditional moments within experimental uncertainty.

We propose “regularizing” the optimization problem (8.15), as proposed by Ref. (Honisch & Friedrich, 2011), with the Kullback-Leibler (KL) divergence \mathcal{D}_{KL} between the empirical PDF $\hat{p}(x)$ and the solution of the steady-state Fokker-Planck equation $p(x, \xi)$, given by Eq. (8.2). The modified cost function is

$$\min_{\xi} \sum_{n=1}^2 \sum_{i=1}^N w_i^{(n)} [m_\tau^{(n)}(x_i, \xi) - \hat{m}_\tau^{(n)}(x_i)]^2 + \eta \mathcal{D}_{\text{KL}}(\hat{p}(x), p(x, \xi)), \quad (8.16)$$

where η is the relative weight of the two contributions to the cost function. The KL divergence is a statistical measure of the difference between two probability distributions p and q , defined as

$$\mathcal{D}_{\text{KL}}(p, q) = \int p(x) \log \left(\frac{p(x)}{q(x)} \right) dx. \quad (8.17)$$

This regularization ensures that the resulting Langevin model is consistent with both the finite-time Kramers-Moyal coefficients and the asymptotic steady-state probability distribution. The optimization problem in Langevin regression is shown schematically in

Fig. 8.5. Because the typical dimension of ξ is relatively small compared to the dimension x , we find that gradient-free optimization methods such as a Nelder-Mead simplex search are more efficient than those designed for large parameter spaces and relatively inexpensive cost function evaluations, such as automatic differentiation.

The resulting optimization problem requires solution of both the forward and adjoint Fokker-Planck equations at each evaluation of the cost function. For a scalar variable x , the steady-state solution to the forward equation can be computed directly with Eq. (8.3). In higher dimensions, there is no analytic steady-state solution to the forward equation; a spectral numerical solver is described at the end of this section. We solve the adjoint equation with a second-order finite difference method. In one and two dimensions, the matrix exponential is a relatively inexpensive method of solving the initial-value problem, while for higher dimensions a time-stepping approach that exploits the sparse operator structure is more efficient.

Unknown model structure: stochastic SINDy If the form of the model can be assumed up to a set of unknown parameters, the Fokker-Planck optimization problem detailed in the previous section is sufficient to estimate the free parameters. However, in many cases we might have some partial prior assumptions about the model structure (e.g. the model consists of polynomials with a particular symmetry, or that one variable is forced by another), but the exact form is unknown. The sparse identification for nonlinear dynamics (SINDy) method has recently shown promise for obtaining nonlinear reduced-order models of laminar flows (Brunton *et al.*, 2016a; Loiseau & Brunton, 2018; Loiseau *et al.*, 2018b; Loiseau, 2020; Deng *et al.*, 2020) from data. However, SINDy typically relies on estimated time derivatives, which is a significant barrier to modeling experimental data or multi-scale systems. In related work, SINDy has recently been leveraged for turbulence closure modeling (Beetham & Capeceletro, 2020).

Boninsegna *et al.* (2018) recently proposed a stochastic SINDy algorithm based on the Kramer-Moyal average without an adjoint correction for finite-time effects. Empirical estimation of conditional moments gives point estimates of drift and diffusion at each histogram bin; if the moments are estimated reliably, then the system identification problem reduces to fitting a curve through these points. In its simplest form, a parsimonious model can be chosen using the SINDy framework, where the model parameters are a coefficient vector for a “library” matrix whose columns consist of candidate functions (Brunton *et al.*, 2016a). For instance, the library $\Theta(x)$ might consist of polynomials in x ; then we look for polynomial representations of the drift $f(x)$ and diffusion $\sigma(x)$, so that

$$f(x) = \Theta_f^T(x)\xi_f \quad \sigma(x) = \Theta_\sigma^T(x)\xi_\sigma, \quad (8.18)$$

where ξ_f and ξ_σ are sparse vectors that have as many zero entries as possible while still capturing the observed dynamics. Standard sparse regression algorithms can be used to select a set of functions balancing parsimony and accuracy.

In the deterministic SINDy algorithm, this regression problem is constructed by concatenating column vectors of an estimated time derivative \dot{x} and the evaluations of the candidate functions. In the present case, however, the regression is performed over the discretized spatial domain rather than a long time series. The conditional finite-time coefficients are estimated by computing Eq. (8.4) over observations that fall into each spatial histogram bin, as visualized in Fig. 8.3. The regression problem is constructed over these bins rather than the direct time series. In practice this typically reduces the length of the column vectors from $\mathcal{O}(10^6)$ to $\mathcal{O}(10^2)$.

One consequence of this formulation is that standard sparse regression algorithms, such as thresholded least squares or forward regression orthogonal least squares, do not work well. A simple alternative is the reverse-greedy stepwise sparse regression (SSR) (Boninsegna *et al.*, 2018). With this method, terms are sequentially removed from the model according to some criteria. In the original SSR algorithm, the coefficients were identified with a simple least squares and terms with the smallest absolute value were removed. However, in general the smallest coefficient does not necessarily imply the least important contribution. For this reason we sequentially remove terms corresponding to the smallest *increase* in cost function. The cost function itself can then serve as a model-selection criterion; for a Pareto-optimal model the cost function should jump significantly from a near-minimum value once important terms begin to be discarded.

When combined with the forward/adjoint Fokker-Planck optimization described in the previous section, this model selection procedure represents a flexible and general framework for identifying stochastic approximations to multiscale nonlinear dynamics, which we refer to as *Langevin regression*. In the following section, the Fokker-Planck parameter estimation is demonstrated on two example systems for which the form of the model is clear. The ability of Langevin regression to simultaneously identify the structure and parameters of a model from data is demonstrated in Sec. 8.3 for experimental measurements of a turbulent wake.

Steady-state solvers for the Fokker-Planck equation Since the optimization problem requires repeated evaluation of the stationary PDF, this section describes an efficient spectral solver for the multi-dimensional steady-state Fokker-Planck equation. The method is based on the Hermite-Galerkin method described by Soize (1988), but in our case the empirical PDF is approximated by a histogram on a uniform grid, for which a function approximation based on Hermite polynomials will suffer from the Runge phenomenon. Instead we propose approximating the PDF with a Fourier series; alternatively, a .

For $x \in \mathbb{R}^d$, the steady-state Fokker-Planck equation is

$$0 = -\frac{\partial}{\partial x_n} f_n(x)p(x) + \frac{\partial^2}{x_n x_m} a_{nm}(x)p(x), \quad (8.19)$$

where the indices imply summation. We use the Fourier representation of the PDF:

$$p(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{p}(k) e^{ikx} dk, \quad \hat{p}(k) = \int_{-\infty}^{\infty} p(x) e^{-ikx} dx. \quad (8.20)$$

The normalization condition implies $\hat{p}(0) = 1$.

In practice, the Fourier representation will be truncated at a finite number of modes. After substituting the approximate representation of $p(x)$ we find the residual

$$R = \frac{1}{2\pi} \int \hat{p}(k) dk \left[-\frac{\partial}{\partial x} (f(x) e^{ikx}) + \frac{\partial^2}{\partial x^2} (a(x) e^{ikx}) \right]. \quad (8.21)$$

Minimizing the residual requires that it is orthogonal to the subspace spanned by the Fourier modes. Project onto an arbitrary wavenumber k' and simplifying:

$$0 = \int \hat{p}(k) dk \left[-ik' \hat{f}(k' - k) - k'^2 \hat{a}(k' - k) \right]. \quad (8.22)$$

This equation must be true for all k' . Practically, the integral must be approximated as a sum over wavenumbers. Finally, the normalization condition can be applied to obtain an inhomogeneous linear equation for \hat{p} :

$$b(k') = ik'_n \hat{f}_n(k') + k'_n k'_m \hat{a}_{nm}(k') \quad (8.23a)$$

$$A(k', k) = -ik'_n \hat{f}_n(k' - k) - k'_n k'_m \hat{a}_{nm}(k' - k) \quad (8.23b)$$

$$b(k') = \sum_{k \neq 0} A(k', k) \hat{p}(k). \quad (8.23c)$$

Constructing the Galerkin operators becomes increasingly difficult in higher dimensions; an alternative is to avoid explicitly constructing the Fokker-Planck operators and instead approximate the steady-state solution with an iterative method (Antoulas, 2005). Assuming the Fokker-Planck equation has a unique steady-state solution, this solution is the eigenvector of the Fokker-Planck operator \mathcal{L} with zero eigenvalue. This eigenvector can be approximated with matrix-free iterative methods, such as Arnoldi iteration. This approach only requires defining the action of the operator \mathcal{L} on a probability distribution p (a “matrix-vector” product). In this method we discretize the Fokker-Planck operator with a pseudospectral approach:

$$\mathcal{L}p = \sum_{n=0}^d \mathcal{F}_n^{-1} \left\{ -ik_n \mathcal{F}_n \{f_n(x)p(x)\} - k_n^2 \mathcal{F}_n \{a_n(x)p(x)\} \right\}, \quad (8.24)$$

where \mathcal{F}_n and \mathcal{F}_n^{-1} denote the forward and inverse Fourier transforms in the n -th dimension.

8.3 Results

In this section we apply Langevin regression to three example problems of increasing complexity. First, we show that the coarse sampling and scale separation ideas discussed in Sec. 8.1 enable the identification of stochastic systems with time-correlated forcing. For this example, we assume knowledge about the structure of the model in order to highlight the effects of colored noise and parameter estimation in the case where the correct structure is known.

Second, we demonstrate the construction of a statistically consistent reduced-order model by approximating the second-order dynamics of a particle in a double-well potential with the corresponding first-order bifurcation normal form. A stochastic normal form model can be derived analytically for this system, although its accuracy quickly degrades away from the bifurcation point. We fix the structure of the model and show that Langevin regression can maintain statistical accuracy even far from the bifurcation point.

Finally, we derive an accurate and efficient stochastic model for the turbulent flow in the wake of an axisymmetric bluff body from experimental measurements. This example presents several challenges, including partial and noisy measurements of a multiscale system, and has relevance to numerous industrial applications (Brunton & Noack, 2015a). Although the structure of the drift dynamics may be inferred from laminar stability analysis, we instead apply sparse model selection to discover this structure entirely from data. Our procedure identifies a simple and interpretable nonlinear model with a multiplicative noise term that improves the correspondence with the empirical power spectrum and probability distribution compared with previous stochastic modeling results.

The synthetic examples are simulated using the SRIW1 stochastic Runge-Kutta method (Rößler, 2010), available in the DifferentialEquations.jl package (Rackauckas & Qing, 2017). Langevin regression is performed on Kramers-Moyal coefficients and empirical PDFs computed from a time series of 10^7 points sampled at $\Delta t = 10^{-2}$. The turbulent wake model is based on the aerodynamic center of pressure, a global integral quantity estimated from 64 evenly spaced pressure taps (Rigas *et al.*, 2014). The model is estimated from a time series of 8.9×10^6 experimental measurements of the center of pressure sampled at 225 Hz. Monte Carlo evaluation of this model is performed in Python using a standard Euler-Maruyama numerical integration scheme at the same sampling rate. The coarse subsampling rate for Kramers-Moyal averaging is chosen for each system according to the criteria discussed in Callahan *et al.* (2021b). The relative weight η of the Kullback-Leibler divergence in the optimization function is a multiple of 10, which is chosen to be roughly equal to the minimum cost function with $\eta = 0$. In other words, the Kramers-Moyal coefficients and the PDF are given roughly equal weight in the optimization.

Pitchfork bifurcation normal form The normal form for a supercritical pitchfork bifurcation provides a canonical example of bistability, where an eigenvalue with zero imaginary part crosses the real axis as a parameter is varied. For example, this normal form de-

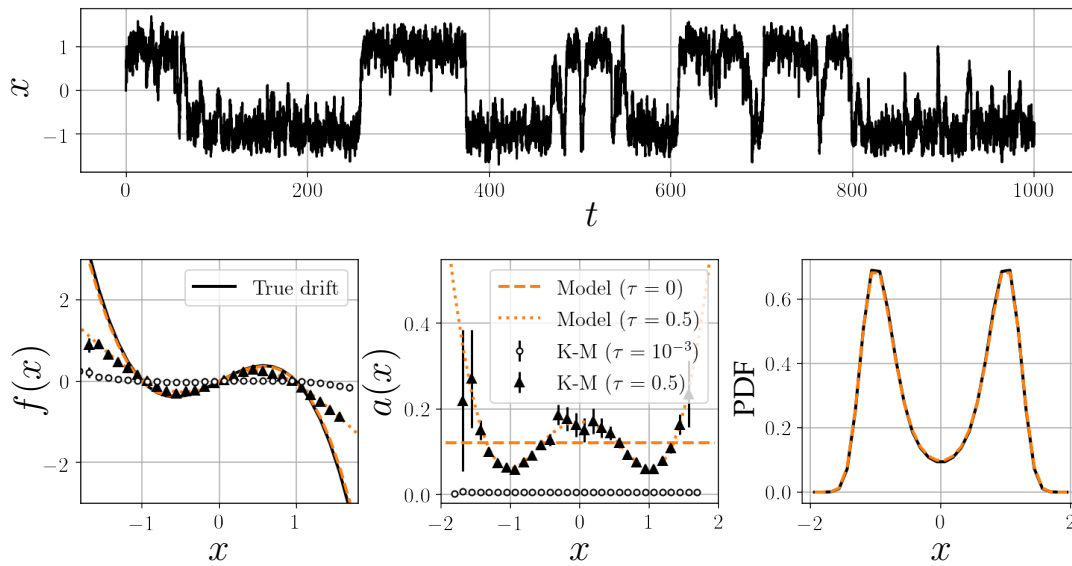


Figure 8.6: **Pitchfork normal form forced by time-correlated (colored) noise.** Top: the noise induces random switching between the metastable equilibria. Bottom: if the Kramers-Moyal coefficients are computed with a sampling rate faster than the decorrelation of the noise (\circ), the drift still appears cubic but the amplitude is underestimated by approximately an order of magnitude. On the other hand, if the noise is allowed to decorrelate (Δ), the estimated Kramers-Moyal coefficients are the right order of magnitude, but are distorted from the zero-time value. The diffusion appears multiplicative and quadratic. The adjoint finite-time corrections recover a consistent Langevin model driven by white noise ($--$).

scribes the amplitude equation governing the symmetry-breaking mode of the wake behind a circular disk, which can be derived by a weakly nonlinear stability analysis (Meliga *et al.*, 2009). This result inspired the use of a stochastically forced pitchfork normal form to model the turbulent evolution of the centroid of the base pressure distribution on the back of an axisymmetric bluff body (Rigas *et al.*, 2014) and the bistability of a three-dimensional Ahmed body wake (Brackston *et al.*, 2016).

However, unresolved degrees of freedom in a turbulent flow do not typically resemble delta-correlated white noise. Here we investigate the impact of correlated noise on stochastic system identification by considering the supercritical pitchfork normal form forced by colored noise:

$$\dot{x} = \lambda x - \mu x^3 + \eta \quad (8.25a)$$

$$\dot{\eta} = -\alpha \eta + \sigma w(t). \quad (8.25b)$$

Here η is an Ornstein-Uhlenbeck process with characteristic relaxation time α^{-1} , which acts as an effective low-pass filter on the white noise process $w(t)$. We choose $\mu = \lambda = \beta = 1$, $\alpha = 10^2$, and $\sigma = 0.5\alpha$ so that the typical amplitude of η is around 0.5.

The nonzero relaxation rate for the Ornstein-Uhlenbeck process introduces temporal correlations that invalidate many of the standard analytic approaches to stochastic modeling if η is not directly observed. However, if the relaxation rate is much larger than any natural dynamics of the slow variable (i.e. $\alpha \gg \lambda$) we can appeal to the dual scale separation idea of Fig. 8.2 and approximate the fast scales as uncorrelated noise.

As demonstrated in Fig. 8.6, the correlated forcing destroys the Kramers-Moyal average as the sampling interval $\tau \rightarrow 0$. This can be mitigated by sampling coarsely enough that the noise decorrelates and appears to whiten. For instance, if we choose $\tau = 0.5 = 50\alpha = 0.5\lambda$ (slower than the noise decorrelation but faster than the drift dynamics), the Kramers-Moyal average is of the correct order of magnitude. However, finite-time sampling rates now significantly deform the observed drift and diffusion, even introducing apparent state dependence in the diffusion (Fig. 8.6, bottom middle). This observation by Ragwitz & Kantz called into question many earlier attempts to use the Kramers-Moyal average for modeling without accounting for finite-time effects (Ragwitz & Kantz, 2001).

Langevin regression accounts for these distortions with the adjoint Fokker-Planck operator, recovering a one-dimensional model nearly identical to Eq. (8.25a), but forced by additive white noise. The identified coefficients (given in Table 8.1) differ from the true values by around 5%, but the model closely matches both the observed finite-time conditional moments and the empirical probability distribution (Fig. 8.6, bottom row). This suggests that the proposed subsample-and-correct approach is capable of identifying statistically consistent Langevin models, even in the presence of correlated noise. This result depends fundamentally on the dual scale separation principle of Fig. 8.2; the success of this approach may be limited when these timescales cannot be clearly separated with the coarse sampling rate.

Model	λ	μ	σ
True system (colored noise)	1.0	1.0	--
Stochastic SINDy (no adjoint)	0.43	0.43	0.44
Langevin regression (white noise)	0.96	0.96	0.49

Table 8.1: True parameters for the pitchfork normal form forced by colored noise along with those estimated from data. Without the adjoint Fokker-Planck corrections for finite sampling rates, the true coefficients are underestimated by a factor of 2. On the other hand, Langevin regression with full adjoint-based optimization identifies a statistically consistent model driven by white noise forcing, where the drift coefficients are a close match to the true system.

Double-well potential In many cases, Langevin-type stochastic models are intended to be reduced-order approximations of the large-scale dynamics of a complex system, rather than faithful representations of first-principles physics. The “microscopic” degrees of freedom in these systems generally have finite correlation times, as with the colored noise in the previous example. Eliminating these variables from the model leads to explicit memory effects in the Langevin equations (Zwanzig, 2001), unless the scale separation principle can be employed to identify a memory-free reduced-order stochastic model from data. Low-order polynomial dynamics, such as normal forms, can arise naturally in this context as a way to describe the macroscopic behavior, even when the underlying physical description appears completely different (Guckenheimer & Holmes, 1983).

For example, the one-dimensional motion of a particle of unit mass in a general potential $U(x)$ subject to thermal fluctuations is given by

$$\ddot{x} + \gamma\dot{x} + U'(x) = \sqrt{2\gamma k_B T} w(t), \quad (8.26)$$

where γ is the damping ratio, k_B is Boltzmann’s constant, T is the temperature, and $w(t)$ is a white noise process (Risken, 1996). We consider the double-well potential

$$U(x) = -\frac{\alpha}{2}x^2 + \frac{\beta}{4}x^4. \quad (8.27)$$

An example trajectory of this system is shown in Fig. 8.7, displaying both small oscillations within each well and random large jumps between wells.

We model a particle of unit mass in a symmetric double-well potential subject to thermal fluctuations with the second-order Langevin dynamics

$$\ddot{x} + 2\gamma\dot{x} = \alpha x - \beta x^3 + \sqrt{2\gamma k_B T} w(t). \quad (8.28)$$

Nondimensionalizing with the relaxation timescale γ^{-1} and the length scale $\sqrt{\gamma^2/\beta}$ given

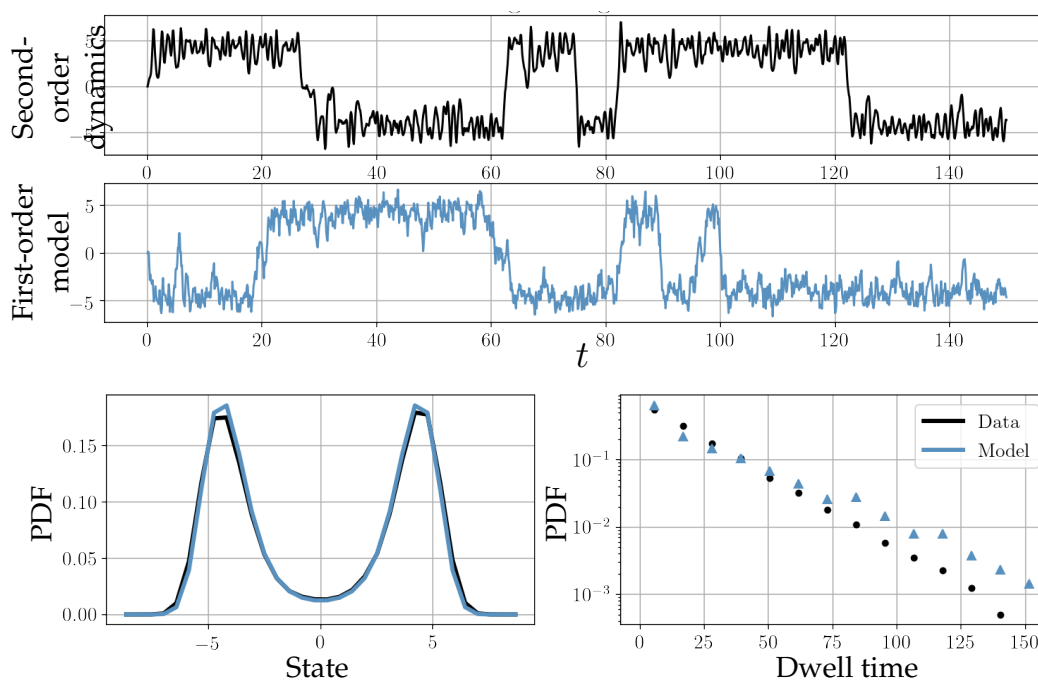


Figure 8.7: **Particle in a one-dimensional double-well potential.** Even when the system is far from the pitchfork bifurcation ($\epsilon = 20$), the dynamics are dominated by bistable switching behavior (top). Langevin regression identified a reduced-order approximation to this system (middle), which is consistent with both the state probability distribution (bottom left) and the distribution of metastable dwell times (bottom right).

by the nonlinear term, the dynamics are

$$\ddot{x} + 2\dot{x} = \epsilon x - x^3 + \sigma_x w(t), \quad (8.29)$$

where $\epsilon = \alpha/\gamma^2$ and $\sigma_x^2/2$ is the dimensionless energy of the thermal fluctuations.

The equation of motion generated by this potential can be nondimensionalized to the form

$$\frac{d}{dt} \begin{bmatrix} x \\ \dot{x} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ \epsilon - x^2 & -2 \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma \end{bmatrix} w. \quad (8.30)$$

The drift undergoes a supercritical pitchfork bifurcation at $\epsilon = 0$, where the origin loses stability to the pair of fixed points $x = \pm\sqrt{\epsilon}$. For small ϵ , an invariant manifold approximation gives a first-order model based on the pitchfork bifurcation normal form:

$$\dot{x} = \lambda(\epsilon)x - \mu(\epsilon)x^3 + \tilde{\sigma}w(t). \quad (8.31)$$

Note that this does not correspond to the large dissipation limit in which the second-order dynamics can be approximately reduced to a first-order equation for the velocity.

The stochastic normal form may be derived as follows. If the system is recast into first-order form and linearized about the origin, the eigenvalues are $\lambda_{1,2} = 1 \pm \sqrt{1 + \epsilon}$, with eigenvector representation

$$\begin{bmatrix} x \\ \dot{x} \end{bmatrix} = v_1\phi_1(t) + v_2\phi_2(t), \quad (8.32)$$

where $v_{1,2}$ are the corresponding eigenvectors. Close to the bifurcation, we assume the following:

1. **Invariant manifold reduction:** The dynamics are restricted to the one-dimensional subspace spanned by the unstable eigenvector v_1 . Since the eigenvectors are not orthogonal, the amplitude of the stable eigenvector is an algebraic function of $\phi_1(t)$ on restriction to this subspace. To leading order, $\phi_2 = h\phi_1$.
2. **Normal form:** The drift dynamics in this subspace are given by the normal form for the pitchfork bifurcation with unknown parameter μ :

$$\dot{\phi}_1 = \lambda_1(\epsilon)\phi_1 - \mu(\epsilon)\phi_1^3 + \sigma_\phi w(t). \quad (8.33)$$

3. **Dynamical consistency:** The reduced-order dynamics preserve the fixed points of the drift function of the full system. That is, $\ddot{x} = \dot{x} = 0$ at $x = 0, \pm\sqrt{\epsilon}$.

These assumptions imply that $h = -\lambda_1/\lambda_2$ and $\mu = (1 + h)^2\lambda_1/\epsilon$. Furthermore, under the similarity transform and Itô's lemma for a change of variables (Risken, 1996) the diffusion coefficient becomes $\sigma_\phi = \sigma_x/2\sqrt{1 + \epsilon}$. Inverting the transformation, $x \approx (1 + h)\phi_1$, which can be used to recast the eigenvector dynamics into the form of Eq. (8.31).

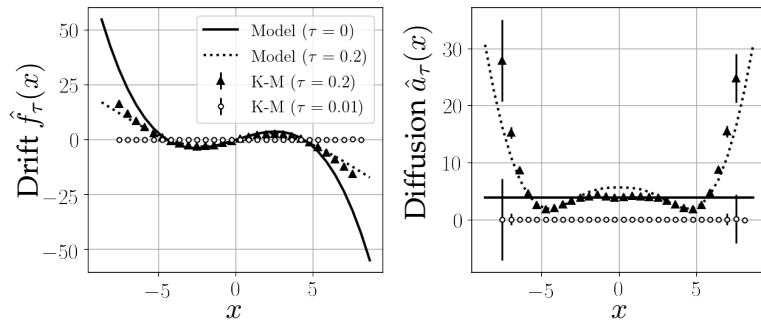


Figure 8.8: **Finite-time effects in Kramers-Moyal average.** Non-Markovian effects from order reduction cause the Kramers-Moyal average to fail at high sampling rates (\circ). The forcing appears uncorrelated when the system is subsampled (\triangle). Langevin regression identifies a model (—) which is consistent with the Kramers-Moyal average at finite sampling rate (- -).

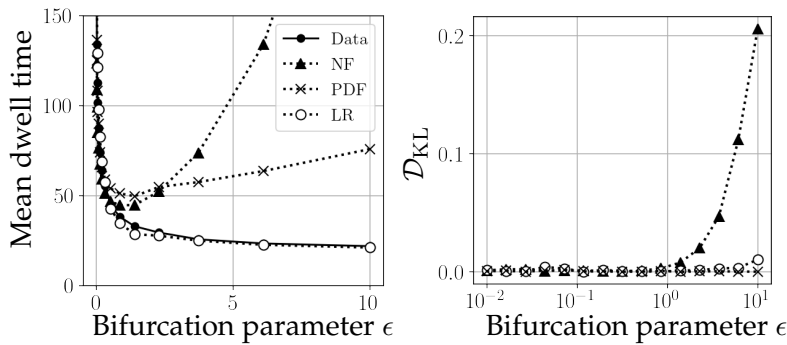


Figure 8.9: **Comparison of models far from bifurcation with Monte Carlo simulation.** Both the normal form (NF) and the PDF fit (Rigas *et al.*, 2015) quickly fail to match the metastable dwell time, while the Langevin regression (LR) model continues to be accurate far from the bifurcation. The Kullback-Leibler divergence \mathcal{D}_{KL} measures the difference between the model and system probability distribution.

As Fig. 8.7 shows for $\epsilon = 20$, even far from the bifurcation the bistability still dominates the dynamics, although the invariant manifold approximation used to reduce the order of the normal form model no longer holds. A first-order equation of the form of Eq. (8.31) can capture this bistability at the cost of ignoring the small oscillations within each potential well. In other words, the goal is to coarse-grain the dynamics while preserving the statistical properties of the system.

Constructing a first-order Langevin model for the position x is made difficult by the fact that the time series is smoothed by integration of the thermal fluctuations forcing \dot{x} in Eq. (8.30). The neglected degree of freedom introduces non-Markovian behavior and confounds the Kramers-Moyal average, which tends towards zero with fast sampling rates (Fig. 8.8). However, by choosing a coarse enough sampling rate so the subsampled dynamics appear Markovian and correcting for the finite-time effects, Langevin regression is able to identify a first-order model that captures both the probability distribution $p(x)$ and the distribution of residence times in each metastable well.

As shown in Fig. 8.9, the Langevin regression model has similar fidelity to the analytic normal form close to the bifurcation, but the data-driven model maintains statistical accuracy well beyond the region where the normal form is valid. Also shown are results




Figure 8.10: **Experimental configuration for the axisymmetric wake.** The bluff body is mounted from the wind tunnel ceiling (left) and the base pressure distribution is measured from 64 pressure taps (middle). The symmetry-breaking instability of the laminar flow persists in the fully turbulent wake, although the center of pressure appears to wander randomly, as seen in the probability distribution $p(x, y)$ (right).

for the same model structure, but with parameters estimated by regression to the empirical PDF. At each value of the bifurcation parameter ϵ , the second-order dynamics are driven by noise $\sigma = (\sqrt{\epsilon} + \epsilon)/2$ so that the mean dwell time maintains a similar order of magnitude throughout the range of comparison.

Turbulent axisymmetric wake Turbulence is a notoriously challenging problem that exemplifies many of the difficulties of stochastic modeling. A turbulent fluid flow is deterministic in principle, but the large and continuous range of spatiotemporal scales often necessitates statistical analysis. However, unlike Brownian motion, turbulence is far enough from equilibrium that it does not obey the principle of detailed balance; the machinery of statistical mechanics cannot easily be applied to viscous turbulence (Kraichnan & Chen, 1989). Nevertheless, many turbulent flows are dominated by large-scale coherent structures whose regular evolution is suggestive of low-dimensional dynamics, despite the unpredictability introduced by strong coupling to the smaller scales in the flow. In particular, high Reynolds number flows are characterized by a wide separation between the slow macroscopic dynamics and the faster turbulent fluctuations (Tennekes & Lumley, 1972).

In this context, Langevin regression is a natural extension of the data-driven modeling methods that have proven successful at identifying low-dimensional dynamics in laminar flows (Loiseau & Brunton, 2018; Loiseau *et al.*, 2018b; Loiseau, 2020; Deng *et al.*, 2020). Just as these methods are capable of generalizing the near-bifurcation results of weakly nonlinear analyses (Sipp & Lebedev, 2007; Meliga *et al.*, 2009) and POD-Galerkin models (Noack *et al.*, 2003), here we aim to further extend this philosophy to turbulent flows by modeling all but the most important degrees of freedom as stochastic forcing.

We demonstrate stochastic model identification on the experimentally measured pres-

sure distribution on an axisymmetric bluff body, visualized in Fig. 8.10 and described in detail by Rigas *et al.* (2014). The Reynolds number based on the body diameter is $Re \sim 2 \times 10^5$. This flow is a stereotypical configuration that exhibits several features important to drag reduction applications.

The low Reynolds number laminar wake behind an axisymmetric bluff body is symmetric, but undergoes a supercritical pitchfork bifurcation at $Re \sim 10^2$ so that the center of pressure is offset at a nonzero radial amplitude (Meliga *et al.*, 2009). Since the unstable symmetric wake has lower drag than the asymmetric configuration, stabilizing the symmetric state is a major goal of flow control studies (Beaudoin *et al.*, 2006; Brackston *et al.*, 2016). Simple low-dimensional models that accurately represent process noise, energy transfers, and frequency dynamics could significantly improve closed-loop control schemes. The symmetry-breaking instability continues to dominate the wake dynamics in the turbulent regime, although the location of the center of pressure tends to wander randomly (Grandemange *et al.*, 2013; Rigas *et al.*, 2014, 2015).

As a macroscopic proxy for the amplitude of the symmetry-breaking, we model the evolution of the center of pressure, or the centroid of the pressure distribution on the back of the body. The base pressure distribution is measured using 64 evenly spaced taps, as shown in Fig. 8.10. Since the center of pressure is a global integral quantity, we expect that fluctuations will be roughly Gaussian, based on the central limit theorem, although time correlations in the forcing still necessitates the use of subsampling and finite-time corrections.

A simple dynamical model that captures the symmetry-breaking behavior is the normal form of the pitchfork bifurcation forced by Gaussian white noise (Rigas *et al.*, 2015). The radial component of the Langevin equation for a symmetric two-dimensional pitchfork bifurcation forced by additive white noise is

$$\dot{r} = \lambda r - \mu r^3 + \frac{\sigma^2}{2r} + \sigma w(t), \quad (8.34)$$

where the $1/r$ term appears as a consequence of Ito's lemma for a change of variables in stochastic systems. In one dimension, the steady-state Fokker-Planck equation can be solved analytically for this model and the free parameters can be identified based on a fit to the empirical probability distribution and mean-square displacement (Rigas *et al.*, 2015).

As shown in Fig. 8.11, this model agrees reasonably well with the observed statistics for the center of pressure, suggesting that the stochastic modeling approach is a promising description for the leading global degrees of freedom. However, it is difficult to extend this modeling methodology to more complex systems. Even in one dimension, the drift and diffusion do not appear independently in the analytic steady-state PDF; additive noise can be estimated from the mean-square displacement, but this poses a challenge for multiplicative noise. In higher dimensions, the Fokker-Planck equation does not have an analytic solution for general drift and diffusion, although model parameters might be es-




Figure 8.11: Statistical evaluation of the axisymmetric wake models. The Langevin regression model (blue) better matches both the power spectral density (left) and tails of the PDF (right) compared with the pitchfork normal form (8.34) with coefficients estimated by PDF fitting, as in Rigas *et al.* (2015). The models are similar, but Langevin regression identifies a quadratic state-dependent noise (Fig. 8.12). The power spectrum is premultiplied by Strouhal number $St = fU/D$, a dimensionless frequency. The large peak at $St \approx 0.2$ corresponds to vortex shedding, which is essentially indistinguishable from the symmetry-breaking instability in the base pressure distribution (Brackston *et al.*, 2016).

timated by optimizing the solution of a numerical solver. This approach cannot resolve oscillatory behavior such as vortex shedding, since temporal information is lost in the steady-state distribution.

Since we do not have a known form of the model in this case, besides the intuition for the pitchfork normal form, we apply the stochastic SINDy procedure described in Sec. 8.2. Based on symmetry considerations, we include only odd polynomials in the library of drift functions. The $1/r$ term in the drift, due to the representation in polar coordinates, is also accounted for in the optimization routine. The model selection criteria shown in Fig. 8.12 show a clear Pareto-optimal model of the form

$$\dot{r} = \lambda r - \mu r^3 + \frac{\sigma^2}{2r} + (\sigma_0 + \sigma_1 r^2)w(t). \quad (8.35)$$

Fig. 8.12 also shows that the finite-time Kramers-Moyal coefficients predicted by this model closely match those estimated by the finite-time conditional average.

The identified model is similar to that proposed by Rigas *et al.* (2015), with the addition of quadratic multiplicative noise. This modification better matches the tails of the probability distribution, as shown in Fig. 8.11. Monte Carlo simulation of the Langevin models also shows that the multiplicative noise leads to a more accurate power spectral density than the model based on fitting the PDF. Quadratic multiplicative noise was previously proposed as an important modification for a spatial Langevin model of turbulence (Ragwitz & Kantz, 2001). Multiplicative terms may be a result of neglecting degrees of freedom with bilinear coupling to the macroscopic variables (Majda *et al.*, 2001).

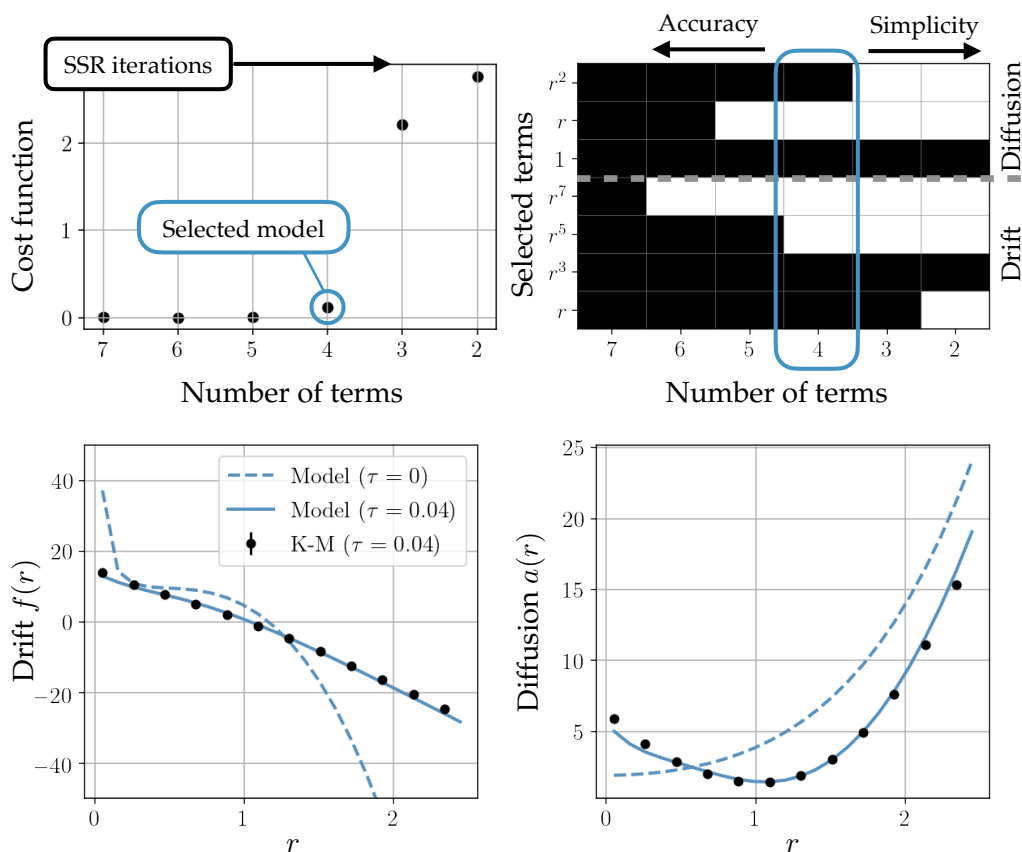


Figure 8.12: **Model selection and Kramers-Moyal coefficients for the axisymmetric wake.** The reverse-greedy Sparse Stepwise Regression identifies a hierarchy of candidate models with varying tradeoffs between accuracy and complexity. The optimal model has the fewest terms before the cost function begins to climb, indicating the remaining terms are essential. In this case, the optimal model is a pitchfork bifurcation normal form forced by quadratic multiplicative noise (top). The model to the right of the optimal model includes only additive noise and corresponds to the model proposed by Rigas *et al.* (2015), while additional terms leads to a higher-order Stuart-Landau equation. The selected model closely matches the empirical finite-time Kramers-Moyal coefficients (bottom), the state PDF, and the power spectral density (Fig. 8.11).

This model for the symmetry-breaking instability does not resolve the peak in the power spectrum near $St \approx 0.2$, which is related to vortex shedding in the wake (Rigas *et al.*, 2014, 2017). Analysis of the corresponding laminar flow indicates that a more complete model of the wake might include three complex amplitudes to capture this periodic component and its interactions with the symmetry-breaking mode (Meliga *et al.*, 2009). However, the vortex shedding is only weakly observable from base pressure sensors, since it mainly takes place downstream from the body (Brackston *et al.*, 2016); for the same reason it is less important aerodynamically than the symmetry-breaking.

Langevin regression is therefore able to clearly identify a simple low-dimensional model from experimental measurements of turbulence. The sparse stochastic model is consistent with both known flow physics and empirical statistics, suggesting that approximating the evolution of global variables with nonlinear Langevin dynamics may be a promising direction in the low-dimensional modeling of turbulent flows.

8.4 Discussion

There is a long history in the physical sciences of approximating complex multiscale systems with reduced-order models driven by stochastic forcing. However, as this approach has spread in popularity, it is not always clear that the assumptions underpinning the rigorous treatment of nonequilibrium statistical mechanics continue to hold. For example, the subscale degrees of freedom in systems like turbulent fluid flows often do not decorrelate fast enough to appear as delta-correlated white noise. Nevertheless, experience suggests that simple stochastic models are often good approximations to systems that violate some of these assumptions. Data-driven modeling has gained significant attention in this context for its ability to construct consistent empirical models without the restrictions of classical analytic approaches.

In this work, we have integrated and generalized three previously disparate approaches to data-driven stochastic modeling, combining sparse model selection based on the Kramers-Moyal average (Boninsegna *et al.*, 2018) with finite sampling-rate corrections (Lade, 2009; Honisch & Friedrich, 2011) and steady-state PDF fitting (Rigas *et al.*, 2015). The proposed modeling framework is designed to identify nonlinear Langevin-type dynamics from noisy experimental data, optimizing parameter estimates with both the forward and adjoint Fokker-Planck equations. Critically, the finite-time corrections allow the method to model systems driven by correlated noise, or fast deterministic degrees of freedom, while sparse stepwise regression can be used to discover the structure of the model when it is unknown *a priori*.

We have demonstrated the flexibility and generality of Langevin regression on three examples. First, the method reconstructs a noisy Stuart-Landau oscillator from partial observations, illustrating the capability to model higher-dimensional systems. Second, we investigated reducing the second-order dynamics of a particle in a one-dimensional double-well potential to the corresponding first-order normal form. The data-driven

model closely matches the statistical behavior of the system far from the bifurcation, well beyond the region where the analytic normal form is valid. Finally, we apply Langevin regression to experimental measurements of a turbulent bluff body wake, identifying a model for the evolution of the base center of pressure. In this case the SSR model selection procedure identifies a model consistent with previous work (Rigas *et al.*, 2015), but modified with a quadratic, state-dependent noise term that better approximates the power spectrum and the long tails of the PDF.

These results indicate that the proposed method is capable of accurately modeling a broad range of systems from limited experimental observations. However, we recognize two limitations of the method as presented here. First, we can currently only construct first-order models. In principle, higher-order dynamics can be recast as a system of first-order equations, provided generalized coordinates and velocities can be measured, but this is often not the case in practice. It may therefore be easiest to model systems with stereotypical macroscopic dynamics reminiscent of a normal form or amplitude equation. Second, although the theory generalizes naturally to higher dimensions (and we have demonstrated a two-dimensional system), the practical limitation lies in constructing n -dimensional histograms for Kramers-Moyal coefficients and in solving the Fokker-Planck equations on the resulting grid. This challenge might be addressed either with more efficient Fokker-Planck solvers or with recently-proposed methods that avoid the need for histograms at the price of enforcing consistency with the PDF (Schneider *et al.*, 2020; Frishman & Ronceray, 2020; Brückner *et al.*, 2020).

Despite these limitations, the proposed method can still be readily applied to a broad range of systems; nonlinear stochastic systems with one or two degrees of freedom have been used to model the global behavior of systems ranging from neuroscience (Laing *et al.*, 2010) to molecular dynamics (Legoll & Lelièvre, 2010) and aerodynamics (Rigas *et al.*, 2015). The success of simple heuristics combined with stochastic models for feedback control (e.g. Brackston *et al.* (2016)) also suggests that these data-driven models could be integrated in a control scheme, although this will require modeling the effects of actuation on the system in addition to the homogeneous behavior.

Another goal of modeling is to uncover the latent low-dimensional structure of macroscopic dynamics. In fluid dynamics, for instance, this is often conceptualized as a small set of global modes whose amplitudes evolve according to low-dimensional nonlinear dynamics (Holmes *et al.*, 1996). Moving beyond integral quantities such as the center of pressure, Langevin regression could be used to identify a data-driven, stochastic counterpart to Galerkin-type reduced-order models and resolve important nonlinear interactions between the large-scale structures of the flow.

In a broader context, the ability to identify reduced-order stochastic models from noisy experimental measurements of multiscale nonlinear dynamics will unlock a powerful set of tools for a much wider range of systems. Data-driven methods allow for the treatment of not only systems which break the strict assumptions of classical stochastic modeling, but also data from ecology, epidemiology, and neuroscience for which first-principles gov-

erning equations are unavailable. Tools such as the Kramers-Moyal average already have a history of success in a variety of fields. It is our hope that the proposed method will build on this legacy and extend these successes to an even more extensive class of complex systems.

Chapter 9

A MEAN-FIELD MODEL OF THE AXISYMMETRIC WAKE

Despite being nominally deterministic, turbulent flows are characterized by multiscale spatiotemporal chaos. Many traditional analyses have therefore relied on statistical descriptions (Monin & Yaglom, 1971). However, it is now known that many inhomogeneous flows are dominated by energetic coherent structures at large scales and low frequencies relative to Kolmogorov's universal equilibrium range (Aubry *et al.*, 1988; Berger *et al.*, 1990; Holmes *et al.*, 1996). By leveraging this intrinsic structure, reduced-order models of turbulent flows promise to advance engineering objectives in design, optimization, and control (Noack *et al.*, 2011; Brunton & Noack, 2015a; Rowley & Dawson, 2017). However, balancing accuracy and efficiency by simultaneously modeling the evolution of the large-scale structures while accounting for the effects of incoherent fluctuations has been notoriously challenging, especially in a non-invasive fashion that is suitable for experimental measurements.

This modeling challenge stems from the multiscale nature of turbulence. In contrast to systems that can be treated by classical statistical mechanics, there is no strict separation of scales between low-frequency coherent dynamics and turbulent fluctuations. Governing equations for filtered or averaged variables can be derived, but the influence of the unresolved scales cannot be eliminated from the coarse-grained equations, leading to the *closure problem*. Numerous deterministic closure models have been proposed in the context of high-dimensional Reynolds-averaged Navier-Stokes (RANS) or large eddy simulation (LES) methods (Pope, 2000), although these approaches are still high-dimensional and computationally expensive.

Alternatively, in order to take advantage of persistent, energetic coherent structures, semi-empirical projection-based methods can be used to derive a compact approximation to the dynamics in the form of amplitude equations governing the evolution of global modes. For example, perhaps the most widely used reduced-order modeling method in fluid dynamics is Galerkin projection of the Navier-Stokes equations onto a modal basis (Aubry *et al.*, 1988; Holmes *et al.*, 1996; Noack *et al.*, 2003). However, because projection-based methods act as a spatial filter, they face the same closure problems as RANS and LES methods. Projection methods also require access to both a high-fidelity numerical solver and information about the flow that is difficult to obtain experimentally.

To address the analytic challenges of turbulence modeling, data-driven methods have long played a role in modal analysis (Taira *et al.*, 2017) and reduced-order modeling (Noack *et al.*, 2011; Rowley & Dawson, 2017). Recent advances in machine learning have generated increased interest in data-driven methods for fluid dynamics (Brenner *et al.*, 2019;

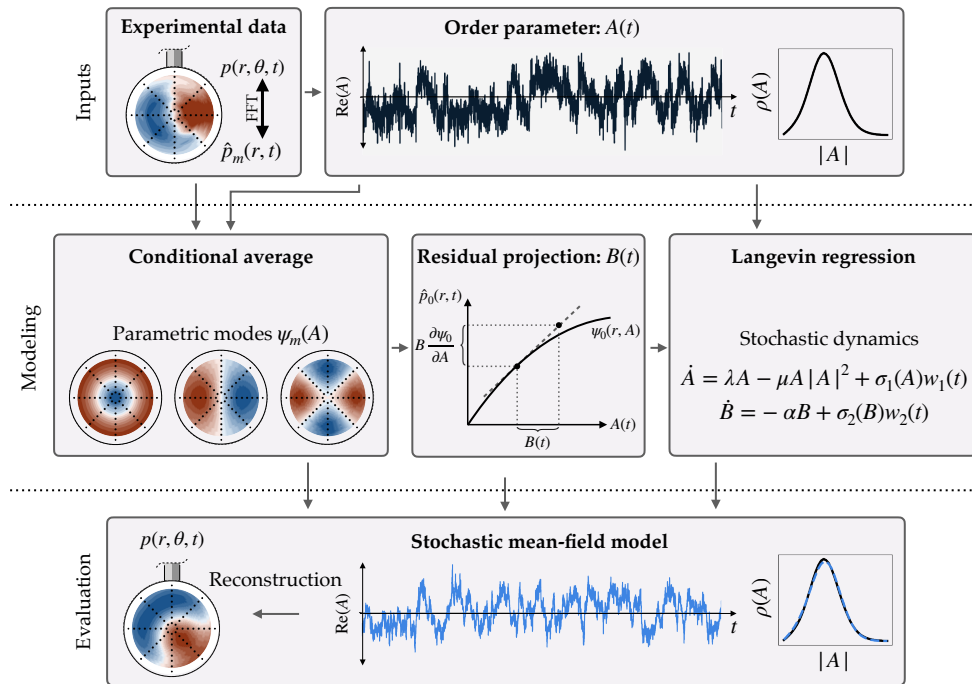


Figure 9.1: Overview of the model development. Beginning with an order parameter $A(t)$ computed from the base pressure measurements, the spatial modes ψ_m are computed at each wavenumber m by conditional averaging. These modes define an approximate slow manifold $\psi(A)$ that captures the dominant anti-symmetric behavior. To fully resolve the symmetric fluctuations, we introduce a generalized shift mode with amplitude $B(t)$, defined by projection onto the tangent space of this manifold. Finally, we identify a nonlinear stochastic dynamical system model with Langevin regression. We can compare the statistics, including the empirical probability distribution ρ , of the model to the experiment with Monte Carlo evaluation of the stochastic system. Because the proposed model (bottom row) is random in nature, it can only reproduce the experimental data (top row) in a statistical sense; in general neither the time series nor the pressure distributions will match on a point-by-point basis.

Bar-Sinai *et al.*, 2019; Brunton *et al.*, 2020; Raissi *et al.*, 2020), including for turbulence modeling (Ling *et al.*, 2016; Duraisamy *et al.*, 2019; Maulik *et al.*, 2019) and forecasting (Novati *et al.*, 2021). Despite the expressive power of modern machine learning methods, it is challenging to develop models that are robust, generalizable, and interpretable. The sparse identification of nonlinear dynamics (SINDy) framework (Brunton *et al.*, 2016a) has promise for interpretable low-dimensional modeling, and it is able to incorporate partial physical knowledge including symmetries, conservation laws, and invariant manifold structure (Loiseau & Brunton, 2018; Loiseau *et al.*, 2018a; Deng *et al.*, 2020). SINDy has mainly been applied to laminar flows with the exception of recent work developing RANS closure models (Beetham & Capecelatro, 2020). Following these successes, our objective is to model coherent structure dynamics with a SINDy-type approach, approximating turbulent fluctuations as noise acting on a few global integral quantities.

In this chapter, we develop a sparse nonlinear model of a fully turbulent wake exper-

iment. Our model takes the form

$$\dot{x} = f(x) + \sigma(x)w(t), \quad (9.1)$$

where x is a generalized state vector and $w(t)$ is Gaussian white noise. Whereas RANS and LES models use a relatively high-dimensional discretization of the flow field, here we assume that x represents a small set of modal coefficients, order parameters, or other integral quantities capturing relevant large-scale structure in the flow. Moreover, this description in terms of deterministic drift dynamics $f(x)$ forced by diffusion $\sigma(x)$ is a fundamentally different approach to turbulence modeling. RANS and LES methods seek a *local, deterministic* closure model for the effects of the unresolved scales in terms of resolved variables, (9.1) approximates the *global, statistical* influence of turbulent fluctuations on the state x .

We learn the model in (9.1) using the Langevin regression method described in Chapter 8 (see Fig. 9.1). Progress in model discovery has advanced our ability to approximate stochastic dynamics from limited experimental data (Ragwitz & Kantz, 2001; Honisch & Friedrich, 2011; Boujo & Cadot, 2019; Boninsegna *et al.*, 2018; Callaham *et al.*, 2021*b*; Sieber *et al.*, 2020; Frishman & Ronceray, 2020; Schneider *et al.*, 2020), even if the model structure is *a priori* unknown (Boninsegna *et al.*, 2018; Callaham *et al.*, 2021*b*; Billings, 2013). These methods extend traditional stochastic modeling beyond near-equilibrium systems, providing a useful approximation despite the absence of a strict scale separation.

We demonstrate this approach to model the turbulent wake behind an axisymmetric bluff body, shown in Fig. 9.3. We reduce the dimension of the experimental measurements by modal decomposition. Traditional methods decompose the field into a fixed set of spatial modes with time-varying amplitudes. The Fourier basis is a convenient representation for coordinates with translational or rotational symmetry. In the present case, the flow is rotationally symmetric, so azimuthal variations can be expanded with Fourier modes $e^{im\theta}$.

Analysis of the Navier-Stokes equations in the wavenumber domain reveals that the nonlinear advective term only admits *triadic* interactions in which both the forced and forcing modes have wavenumbers that sum to zero. In particular, the axisymmetric fluctuations at $m = 0$ are driven by the “self-interaction” of the complex-conjugate components at $m = \pm 1, \pm 2$, etc., as shown in Fig. 9.2. The full nonlinear flow is made up of a complicated network of such interactions across all scales, making it notoriously difficult to construct simplified models with reduced degrees of freedom.

One approach to circumventing this problem, typically near the threshold of some instability, is to perform an asymptotic multiscale expansion and assume that higher-order nonlinear interactions may be neglected. This approximation can resolve the leading-order mean flow deformation due to the self-interaction of the instability mode, a central feature of the Stuart-Landau nonlinear stability mechanism (Stuart, 1958; Landau & Lifshitz, 1959). The triadic interactions are truncated at leading order (Fig. 9.2). In this weakly nonlinear regime, the mean flow deformation can be treated as another fixed spa-

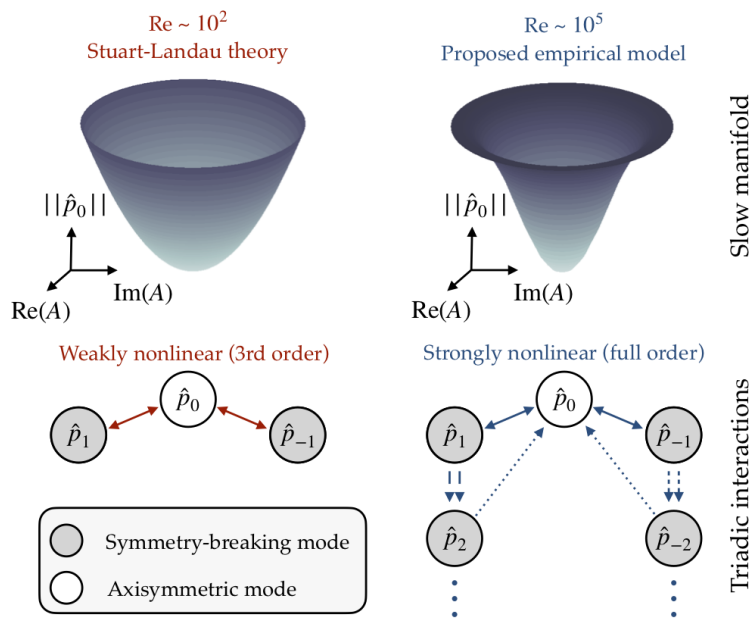


Figure 9.2: Visualization of approximate slow manifolds and the nonlinear interactions responsible for generating them. Weakly nonlinear analyses of laminar flows typically neglect higher-order interactions for small fluctuations, leading to a parabolic manifold for the low-dimensional dynamics (left). Using conditional averaging, we show that these interactions are necessary to explain observations of the turbulent wake (right). Pairs of lines with similar styles indicate the structure of the leading-order triadic interactions. The physical interactions are between velocity components, but here we use the base pressure field as a proxy for coherent structures in the wake.

tial mode with an amplitude that depends on the strength of the instability (Noack *et al.*, 2003; Sipp & Lebedev, 2007; Meliga *et al.*, 2009).

However, these tools are largely theoretical and numerical; here the flow is fully turbulent and thus strongly nonlinear. In order to avoid the assumption of weak nonlinearity and explore the spatial structure of the mean flow deformation, we estimate the modes from measurements of the base pressure distribution by reducing the symmetry via phase alignment and averaging conditioned on the center of pressure. The amplitude dependence of the symmetric mode deviates significantly from the polynomial scaling predicted by weakly nonlinear analysis, confirming the need for a parametric basis to capture the effect of strongly nonlinear interactions.

Following the conditional average, the method proceeds as shown schematically in Fig. 9.1. The modal expansion yields physically meaningful order parameters $A(t)$ and $B(t)$ related to symmetry-breaking and mean-field deformation in the wake, respectively. We apply Langevin regression to identify an interpretable dynamical system that models the broadband turbulence as stochastic forcing of the low-dimensional symmetry-breaking dynamics. The resulting model is a stochastic Stuart-Landau equation similar to those proposed in previous studies of similar configurations (Rigas *et al.*, 2015; Brackston *et al.*, 2016; Boujo & Cadot, 2019; Herrmann *et al.*, 2020), but with the addition of state-dependent forcing and an additional degree of freedom. Monte Carlo evaluation of the Langevin system compares favorably to the experimental statistics, and the model provides new physical insights into low-frequency fluctuations of the axisymmetric recirculation bubble. Each of these will be investigated below.

9.1 Symmetry-breaking in the bluff body wake

The turbulent wake in Fig. 9.3 provides a rich test system for model discovery. Flow with free-stream velocity U_∞ and kinematic viscosity ν is deflected around a cylindrical blunt-nosed body with diameter D . The spatiotemporal dynamics of the wake are determined by the Reynolds number $Re = DU_\infty/\nu$. Flows past bluff bodies at low Reynolds number often exhibit stereotypical global instabilities such as von Kàrmàn vortex streets or steady symmetry-breaking wake deflection (Berger *et al.*, 1990; Noack *et al.*, 2003; Meliga *et al.*, 2009; Grandemange *et al.*, 2012; Östh *et al.*, 2014). In the laminar regime, nonlinear saturation of the exponential growth of these instabilities can often be described by Stuart-Landau theory, in which the fluctuations deform the mean flow in a stabilizing feedback loop Stuart (1958); Landau & Lifshitz (1959). The amplitude $A(t)$ of the instability mode is governed by the cubic Stuart-Landau equation:

$$\frac{dA}{dt} = (\lambda - \mu|A|^2) A. \quad (9.2)$$

The term in parentheses is the effective eigenvalue of the instability mode, as modified by the mean flow deformation. When the real part of λ is positive, small perturbations grow exponentially until the instantaneous growth rate reaches a balance with $\mu|A|^2$.

Global modes that are qualitatively similar to the laminar instabilities appear to dominate these flows well into the turbulent regime (Grandemange *et al.*, 2013; Rigas *et al.*, 2014, 2016). Although the Stuart-Landau mechanism is typically derived via an asymptotic weakly nonlinear analysis, the resulting amplitude equations are often assumed to approximately describe the evolution of turbulent coherent structures as well (Rigas *et al.*, 2015; Brackston *et al.*, 2016; Boujo & Cadot, 2019; Herrmann *et al.*, 2020; Sieber *et al.*, 2020).

At very low Reynolds number, the wake is steady, axisymmetric and laminar. The flow undergoes two bifurcations, becoming linearly unstable to a steady symmetry-breaking mode ($Re_c^1 \approx 424$) and a second pair of unsteady vortex-shedding modes ($Re_c^2 \approx 605$), both with azimuthal wavenumber $m = \pm 1$ (Rigas *et al.*, 2016). By approximating these as a single codimension-2 bifurcation, the weakly supercritical flow can be approximated with an asymptotic expansion and normal form dynamics (Meliga *et al.*, 2009). These instability modes continue to dominate the coherent part of the flow even in fully developed turbulence, as has been shown for the present experimental data at $Re \approx 2 \times 10^5$ (Rigas *et al.*, 2014). The base pressure distribution is measured from a set of 64 equally spaced pressure taps from which a time series of 8.9×10^6 measurements are sampled at 225 Hz; further details are given in (Rigas *et al.*, 2014). Equally spacing the pressure taps allows us taking advantage of the geometric symmetry with a Fourier series representation of the pressure field. Hence, the sensor distribution determines the wavenumber resolution of the analysis. In particular, eight sensor stations restricts the Fourier representation to wavenumbers $m = \pm 3$, although we find that only $m = 0, \pm 1, \pm 2$ are strongly correlated with the symmetry-breaking.

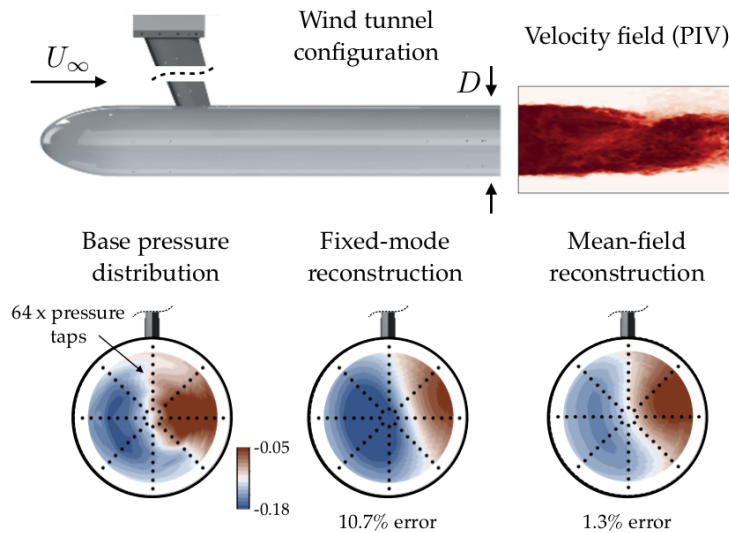


Figure 9.3: Experimental configuration for the axisymmetric wake. The signature of the dominant symmetry-breaking instability is captured by the complex order parameter defined in (9.5). The coherent fluctuations in the pressure distribution can be approximated with the mean-field model (bottom). The proposed model significantly improves the reconstruction over a traditional fixed-mode decomposition.

Even though the time series of measurements appears stochastic at this Reynolds number, the flow is characterized by semi-regular energetic structures, including vortex shedding and a symmetry-breaking wake deflection (Berger *et al.*, 1990; Grandemange *et al.*, 2013; Rigas *et al.*, 2014). These structures can be directly linked to the instability modes and weakly nonlinear dynamics of the corresponding laminar flow (Meliga *et al.*, 2009; Rigas *et al.*, 2017). Although the vortex shedding is dynamically important in the wake, it is only weakly observable from the pressure signal on the bluff body itself. This also suggests that the vortex shedding is potentially less important to practical drag reduction than the symmetry-breaking wake deflection. In this work we therefore restrict our attention to the steady symmetry-breaking instability.

This wake deflection is particularly important as it represents generic spatial symmetry-breaking behavior that occurs in a wide variety of three-dimensional bluff body wakes (Meliga *et al.*, 2009; Grandemange *et al.*, 2012; Boujo & Cadot, 2019). Moreover, this symmetry-breaking is associated with increased pressure drag, making it the target of a variety of active flow control investigations (Östh *et al.*, 2014; Barros *et al.*, 2016; Brackston *et al.*, 2016; Rigas *et al.*, 2017; Herrmann *et al.*, 2020). In several of these studies, reduced-order models have played a key role in designing and understanding the actuated system.

9.2 Mean-field theory of symmetry-breaking transitions

The Stuart-Landau nonlinear stability theory is typically derived as an asymptotic expansion in multiple timescales (see example in Section 2.2, for instance). It therefore falls into the category of weakly nonlinear analysis and is only strictly applicable near the threshold of instability, although the effect is generally understood to persist much beyond the asymptotic regime in many cases (Noack *et al.*, 2003; Luchtenburg *et al.*, 2009; Mantič-Lugo *et al.*, 2014). However, Landau also considered another limiting case of dynamics:

symmetry-breaking phase transitions of systems in thermodynamic equilibrium. Turbulence is both strongly nonlinear and far from thermodynamic equilibrium, but the observation that both limiting regimes can be described with similar equations motivates the development of phenomenological Stuart-Landau-type models for symmetry-breaking behavior in the turbulent axisymmetric wake.

A generic system in thermodynamic equilibrium that undergoes a continuous symmetry-breaking phase transition at critical temperature T_c (or temperature analogue such as Reynolds number) has an effective potential $V(T, A)$ which we assume can be expanded in the magnitude of a (generally complex) order parameter A . A canonical example of this is the Ising model, in which the statistically symmetric disorder of the high-temperature system is broken in a phase transition to a ferromagnetic state below a critical temperature. Based on physical symmetries, the effective potential can be expanded as:

$$V(T) = V_0(T) + V_1(T)|A|^2 + V_2(T)|A|^4 + \dots \quad (9.3)$$

This order parameter, which here we assume to be small when suitably nondimensionalized, quantifies the degree of symmetry-breaking in the system. The equilibrium condition A_* at a given temperature is determined by the minimum free energy with respect to $|A|$. To leading order, $A_*(T) = \sqrt{-V_1(T)/2V_2(T)}$.

For the high Reynolds number axisymmetric wake studied in this work, the unsteady aerodynamic center of pressure serves as an order parameter capturing the symmetry-breaking wake deflection. When nondimensionalized by the body diameter, its value is small even far from the bifurcation with mean value $\bar{A} \approx 0.032$.

The system is unsteady but statistically stationary. In the thermodynamic perspective, the instantaneous field is disturbed from the minimum-potential state by broadband turbulence. We model this as near-equilibrium thermal fluctuations in overdamped Langevin dynamics (Risken, 1996):

$$\frac{dA}{dt} = -\nabla V(A) + \Sigma(A)w(t), \quad (9.4)$$

where $w(t)$ is a white noise process and Σ is the diffusion function. Expanding $V(A)$ in $|A|$ to third order, the Langevin model would take the form of the Stuart-Landau equation (9.2) forced by white noise. Although this qualitative symmetry-based argument suggests the expected structure of the dynamics, we use the Langevin regression method for identifying nonlinear stochastic models (Callaham *et al.*, 2021b) to identify the drift and diffusion functions rather than presuppose this form, as described below.

9.3 Parametric modal expansion

Reduced-order models such as (9.1) rely on dominant low-dimensional structure in order to approximate salient features of the flow with a small set of variables; such low-

dimensionality in fluid flows often arises as a consequence of global instability. The instability mode responsible for the steady wake deflection in the laminar flow stabilizes at finite amplitude due to mean flow deformation, with temporal evolution given by compact normal form dynamics (Meliga *et al.*, 2009). The turbulent wake exhibits qualitatively similar behavior; recent studies have modeled the aerodynamic center of pressure with stochastic Stuart-Landau equations (Rigas *et al.*, 2015; Boujo & Cadot, 2019). However, while the relationship between the amplitude equations and spatial mean flow deformation is clear for the weakly nonlinear laminar case, it has been unexplored in the phenomenological models developed for turbulent flows.

Standard model reduction methods decompose the field into a set of spatial modes (Taira *et al.*, 2017) with coefficients whose time evolution is governed by the amplitude equations. Within this framework, one way to resolve mean flow deformation is with the addition of a spatial mode parameterizing the difference between the unstable steady state and the mean flow. This additional *shift mode* can either be derived empirically (Noack *et al.*, 2003) or as a natural product of an asymptotic expansion (Sipp & Lebedev, 2007). In either case, the assumption of weakly nonlinear interactions implies polynomial scaling of the deformation amplitude with respect to the amplitude of the dominant instability. This approach has proven successful in low-dimensional models of a variety of laminar flows (Meliga *et al.*, 2009; Luchtenburg *et al.*, 2009; Loiseau & Brunton, 2018; Deng *et al.*, 2020). However, in strongly nonlinear turbulent flows this fixed modal basis cannot generally be expected to resolve the mean flow deformation.

In order to describe the symmetry-breaking and associated mean-field deformation, we model the evolution of the base pressure distribution $p(r, \theta, t)$, where r and θ are polar coordinates on the circular bluff body base. Although the pressure is not a dynamic variable in incompressible flow, the base pressure can be used as a convenient and experimentally accessible proxy for the energetic coherent structures in the wake, since these limited observations can be clearly connected to previously observed wake structures based on symmetries and spectral energy content (Rigas *et al.*, 2014). The radial coordinate of the unsteady aerodynamic center of pressure is also a natural order parameter for the degree of symmetry-breaking.

Let $p^0(r)$ be the pressure associated with the unstable axisymmetric steady state, which is unknown and experimentally inaccessible, and $\bar{p}(r)$ be the temporal mean field estimated from the pressure taps. The self-interaction of the velocity fluctuations associated with p' deforms the unstable steady state to the mean flow via the Reynolds stresses (Pope, 2000).

We define the unsteady aerodynamic center of pressure on the bluff body base as a complex-valued order parameter $A(t)$:

$$A(t) = \frac{1}{\int p(r, \theta, t) r dr d\theta} \int p(r, \theta, t) r e^{i\theta} dr d\theta. \quad (9.5)$$

The amplitude $|A(t)|$ is a measure of the degree of asymmetry in the wake, while its phase




Figure 9.4: **Illustration of phase-aligned conditional averaging.** We assume that the average at small amplitudes is representative of the pressure distribution associated with the unstable steady state q^B . The field at any other point can be approximated with spline interpolation, allowing us to explore the amplitude dependence of the coherent fields. Also shown for reference is the unconditional phase-aligned mean field and the two-dimensional log-probability distribution, which is roughly analogous to a potential field.

gives the azimuthal orientation of the wake deflection. We approximate this integral with a Riemann sum over the 64 pressure sensor locations.

The instantaneous strength of the coherent antisymmetric fluctuations associated with $|A(t)|$ are responsible for the axisymmetric mean flow deformation and nonlinear equilibrium of the instability mode (Meliga *et al.*, 2009). The mean field \bar{p} and steady state p^0 are therefore associated with the mean amplitude $\bar{A} \equiv \overline{|A(t)|}$ and $A = 0$, respectively, although both fields are themselves axisymmetric. Similarly, an instantaneous amplitude $|A(t)|$ between 0 and \bar{A} is associated with an axisymmetric field interpolating between $p^0(r)$ and $\bar{p}(r)$, though the amplitude itself only directly represents antisymmetric fluctuations. In other words, the part of the instantaneous $m = 0$ field that resolves the Stuart-Landau deformation mechanism is a direct function of the order parameter. More broadly, we expect that the part of the field that is coherent with the order parameter can be revealed with an average conditioned on its amplitude.

Coherent fields via conditional averaging We expect that the higher-order nonlinear interactions in the turbulent wake may lead to more complicated amplitude dependence compared to the weakly nonlinear laminar regime. We therefore propose identifying parametric modes with a phase-aligned conditional average on the order parameter. The phase alignment reduces the symmetry of the fields; without this step all asymmetry would vanish on averaging. On the other hand, the conditional average captures the natural variation of the field with the order parameter amplitude without any assumptions on the functional form of the A -dependence.

Beginning with the Fourier decomposition into modes $\hat{p}_m(r, t)$, we compute the order

parameter with (9.5) in amplitude-phase representation $A = |A|e^{i\phi}$. The phase of the order parameter can then be removed from each field:

$$p'(r, \theta, t) = \sum_m \hat{p}_m(r, t) e^{im(\theta - \phi(t))} = \sum_m \hat{p}'_m(r, t) e^{im\theta}. \quad (9.6)$$

We divide the space of observed order parameter amplitudes $|A|$ into histogram bins centered on A_i with width $2\Delta A$. For each wavenumber m and histogram bin i , the radial component of $\psi_m(r, A_i)$ is approximated with

$$\psi_m(r, |A_i|) = \left\langle \hat{p}'_m(r, t) \middle| |A(t)| - |A_i| < \Delta A \right\rangle. \quad (9.7)$$

We also estimate the base field $p^0(r)$ as the conditional mean at $m = 0$ for the smallest histogram, i.e. over fields for which $|A(t)| < \Delta A$.

Based on the symmetry of the flow, we expand the pressure field with a Fourier series $p(r, \theta, t) = \sum_m \hat{p}_m(r, t) e^{im\theta}$. The preceding discussion suggests that the part of the field that is coherent with the symmetry-breaking might be approximated with a parametric modal decomposition $\hat{p}_m(r, t) \approx \psi_m(r, A(t))$. In contrast to a fixed space-time separation of variables, such as is used in proper orthogonal decomposition or dynamic mode decomposition, this approach allows the modes to naturally deform with the instantaneous order parameter amplitude. Once the modes $\psi_m(r, A)$ are known, such a representation reduces the field to a function of this single complex degree of freedom.

For example, if the self-interaction of the fluctuations are assumed to be weakly nonlinear and higher-order nonlinearity is neglected, the axisymmetric component can be approximated with the leading term in a polynomial expansion, i.e. $\psi_0(r, A) = p^0(r) + |A|^2 p_\Delta(r) + \mathcal{O}(|A|^4)$. In a numerical setting, the “shift mode” $p_\Delta(r)$ can be determined either through an asymptotic expansion about the unstable base flow (Sipp & Lebedev, 2007; Meliga *et al.*, 2009), or from the difference between the base and mean flows (Noack *et al.*, 2003).

Experimentally, neither of these approaches is viable, since the unstable steady state is typically unavailable. Instead, we propose identifying the parametric modes with a phase-aligned conditional average on the order parameter, visualized in Fig. 9.4. This procedure, described in detail in the Materials & Methods section, estimates the part of the field at each wavenumber that is correlated with the order parameter amplitude $|A(t)|$. A continuous estimate of the modes can then be constructed with a spline interpolation of the conditional averages $\psi_m(r, |A_i|)$.

Fig. 9.5 shows the radially integrated modes, along with the axisymmetric mode at each radial sensor location. The conditional average is compared to a fixed mode approximation where the spatial structure is fixed at its value at the mean amplitude \bar{A} . We draw several conclusions from the amplitude scaling of these modes. First, the axisymmetric field at $m = 0$ cannot be described by a fixed mode. The weakly nonlinear scaling

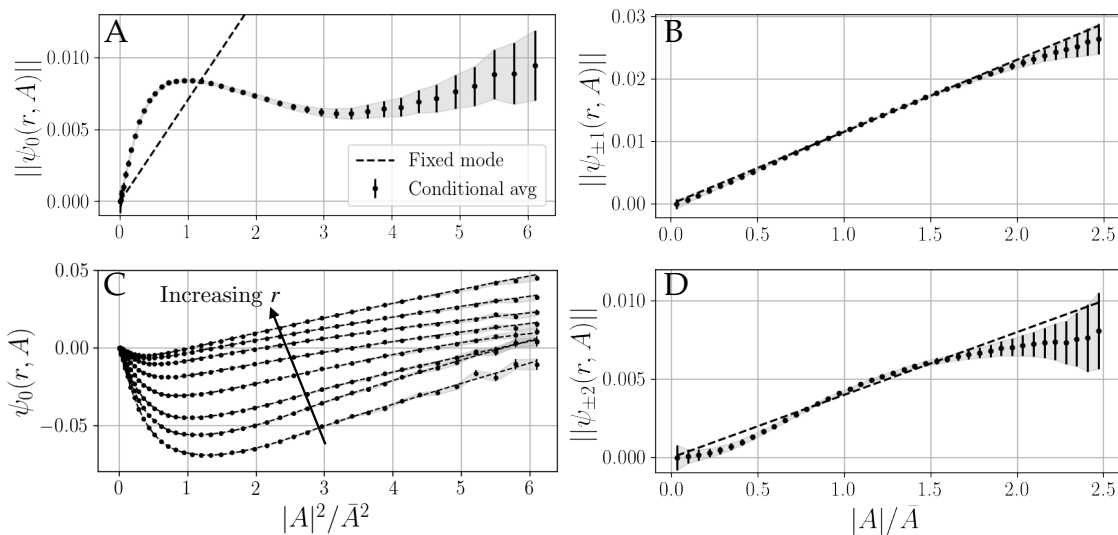


Figure 9.5: Field deformation in the conditional average. The phase-aligned average gives the coherent fields as a function of instability amplitude $|A|$ for various azimuthal wavenumbers m . The deformation of the axisymmetric part of the field cannot be explained by a single fixed mode or weakly nonlinear scaling (A), even though the deformation has a smooth dependence on the amplitude (C). We approximate this amplitude dependence with a spline fit (C, dashed lines). On the other hand, the symmetry-breaking fields at $m = \pm 1, \pm 2$ are consistent with the fixed-mode assumption of Stuart-Landau theory (B, D).

$\psi_0 \sim |A|^2$ clearly does not hold for typical amplitudes in this case. Fig. 9.5C also shows the value of the $m = 0$ field at each of the 8 radial sensor locations as a function of $|A|^2$; a single fixed mode cannot explain this behavior even if its integrated amplitude is a complicated function of $|A|$. This indicates that the nonlinear axisymmetric self-interaction and higher antisymmetric harmonics play a significant role in altering the spatial structure of the axisymmetric deformation as the fluctuation amplitude changes. However, to a good approximation both the $m = \pm 1$ and $m = \pm 2$ Fourier components can be reasonably well-described by a single fixed mode with linear dependence on $|A|$. This is consistent with the typical assumption of the Stuart-Landau theory that the instability is a fixed eigenmode of the linear operator with variable eigenvalue. Higher wavenumbers show weak coherence with the order parameter; we therefore truncate the reconstruction at $|m| = 2$.

The conditional average can be viewed as an empirical approximation of the slow manifold related to the symmetry-breaking behavior. Fig. 9.2 visualizes this manifold by revolving the radially integrated axisymmetric deformation about the $\|\hat{p}_0\|$ -axis, while the parabolic “theoretical” manifold is generated from the weakly nonlinear $|A|^2$ scaling.

This conditional averaging approach can be viewed as an empirical approximation to the self-consistent mean-field model (Mantić-Lugo *et al.*, 2014), since it gives the fields at arbitrary fluctuation amplitudes without assuming a fixed spatial structure as in standard




Figure 9.6: Model for axisymmetric fluctuations. The conditional average on the order parameter A defines a slow manifold (left), but including an additional degree of freedom to account for fluctuations significantly improves the resolution of the $m = 0$ component (right).

modal analyses. In contrast to the self-consistent model, however, we do not neglect the influence of higher harmonics on the base flow. Since the modes are derived directly from experimental data, they naturally account for nonlinear interactions at all orders and can resolve arbitrary deformations of the axisymmetric component.

Residual projection for mean flow modification Since the symmetry-breaking is linked to an instability mode, we expect that antisymmetric modes will not significantly deform with amplitude, so $\psi_m(r, |A(t)|) \approx |A(t)|\psi_m(r)$ for $m \neq 0$. This is confirmed by the conditional averaging analysis above. However, the coherent part of the axisymmetric component is dominated by nonlinear deformations induced by Reynolds stresses, so it may have a complicated amplitude dependence.

For weakly nonlinear laminar flows, an asymptotic analysis indicates that axisymmetric deformations can be expressed as a function of the instability amplitude $|A(t)|^2$, pinning the state to the slow manifold (Sipp & Lebedev, 2007; Meliga *et al.*, 2009). However, the structure of models based on Galerkin projection onto a set of fixed spatial modes suggests that there may be some finite relaxation time to the equilibrium state determined by $|A(t)|$ (Noack *et al.*, 2003).

In other words, the conditional average defines an approximate slow manifold (illustrated in Fig. 9.2), though we would like to avoid assuming the state always resides on this surface. We address this by introducing an additional (real) degree of freedom $B(t)$ representing the difference between the axisymmetric field $\hat{p}_0(r, t)$ and the manifold field $\psi_0(r, |A(t)|^2)$. If this difference is typically small, the axisymmetric field can be expressed as a linearization about $|A(t)|$ by defining a new $m = 0$ mode $\psi_B(r, |A(t)|^2)$ from the gradient of the manifold:

$$\hat{p}_0(r, A, B) \approx p^0(r) + \psi_0(r, |A|^2) + 2|A|B\psi_B(r, |A|^2) \quad (9.8a)$$

$$\psi_B(r, |A|^2) = \left. \frac{\partial \psi_0}{\partial |A|^2} \right|_{|A(t)|^2}. \quad (9.8b)$$

This model for the axisymmetric part of the field is a generalization of the fixed “shift mode” model proposed by (Noack *et al.*, 2003). In particular, if the axisymmetric field

does have the weakly nonlinear scaling $\psi_0(r, |A(t)|^2) = |A(t)|^2 \psi_0(r)$ and the linearization is about the unstable fixed point $A = 0$, then the tangent field in (9.8) is just $\psi_0(r)$. For models based on Galerkin projection it is more natural to define a single coefficient for each mode. That is, in the models introduced by Noack *et al.* (2003) and Sieber *et al.* (2020), B is defined so that $\hat{p}_0(r, B(t)) \approx B(t)\psi_0(r)$. However, the proposed parametric expansion in the present work allows the model more flexibility to capture the natural variations of the flow, without assuming any scaling behavior.

The instantaneous value of this secondary order parameter $B(t)$ can be estimated by projecting the part of the axisymmetric field not correlated with $|A|$ onto the tangent space of the slow manifold (see Fig. 9.6):

$$B(t) \approx \frac{\int (\hat{p}_0(t) - \psi_0(t)) \psi_B(t) r dr}{\int \psi_B(t) \psi_B(t) r dr}, \quad (9.9)$$

where the explicit dependence on r and $|A(t)|$ has been omitted everywhere for notational clarity. In practice, if the axisymmetric mode $\psi_0(r, |A(t)|^2)$ is interpolated in $|A|^2$ with a spline, the tangent field $\psi_B(r, |A(t)|^2)$ can be computed with a derivative of the spline at each radial position r .

We emphasize that the conditional averaging has already resolved the axisymmetric deformations associated with the Stuart-Landau nonlinear stability mechanism. This secondary order parameter $B(t)$ is therefore not strictly necessary to describe the symmetry-breaking behavior, but it significantly improves the resolution of the model for the axisymmetric fluctuations, reducing the error in the modal approximation by 37%. Moreover, as shown below, the new degree of freedom $B(t)$ that allows for deviations from the slow manifold introduces the a finite relaxation timescale that can account for a previously unexplained peak in the symmetric power spectrum.

With this modal expansion, the base pressure field can be reconstructed with

$$p(r, \theta, t) = \underbrace{p^0(r) + \psi_0(r, |A|^2) + 2|A|B\psi_B(r, |A|^2)}_{\text{axisymmetric}} + \underbrace{|A| \sum_m \psi_m(r) e^{im(\theta + \phi(t))}}_{\text{antisymmetric}}, \quad (9.10)$$

where the time dependence of $A(t)$ and $B(t)$ has been omitted for brevity. The full model consists of seven modes: the unstable steady state p^0 , the symmetric deformation ψ_0 , the linearization of the manifold ψ_B , and the instability mode at wavenumber $m = \pm 1, \pm 2$. However, the model only has three degrees of freedom: the complex instability amplitude $A(t)$ and the real axisymmetric deviation from the manifold $B(t)$. The modal expansion (9.10) has a straightforward dependence on these degrees of freedom, so it will behave predictably even when evaluated for a state not in the original data set. All of

the modes, along with their parametric variation in A , can be directly computed from experimental observations. An example reconstructed field is shown in Fig. 9.3.

9.4 Reduced-order model of symmetry-breaking

We identify a stochastic model of the form of (9.1) by applying Langevin regression to the generalized modal coefficients $A(t)$ and $B(t)$, approximating the incoherent turbulent fluctuations as Gaussian white noise. This method optimizes unknown parameters of a proposed nonlinear stochastic model using forward and adjoint Fokker-Planck equations to enforce consistency with experimental statistics. If the form of the model is unknown, a simple reverse-greedy model selection procedure can identify the simplest model that explains the observed data (Callahan *et al.*, 2021b). This method explicitly treats inherent fluctuations of the system as process noise.

Reduced-order models of turbulent flows are generally challenging to construct for the same reason that the statistical perspective was so successful historically; signals are dominated by apparently random turbulent fluctuations. The most common approach to dealing with the fluctuations is to introduce a closure model that approximately captures their effects, for instance via an eddy viscosity term to achieve the correct dissipation, without explicitly treating the fast scales (Aubry *et al.*, 1988; Holmes *et al.*, 1996; Noack *et al.*, 2008).

A major difficulty in treating the fast fluctuations directly is that there is not a strict enough separation of scales to apply the machinery of near-equilibrium statistical mechanics. Although they cannot be formally justified, in practice some approximations based on statistical mechanics have been successful. For example, while Eulerian statistics are usually non-Gaussian, central limit theorem arguments suggest that both Lagrangian variables and integral quantities have near-normal distributions (Monin & Yaglom, 1971; Kraichnan & Chen, 1989). Similarly, although the turbulent fluctuations are correlated in time, various modeling methods have nevertheless been successful by approximating them as white noise (Majda *et al.*, 2001; McKeon & Sharma, 2010; Rigas *et al.*, 2015). Alternatively, recent work has investigated the use of colored noise in linearized flow models (Zare *et al.*, 2017; Towne, 2021), although classical statistical physics tools such as the Fokker-Planck equation, which is integral to our system identification method, cannot be applied in this case.

In previous work the evolution of the order parameter magnitude $|A(t)|$ has been successfully modeled by a stochastic cubic amplitude equation, inspired by the weakly nonlinear normal form (Rigas *et al.*, 2015). However, the weakly nonlinear analysis is predicated on a fixed-mode decomposition, which is at odds with the proposed amplitude-dependent spatial modes. Nevertheless, as described above a dynamical model with similar structure can also arise from the mean-field theory of symmetry-breaking phase transitions, which does not rely on fixed spatial modes or weak nonlinearity.

Nonlinear stochastic system identification We use the recently proposed Langevin regression method for identifying nonlinear stochastic models (Callaham *et al.*, 2021b) to identify the drift and diffusion functions rather than presuppose a form based on qualitative arguments. This method enables identification of a statistically consistent low-dimensional nonlinear generalized Langevin equation of the form (9.1). Langevin regression, a stochastic variant of the sparse identification of nonlinear dynamics (SINDy) approach (Brunton *et al.*, 2016a), optimizes free parameters of a Langevin model via solutions of both the forward and adjoint Fokker-Planck equations. The steady-state solution of the forward equation gives the steady-state probability distribution, while the adjoint solution gives the finite-time propagation statistics of the Langevin model (Risken, 1996; Zwanzig, 2001). Langevin regression simultaneously minimizes the discrepancy between the Fokker-Planck solutions and statistics computed from the experimental data. This system identification approach does not require access to the governing equations and can be applied to arbitrary quantities of interest. Details about the algorithm and numerical methods for solving the Fokker-Planck equations are given in (Soize, 1988) and the supplementary information of Callaham *et al.* (2021b).

Moreover, this approach can be combined with the simple reverse-greedy stepwise sparse regression (SSR) procedure for selecting a parsimonious, but maximally descriptive model from a set of candidates (Boninsegna *et al.*, 2018; Callaham *et al.*, 2021b). This algorithm sequentially discards functions whose exclusion is associated with the smallest increase in cost function. The Pareto-optimal model can then be chosen as the simplest model with a small cost function, as shown in Fig. 9.7. We use a library of polynomials consistent with the expected symmetries as candidate functions. The drift and diffusion functions are approximated by sparse linear combinations of these functions. For instance, the drift and diffusion libraries for A are

$$\Theta_f^A(A) = [A \quad A|A|^2 \quad A|A|^4 \quad A|A|^6] \quad (9.11a)$$

$$\Theta_\sigma^A(A) = [1 \quad |A| \quad |A|^2]. \quad (9.11b)$$

Then the approximate drift function is $f_A(A) = \Theta_f^A(A)\xi_f^A$, where ξ_f^A is a relatively sparse coefficient vector that identifies the library terms that are active in the dynamics. We do not include the axisymmetric deviation amplitude B because it is not associated with any symmetry-breaking in the flow and therefore would not appear in the effective potential for A . The rotational symmetry of the physical system implies that ξ_f^A should be purely real, but since Langevin regression is based on a least-squares method a small imaginary part will generally be retained. For the sake of a minimum-complexity model we enforce that ξ_f^A is real, although allowing complex-valued coefficients does not noticeably change the results.

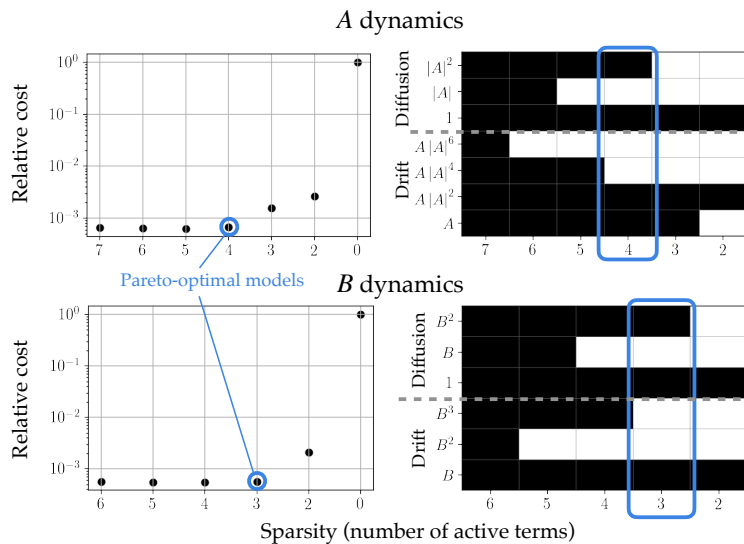


Figure 9.7: **Model selection with Langevin regression.** The drift and diffusion functions are sparse linear combinations from a library of candidates. We select the simplest models that are statistically consistent (small cost function). The vertical axes are scaled by the value of the cost function with an identically zero coefficient vector.

Likewise, the libraries for B are

$$\Theta_f^B(B) = [B \quad B^2 \quad B^3] \quad (9.12a)$$

$$\Theta_\sigma^B(B) = [1 \quad B \quad B^2]. \quad (9.12b)$$

Again, we do not expect strong coupling between the order parameters A and B , since B was defined primarily to capture fluctuations that were uncorrelated with A .

Decoupling the dynamics also reduces the maximum dimensionality of the regression problems from three dimensions to two, since the real and imaginary parts of A must be treated separately. Since Langevin regression is a histogram-based method where the Fokker-Planck equation is solved and evaluated on a grid, it does not scale well to higher dimensions. For problems with multiple coupled instability modes, histogram-free approaches such as Langevin inference (Brückner *et al.*, 2020; Frishman & Ronceray, 2020) may be more appropriate, although this does not enforce consistency with the steady-state probability distribution.

The two key parameters in this method are the finite sampling rate, which allows the fast turbulent fluctuations to appear approximately uncorrelated in the time series, and the weight of the Kullback-Leibler (KL) divergence between the empirical probability distribution and the steady-state solution of the Fokker-Planck equation. The latter controls the relative weight factor of the forward and adjoint Fokker-Planck solutions in the objective function. We choose sampling rates 200 times slower than the experimental sampling rate for A and 50 times slower for B and select the KL divergence weight so that the forward and adjoint Fokker-Planck solutions have approximately equal contributions in the optimization. The values are 10^{-1} and 10^2 for the A and B optimizations, respectively.

λ	μ	σ_A	γ_A	α	σ_B	γ_B
1.9	-1.9	$0.8 + 0.7i$	$0.3 + 0.3i$	-26.7	6.7	0.2

Table 9.1: Identified coefficients for the Langevin models (9.13).

Model analysis and evaluation Langevin regression identifies clear Pareto-optimal models for both the A and B dynamics, as shown in Fig. 9.7. That is, we can clearly select a model with minimal complexity in the sense that neglecting any more terms would cause the optimization loss function to significantly increase. The full identified model is

$$\dot{A} = \lambda A - \mu A |A|^2 + (\sigma_A + \gamma_A |A|^2) w_1(t) \quad (9.13a)$$

$$\dot{B} = -\alpha B + (\sigma_B + \gamma_B B^2) w_2(t), \quad (9.13b)$$

with coefficients given in Table 9.1.

The drift function for the order parameter A identified by the model selection has the form of a Langevin-Stuart-Landau equation, while the drift of the axisymmetric deformation B is linear relaxation. Physically, since B is defined as the difference between the instantaneous axisymmetric component of the field and that given by the conditional average on $|A|$, this models noisy relaxation towards the location on the slow manifold defined by the instantaneous value of A . This is consistent with the Fourier-decomposed Navier-Stokes equations; the axisymmetric field does not instantaneously reach the equilibrium corresponding to the amplitude of the antisymmetric instability, although the relaxation timescale is often considered negligible in weakly nonlinear laminar flows (Noack *et al.*, 2003; Meliga *et al.*, 2009).

The method also identifies quadratic state-dependent diffusion for both A and B . This is perhaps surprising, since the turbulent fluctuations are expected to be locally isotropic and therefore approximately independent of the large-scale motions (Pope, 2000). However, similar state-dependent diffusion arises in the case that slow macroscopic dynamics are averaged over fast degrees of freedom that are excited by stochastic forcing due to nonlinear coupling across scales (Majda *et al.*, 2001; Pradas *et al.*, 2012). In particular, the diffusion functions in (9.13) have the same form as a Taylor expansion of the diffusion derived by (Pradas *et al.*, 2012) for an unstable mode with quadratic coupling to stable modes driven by additive white noise. The state-dependent diffusion also allows the Langevin model to better resolve the long tails of the probability distributions, as observed by Callahan *et al.* (2021b).

To evaluate the model, we simulate the Langevin models and construct approximate pressure fields using the parametric modal expansion (9.8). This reconstruction returns from the three-dimensional state space of the Langevin model to the 64-dimensional space of the pressure measurements. We compare the full base pressure distributions in order to test the ability of the entire model, including both Eq. (9.10) and (9.13), to capture the

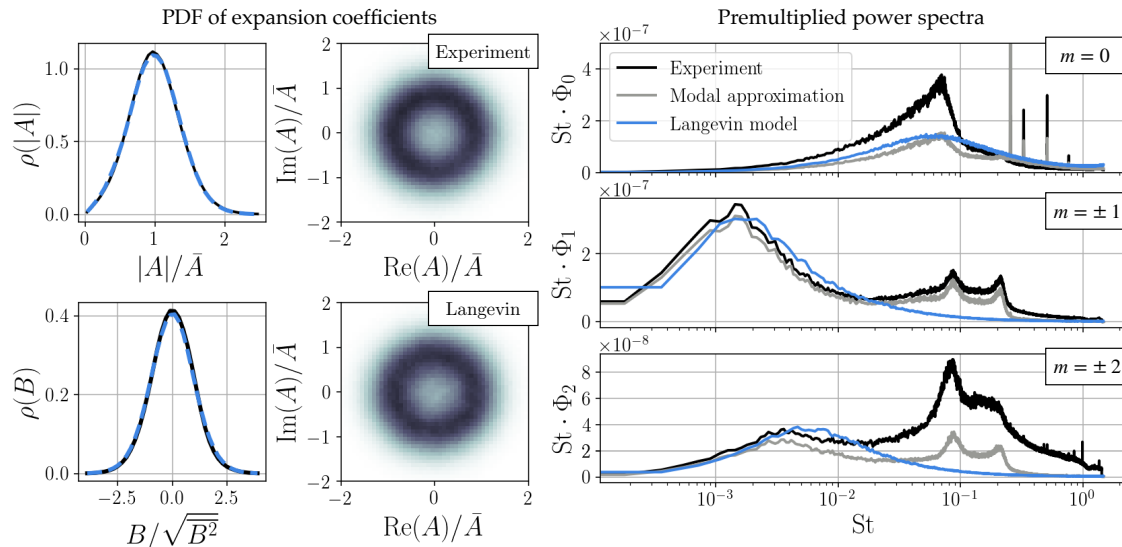


Figure 9.8: **Statistical evaluation of the model.** The mean-field modal expansion accurately captures most features of the radially averaged premultiplied power spectrum based only on the order parameters $A(t)$ and $B(t)$ (gray). Monte Carlo evaluation of the Langevin model (blue) shows that it reproduces the the empirical probability distributions and dominant frequency content at the leading azimuthal wavenumbers.

physical behavior.

The results are compared to both the experimental data and the reconstruction based on the modal expansion and experimental values of A and B in Fig. 9.8 using both the empirical probability distribution and the radially averaged premultiplied power spectra. We compare the premultiplied spectra $St \cdot \Phi_m(St)$, where $\Phi_m(St)$ is the estimated power spectral density of the radially integrated m -th Fourier component, rather than the power spectral density itself since the area under the premultiplied spectrum directly corresponds to signal energy. This makes it better suited for comparison of the dominant energy scales. Although the model is clearly too simplified to capture the full power spectrum, it does reproduce the dominant peak for the leading wavenumbers and accurately approximates the probability distributions.

The compact empirical model (9.13) therefore resolves the dominant physical mechanisms associated with symmetry-breaking in the turbulent wake, including linear instability, mean-field deformation, and the influence of higher harmonics. Since Eq. (9.13) gives the drift and diffusion in terms of simple analytic expressions, the behavior of the model can be fully characterized throughout the low-dimensional state space. For example, the Fokker-Planck equation can be used to either numerically compute the steady-state probability density or predict the evolution of an uncertain distribution of states. Thus, although the model is fit to a finite data set, its behavior in “unseen” regions (e.g. $|A| \gg \bar{A}$) can be completely quantified.

Critically, the model also reproduces the dominant frequency peak in the axisymmetric power spectrum, previously described as a “bubble-pumping mode” (Berger *et al.*, 1990; Rigas *et al.*, 2014; Boujo & Cadot, 2019), although it does not appear in any stability analysis (Rigas *et al.*, 2016). This model suggests that this spectral peak may instead be viewed as the result of a finite relaxation timescale towards the axisymmetric deformation induced by the amplitude of the instability mode. The spatial field associated with these dynamics is the gradient of the manifold defined by the conditional average. These relaxation dynamics are consistent with the underlying Navier-Stokes equations, but this timescale is often considered negligible for weakly nonlinear laminar flows.

9.5 Discussion

We have demonstrated a data-driven modeling approach to learning interpretable nonlinear models for fluid coherent structures, where multiscale turbulence is treated as state-dependent stochastic forcing. This approach has been used to develop a simplified mean-field model of the symmetry-breaking behavior in a turbulent axisymmetric wake. The empirical model comprises seven parametrically varying spatial modes and three dynamical degrees of freedom, and was constructed entirely from experimental observations. Using a phase-aligned conditional average with respect to the center of pressure, we have shown that the fixed-mode ansatz of standard modal decompositions cannot explain the mean-field deformation related to the symmetry-breaking instability of the turbulent axisymmetric wake. This reflects the higher-order influences as well as nonlinear self-interaction of the axisymmetric part of the flow, both of which are typically negligible for weakly supercritical laminar flows.

Modeling approaches based on weakly nonlinear approximations have proven highly successful in laminar flows. However, in this work we have shown that the extension of this perspective to turbulent flows is more subtle than adding stochastic forcing to the weakly nonlinear model. The amplitude scaling and structure of the axisymmetric deformation is inconsistent with the quadratic dependence implied by such an approach, although the proposed conditional averaging procedure can capture the natural behavior of the mean flow deformation.

Langevin regression identifies a simple Stuart-Landau-type model for the modal coefficients, as well as a quadratic nonlinear state-dependent noise model. This form of diffusion is consistent with analysis of fast-slow systems with quadratic nonlinearities where only the fast scales are stochastically forced (Majda *et al.*, 2001; Pradas *et al.*, 2012). Monte Carlo evaluation of the model matches the stationary probability distributions of the experimental data and resolves the dominant peak in the power spectrum at the leading azimuthal wavenumbers. Moreover, the model is simple and interpretable, yielding physical insights into physical mechanisms including the mean flow deformation and low-frequency modulation of the recirculation bubble.

Beyond the context of the axisymmetric wake, these results support the parameteri-

zation of turbulent fluctuations as stochastic forcing of the quasi-deterministic coherent structures evolving near a slow manifold, at least as an approximation for empirical models. We emphasize that this description relies on a strict separation of scales, which is known to be absent in turbulent flows. Even with this caveat, this simplification is appealing enough for engineering applications such as closed-loop flow control that it may be useful even if it only holds in an approximate sense. For example, a feedback controller based on a nonlinear Langevin model has been shown to produce power-efficient drag reduction on an Ahmed body with a similar symmetry-breaking instability (Brackston *et al.*, 2016).

More broadly, we hope that this work illustrates a principled approach to constructing stochastic reduced-order models from limited experimental observations of a turbulent flow. Although much progress has recently been made in developing stochastic models of turbulent flows using the linearized governing equations (Zare *et al.*, 2017; Towne, 2021) and with empirical models of experimental data (Rigas *et al.*, 2015; Sieber *et al.*, 2020; Herrmann *et al.*, 2020), there are many open questions about how the heuristics of low-dimensional models of weakly nonlinear flows will translate to fully-developed turbulence. In this work we have chiefly focused on the mean-field deformation associated with the symmetry-breaking bifurcation, but recent studies have highlighted the role of higher-order triadic nonlinear interactions in capturing the dynamics of both natural (Schmidt, 2020) and actuated (Herrmann *et al.*, 2020) turbulent wakes. Capturing the interactions between instability modes (for instance, symmetry-breaking and vortex-shedding) may also prove critical in developing models suitable for active flow control.

As fully empirical data-driven methods continue to grow in popularity and utility, ensuring consistency with the underlying physics remains a central challenge. The model proposed in this work captures many of the essential statistical features of the data and leads to new hypotheses about the behavior of the axisymmetric wake in particular and globally unstable turbulent flows in general.

Part V

DOMINANT BALANCE

Chapter 10

DATA-DRIVEN DOMINANT BALANCE MODELING

It is well known across the engineering and physical sciences that persistent behaviors in complex systems are often determined by the balance of just a few dominant physical processes. This heuristic, which we refer to as dominant balance, has played a pivotal role in our study of systems as diverse as turbulence (Holmes *et al.*, 1996), geophysical fluid dynamics (Gill, 1982; Lighthill, 1966), and fiber optics (Blow & Wood, 1989). It is also thought to play a role in the emerging fields of pattern formation (Cross & Hohenberg, 1993; Morris *et al.*, 1993; Grzybowski *et al.*, 2000), wrinkling (Cerde & Mahadevan, 2003), droplet formation (Shi *et al.*, 1994), and biofilm dynamics (Seminara *et al.*, 2012). These balance relations provide reduced-order mechanistic models to approximate the full complexity of the system with a tractable subset of the physics.

The success of dominant balance models is particularly evident in the field of fluid mechanics. The Navier-Stokes equations describe behavior across a tremendous range of scales, from water droplets to supersonic aircraft and hurricanes. Thus, much of our progress has required simplifying the physics with nondimensional parameters that determine which terms are important for a specific problem. Perhaps the most well-known dimensionless quantity, the Reynolds number, embodies the balance between inertial and viscous forces in a fluid. Other nondimensional numbers capture the relative importance of inertial and Coriolis forces (Rossby number), inertia and buoyancy (Froude number), and thermal diffusion and convection (Rayleigh number), among dozens of other possible effects. In many situations, the magnitude of these coefficients determine the important mechanisms at work in a flow; conversely, they determine which mechanisms may be safely neglected. In geophysical flows, balance arguments bypass the incredible complexity of the ocean and atmosphere to identify driving mechanisms such as geostrophy, the thermal wind, Ekman layers, and western boundary currents (Lighthill, 1966; Gill, 1982). Lighthill, one of the most influential fluid dynamicists of the 20th century, often relied on dominant balance arguments as physical motivation for his mathematical analyses (Lighthill, 1952, 1966). Beyond fluid mechanics, asymptotic methods have been crucial in characterizing a diverse range of physical behavior.

Advanced statistical tools now allow analysis of the increasing wealth of data from modern experimental and numerical methods, but to date there is no direct link between this data and the powerful insights of asymptotic scaling analysis. This presents an exciting opportunity to leverage data-driven methods, which are driving changes in a wide range of fields, from control (Pastoor *et al.*, 2008; Verma *et al.*, 2018) to turbulence modeling (Duraismy *et al.*, 2019), forecasting (Lguensat *et al.*, 2017), and extreme event predic-

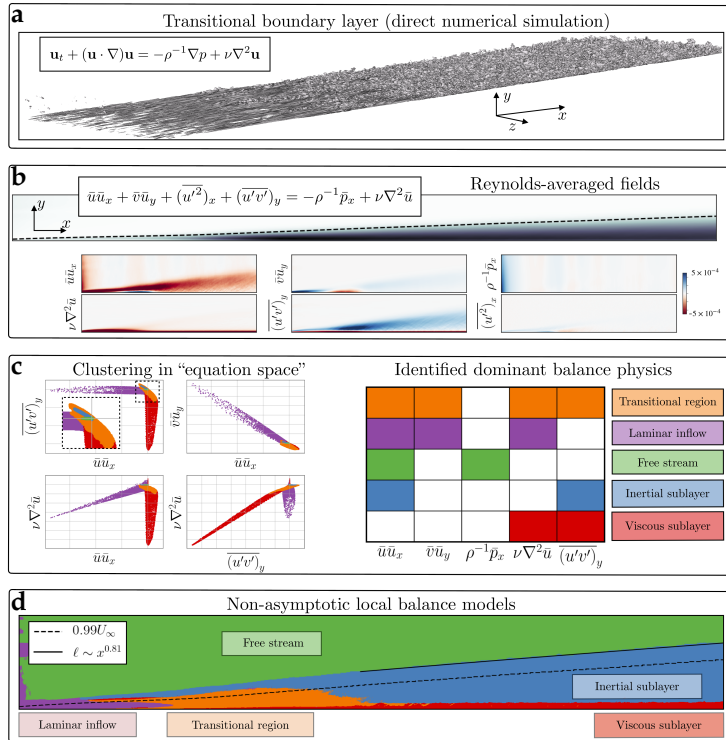


Figure 10.1: Schematic of the dominant balance identification procedure applied to a turbulent boundary layer. High-resolution direct numerical simulation results (**a**, visualized with a turbulent kinetic energy isosurface) are averaged to compute the terms in the Reynolds-averaged Navier-Stokes equations (**b**). The equation space representation of the field enables clustering and sparse approximation methods to extract the distinct geometrical structures in the six-dimensional space corresponding to dominant balance physics (**c**). Finally, the entire domain can be segmented according to these interpretable balance models, identifying distinct physical regimes (**d**). A curve fit to the wall-normal extent of the post-transition region of the identified inertial sublayer shows an approximate scaling of $\ell \sim x^{0.81}$, consistent with the theoretical prediction of $x^{4/5}$ from boundary layer theory. The 99 % free-stream velocity (U_∞) contour is also shown for reference.

tion (Wan *et al.*, 2018). Although some studies have addressed the dominant balance problem by using expert knowledge to design application-specific clustering algorithms (Portwood *et al.*, 2016; Lee & Zaki, 2018) or a post hoc interpretation of unsupervised clustering in terms of dominant balance (Sonnewald *et al.*, 2019), to our knowledge the general challenge of identifying local dominant balance regimes directly from data remains open.

In this work, we develop a generalized data-driven method to identify dominant balance regimes in complex physical systems. Figure 10.1 demonstrates the method applied to fluid flow over a flat plate in transition to turbulence. We introduce a geometric perspective on dominant balance in which standard machine learning tools can automatically identify dominant physical processes. The geometric approach naturally links the analysis to the underlying equation so that the entire procedure can be easily interpreted and visualized. This data-driven method is designed to be applied in tandem with, rather than supplant, classical asymptotic analysis; the flexibility and generality of this combination extends balance modeling to a broader range of systems.

Our approach begins with a governing equation, which might be derived from fundamental physics (e.g. Maxwell’s equations or the Navier-Stokes equations) but could also result from a model discovery procedure (Schmidt & Lipson, 2009; Brunton *et al.*, 2016a; Rudy *et al.*, 2017). These governing equations are physical models capable of describing

a wide range of phenomena. However, it is well understood that the full complexity of such models is not always necessary to describe the local behavior of a system. In many regimes the dynamics are governed by just a subset of the terms involved in the global description.

We introduce the idea of an equation space, where each coordinate is defined by one of the terms in the governing equation. Each term may be evaluated individually at any point in space and time, resulting in a vector with each entry corresponding to a term in the governing equation. We define a dominant balance regime as a region where the evolution equation is approximately satisfied by a subset of the original terms in the equation; the remaining terms may be safely neglected. When a point in the field is approximately in dominant balance, the equation space representation of the field will have near-zero entries corresponding to negligible terms. Clearly, the equation space representation of a field is not unique; a fluid flow might be represented by velocity, vorticity, or streamfunction, for example. The interpretation of the dominant physics therefore depends on the choice of an appropriate governing equation for the application.

Dominant balance physics thus has a natural geometric interpretation in equation space, allowing standard machine learning tools to automatically identify regions where groups of terms have negligible contributions to the local dynamics. From this perspective, a dominant balance regime is characterized by a cluster of points that have significant covariance in directions of equation space corresponding to active physical processes. The covariance structure of this cluster is sparse in the sense that there is weak variation in directions that represent the negligible terms. This corresponds to the mathematical condition that the governing equation is approximately satisfied by a subset of its terms in a local region.

While such dominant balance regimes might be identified by many possible algorithms, we choose to cluster the data using Gaussian mixture models (GMM) (Bishop, 2006) and then extract a sparse approximation to the direction of maximum variance for each cluster using sparse principal components analysis (SPCA) (Zou *et al.*, 2012). We take the active terms in each cluster to be those that correspond to nonzero entries in the sparse approximation to the leading principal component.

In simple cases, this two-step GMM-SPCA procedure may be equivalent to applying a hard threshold, where a term is considered active if it exceeds some small value. However, our approach considers the local, relative importance of terms, whereas thresholding describes global, absolute importance. This distinction is important in multiscale systems where the scale of the dynamics varies significantly throughout the domain.

The data-driven approach to dominant balance analysis generalizes traditional methods in several critical directions. First, it does not rely on any explicit assumption of asymptotic scaling. Second, the clustering method yields pointwise estimates of the spatiotemporally local dominant balance not afforded by traditional scaling analysis in complex geometries. Third, while many dominant balance regimes have been proposed or assumed based on heuristic or intuitive arguments, this method provides an objective,

reproducible approach to testing these hypotheses. Finally, the probabilistic Gaussian mixture modeling framework is fully compatible with the relative nature of dominant balance analysis, providing natural estimates of uncertainty in the identified balance

This chapter includes an overview of the proposed method (Sec. 10.1) along with example applications to the unsteady cylinder wake (Sec. 10.2) and the mean profile of a boundary layer in transition to turbulence (Sec. 10.3). For further analysis and additional examples, including a geophysical flow and a combustion analog, see the original work by Callahan *et al.* (2021a).

10.1 Unsupervised dominant balance identification

In many fields of physics, painstaking analyses have produced models that are capable of describing a wide range of physical phenomena. However, it is well understood that the full complexity of such models is not always necessary to describe the local behavior of a system. We find that in many regimes the dynamics are governed by just a subset of the terms involved in the global description. A general evolution equation for the field $u(x, t)$ on the domain $(x, t) \in \mathcal{D}$ can be written as

$$\mathcal{N}(u) = \sum_{i=1}^K f_i(u, u_x, u_{xx}, \dots, u_t, \dots) = 0. \quad (10.1)$$

For example, the viscous Burgers' equation is

$$\mathcal{N}(u) = u_t + uu_x - \nu u_{xx} = 0 \quad (10.2)$$

We represent the equation in this implicit form rather than the typical $u_t = \mathcal{F}(u)$ form for two reasons. First, it includes arbitrary PDEs which are not easily expressed in the standard form. Second, this form highlights the fundamental balance of the equation; all terms must sum to zero. If some subset is dominant, the rest must be relatively small. We assume the total number of terms K is known, either from the physical model or as the result of a model selection procedure (Mangan *et al.*, 2017).

Consider an "equation space" where each coordinate is defined by one of the K terms in Eq. (10.1). At each point (x, t) in space and time, each of the K terms f_i in the governing equations (10.1) may be evaluated at $u(x, t)$, resulting in a vector $\mathbf{f} \in \mathbb{R}^K$:

$$\mathbf{f}(x, t) = [f_1(u(x, t), \dots) \quad f_2(u(x, t), \dots) \quad \cdots \quad f_K(u(x, t), \dots)]^T. \quad (10.3)$$

By construction, $\mathbf{1}^T \mathbf{f}(x, t) = \mathcal{N}(u) = 0$ for all $(x, t) \in \mathcal{D}$. Simulated or measured field data is typically discretized, so the domain is approximated by N spacetime points: $\mathcal{D} \approx \{(x, t)^j \mid j = 1, 2, \dots\}$. The field at each of these points corresponds to a point in equation space.

We define a dominant balance regime as a region $\mathcal{R} \subset \mathcal{D}$ where the evolution equation

is approximately satisfied by a subset of $p < K$ of the original terms in the equation; the remaining terms may be neglected. In this case $f(x, t)$ will have near-zero entries corresponding to negligible terms when $(x, t) \in \mathcal{R}$. Geometrically, the field is approximately restricted to p of the original K dimensions of the equation space, resulting in a subspace that is aligned with the active p terms.

This geometric perspective on dominant balance physics leads naturally to segmentation via unsupervised clustering. For example, the Gaussian mixture model (GMM) framework learns a probabilistic model by assuming the data are generated from a mixture of Gaussian distributions with different means and covariances (Bishop, 2006). The learned covariances for each cluster can then be interpreted in terms of active and inactive terms in the evolution equation. The N spacetime points in \mathcal{D} are used to train a mixture model; the algorithm treats points from a dominant balance regime as if they were generated from a distribution with near-zero variance in the directions corresponding to negligible terms. Data beyond the original inputs can efficiently be assigned to a balance model using the trained GMM.

In practice, there is no reason to expect the points will even approximate a mixture of Gaussian distributions. We therefore expect that the number of clusters required to capture all of the relevant physics will exceed the number of distinct balance regimes, resulting in redundant clusters. Furthermore, there is some ambiguity in the interpretation of “near-zero variance”. We address both of these issues using sparse principal components analysis (SPCA) (Zou *et al.*, 2012), which uses ℓ_1 regularization to extract a sparse approximation to the leading principal component. If a cluster describes a dominant balance regime, it should be well-described by its direction of maximum variance. Moreover, this leading principal component should have many near-zero entries. We apply SPCA to the set of points in each GMM cluster and take the active terms in the cluster to be those which correspond to nonzero entries in the sparse approximation to the leading principal component. The number of models can then be reduced by grouping clusters with the same set of active terms (or equivalently, the same sparsity pattern in the SPCA approximation).

Dominant balance identification can be seen as a localized active subspace analysis in equation space (Constantine *et al.*, 2014). Rather than assuming that there is a global decomposition into approximately active and inactive subspaces, we simultaneously search for subspaces corresponding to different balance relations and the regions of the domain where the dynamics are well-described by this subspace.

Burgers’ equation One of the simplest models that demonstrates dominant balance is the viscous Burgers’ equation, shown in Fig. 10.2. Shocks form from the nonlinear advection and are dissipated by the viscous term. Away from the shock front, however, the gradients of the field are relatively weak, so viscosity does not contribute significantly to the dynamics. Figure 10.2 demonstrates the balance identification procedure applied to a snapshot of the viscous Burgers’ equation example. Most of the field is classified into two

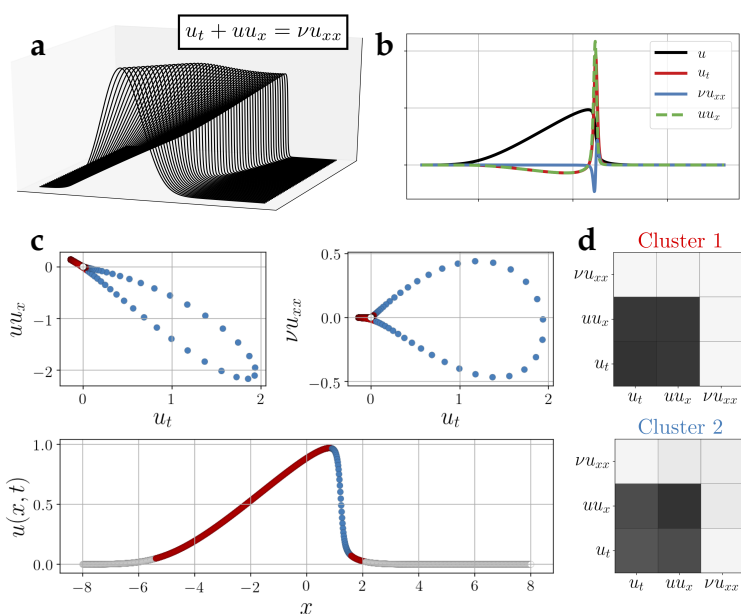


Figure 10.2: **Dominant balance of the viscous Burgers' equation.** The fields highlight identified dominant balance (a), with constituent terms shown in (b). The viscous term acts to diffuse sharp gradients and prevent formation of a discontinuous shock, but away from the shock front the dynamics are essentially inviscid. Away from the shock front, the field is approximately restricted to the $\nu u_{xx} = 0$ plane (c). This is reflected in the covariance matrices learned by the Gaussian mixture model (d).

clusters, corresponding to either no dynamics or an inviscid balance between acceleration and advection. Only a narrow slice along the shock front belongs to a cluster in which viscosity is active.

In simple cases, this two-step GMM-SPCA procedure might be replaced with a hard threshold; if a term exceeds some value ϵ it is “on”. However, the proposed method offers two main advantages over thresholding. First, the idea of dominant balance has a natural geometric interpretation in equation space, thereby avoiding setting an arbitrary threshold for which diagnostics and interpretation may not be straightforward. Second, our method considers the *local, relative* importance of terms, whereas thresholding describes *global, absolute* importance. For example, this distinction is significant in multi-scale systems with some background process underlying intermittent bursts of activity. The intermittency is dominated by a balance between terms which may be much larger than the background process, although the dynamics during quiescent periods would be determined primarily by the background process. In this case an absolute thresholding method would either choose the background process to be always on or always off, whereas a relative approach recognizes that the dominant local balance simply changes during the intermittent activity.

10.2 The unsteady cylinder wake

Flow past a cylinder at moderate Reynolds number is a prototypical flow configuration for bluff body wakes. The wake transitions from steady laminar flow to periodic vortex shedding via a Hopf bifurcation at $Re \approx 46$. The transition from linear instability to a stable limit cycle is itself a fascinating example of dominant balance in fluid mechanics and dynamical systems. The quadratic nonlinearity, initially inactive in the linear regime, mediates energy transfer between the mean flow and instability modes, deforming both until an energy balance is reached in the periodic limit cycle. This nonlinear stability mechanism was first described by Stuart (1958); Landau & Lifshitz (1959) and later employed for reduced-order modeling (Noack *et al.*, 2003). Even in the stable limit cycle, however, the local dynamics of the flow vary widely throughout the domain, highlighting mechanisms that give rise to von Kàrmàn-type vortex streets in a wide variety of flows.

We simulate this configuration at $Re = 100$ with unsteady incompressible DNS using the open source spectral element solver Nek5000 (Fischer *et al.*, 2008). The domain consisted of 17,432 seventh order spectral elements on $x, y \in (-20, 50) \times (-20, 20)$, refined close to a cylinder of unit diameter centered at the origin. Diffusive terms are integrated with third order backwards differentiation, while convective terms are advanced with a third order extrapolation. The results of this simulation have been validated against those of the immersed boundary projection method (Taira & Colonius, 2007) by comparing aerodynamic coefficients and vortex shedding frequency. We extract the vorticity field and spatial terms in equation (10.5) directly from the solver for further analysis. Time derivatives for dominant balance identification were estimated with a second order

central difference.

Governing equations and analytic scaling This unsteady, incompressible, viscous flow is governed by the two-dimensional Navier-Stokes equations:

$$\tilde{\mathbf{u}}_t + (\tilde{\mathbf{u}} \cdot \nabla)\tilde{\mathbf{u}} = -\frac{1}{\rho}\nabla\tilde{p} + \nu\nabla^2\tilde{\mathbf{u}} \quad (10.4a)$$

$$\nabla \cdot \tilde{\mathbf{u}} = 0, \quad (10.4b)$$

where $\tilde{\mathbf{u}}$ is the velocity field, \tilde{p} is the pressure, ρ is the density, and ν is kinematic viscosity. Of course, these equations themselves involve some degree of approximation, ignoring effects such as compressibility and gravity, making use of the Newtonian form of the stress tensor, and assuming Fickian diffusion, though they have proven highly accurate when applied in the correct regime. Nevertheless, there are distinct regimes in this simple wake flow.

For the wake behind a circular cylinder, the most relevant scales are the cylinder diameter L and free-stream velocity U . Dimensional analysis then suggests that

$$\tilde{\mathbf{u}} \sim U, \quad \tilde{p} \sim \nu U^2, \quad \nabla(\cdot) \sim \frac{1}{L}, \quad \frac{\partial}{\partial t}(\cdot) \sim \frac{U}{L}.$$

Nondimensionalizing with respect to these scales, we find that the viscous term is smaller than the others by a factor of the Reynolds number, $Re = UL/\nu$, resulting in the familiar nondimensional form of the Navier-Stokes equations:

$$\mathbf{u}_t + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\nabla p + \frac{1}{Re}\nabla^2\mathbf{u} \quad (10.5a)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (10.5b)$$

The variables and operators have been nondimensionalized according to the previous scales. For even moderately large Reynolds numbers, we would expect the flow to behave in an approximately inviscid manner away from the cylinder. Thus, structures formed in the near-wake region will be advected downstream by the mean flow with only weak dissipation, as observed in the vortex street.

Near the cylinder, the no-slip boundary conditions due to viscosity change the behavior qualitatively. If we examine the flow at a point a distance $\delta \ll L$ from the wall, then δ is a more appropriate length scale for the gradients. However, since the near-wall flow varies on a similar timescale to the wake, suppose that U/L is still a good scale for the time derivative. The various terms then scale as

$$\tilde{\mathbf{u}}_t \sim \frac{U^2}{L}, \quad (\tilde{\mathbf{u}} \cdot \nabla)\tilde{\mathbf{u}} \sim \frac{U^2}{\delta}, \quad -\frac{1}{\rho}\nabla\tilde{p} \sim \frac{U^2}{\delta}, \quad \nu\nabla^2\tilde{\mathbf{u}} \sim \frac{\nu U}{\delta^2}.$$

We find that the acceleration term is now smaller by a factor of δ/L , and expect the viscous term to be balanced by advection and the pressure gradient. The relatively strong gradients near the wall give rise to the vortex structures which characterize the wake.

Identified dominant balance Figure 10.3 shows an example vorticity field along with views of the 4D equation space corresponding to Eq. (10.5). Although the method treats space and time equivalently, here we freeze time and explore a single snapshot; since the flow is periodic we expect the results to be representative. The visualization in equation space clearly reveals signatures of balance relations. One set of GMM clusters is nearly restricted to the the zero-viscosity plane, while another has reduced variance in the acceleration direction. The sparse approximations to the leading principal components of each cluster confirms this intuition; we use SPCA to construct balance models by grouping the Gaussian models with non-negligible variance in the same directions. As expected, the far wake is approximately inviscid, while the region near the cylinder is dominated by a balance between viscosity, pressure, and advection. This method also identifies other approximate regions, such as a low-pressure-gradient balance between acceleration and advection (blue), slowly varying potential flow (green), and a far-field region with near-zero dynamics (white).

Nonlinear stability The cylinder wake at moderate Reynolds number is of particular interest in the reduced-order modeling community because it is a canonical example of a self-limiting instability exhibiting the Stuart-Landau nonlinear stability mechanism (Noack *et al.*, 2003; Sipp & Lebedev, 2007; Mantič-Lugo *et al.*, 2014). The steady-state solution (\mathbf{u}_0, p_0) , defined as¹

$$(\mathbf{u}_0 \cdot \nabla)\mathbf{u}_0 = -\nabla p_0 + Re^{-1}\nabla^2\mathbf{u}_0, \quad (10.6)$$

becomes unstable to infinitesimal perturbations at $Re_c \approx 46$. At this critical parameter value, the flow undergoes a Hopf bifurcation (Jackson, 1987; Provansal *et al.*, 1987); the asymptotic solution is the limit-cycle vortex shedding explored above.

The mechanism by which the exponential energetic growth of the linear instability transitions to a stable limit cycle may be understood in terms of the Stuart-Landau mean field theory. An arbitrary flow field may be decomposed into the steady-state “base” flow \mathbf{u}_0 and a time-varying perturbation \mathbf{u}' :

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{u}_0(\mathbf{x}) + \mathbf{u}'(\mathbf{x}, t). \quad (10.7)$$

A similar expansion can be applied to the pressure field. Expanding the momentum equa-

¹All fields are understood to be divergence-free; here we give only the momentum equations for brevity.

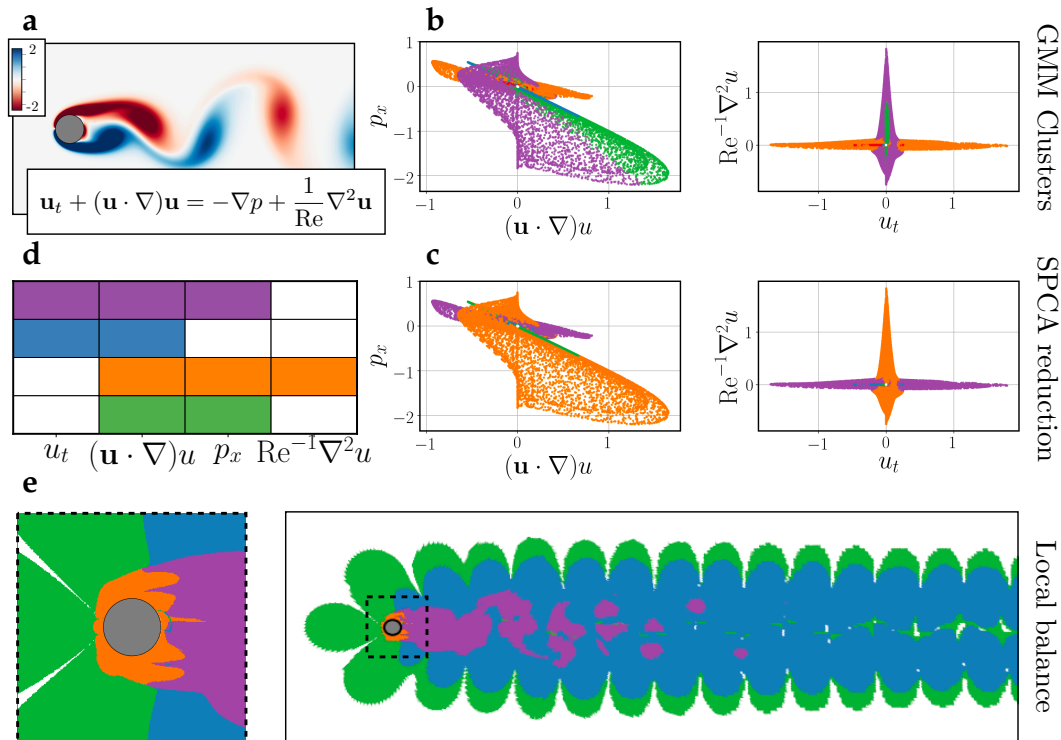


Figure 10.3: **Vorticity snapshot for the wake behind a cylinder at $Re = 100$ (a).** A Gaussian mixture model (GMM) assigns field points to clusters by looking for groups with distinct mean and covariance (b). For instance, some clusters vary mainly in the acceleration-advection directions, while others vary principally in the viscous-advection directions. We would expect these to represent the far-field and boundary regions, respectively. This is confirmed by the sparse principal components analysis (SPCA) reduction, where clusters with significant nonzero variance in the same directions are grouped together (c). These directions can be interpreted as active terms in the balance relation (d). As anticipated, the region near the cylinder is dominated by a balance between viscosity and advection and pressure forces, while the far wake is approximately inviscid (e).

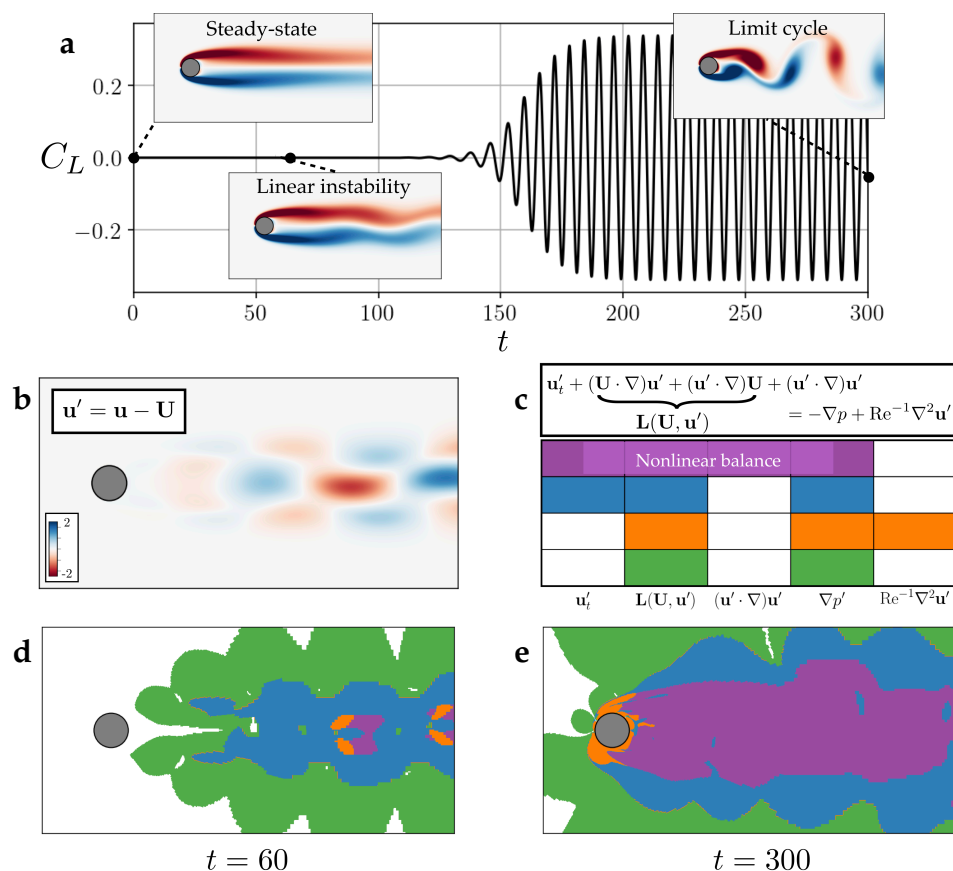


Figure 10.4: The role of nonlinearity in stabilizing a von Kàrmàn vortex street. The transient flow evolves from an unstable steady-state through the exponential growth of a linear instability mode to the vortex shedding limit cycle (a, visualized with lift coefficient C_L). Only one balance is identified that includes the nonlinear term $(\mathbf{u}' \cdot \nabla)\mathbf{u}'$. This balance does not appear significantly in the linear growth regime (d), consistent with the correspondence between the fluctuations and the instability mode. The fully saturated limit cycle is dominated by this balance, however, confirming the interpretation of the stabilizing feedback loop mediated by the nonlinearity.

tions, we find that the fluctuations evolve according to

$$\mathbf{u}'_t + (\mathbf{u}_0 \cdot \nabla)\mathbf{u}' + (\mathbf{u}' \cdot \nabla)\mathbf{u}_0 + (\mathbf{u}' \cdot \nabla)\mathbf{u}' = -\nabla p' + Re^{-1}\nabla^2\mathbf{u}'. \quad (10.8)$$

The only part of these equations that is nonlinear in \mathbf{u}' is the advection term $(\mathbf{u}' \cdot \nabla)\mathbf{u}'$. Linear stability analysis proceeds by assuming the fluctuations are weak enough that this term is negligible.

However, the exponential growth of unstable modes implies that eventually this assumption cannot hold; in order for the energy of the flow to be bounded at long times the nonlinearity must play a stabilizing role. This is typically conceptualized as a mean field deformation. The mean of the limit cycle has a much shorter recirculation region than the steady-state solution and is approximately neutrally stable (Barkley, 2006). The role of the nonlinearity is thus understood as a feedback mechanism. The base flow is deformed in a manner that reduces the growth rate of the instability. The two come into energetic balance on the limit cycle, where the mean flow becomes neutrally stable; this is the basis for the “self-consistent” mean field modeling approach (Mantić-Lugo *et al.*, 2014). Although this argument is based on a linear perturbation analysis of a smoothly deforming base flow, it can be confirmed more rigorously close to the critical Reynolds number with a weakly nonlinear expansion (Sipp & Lebedev, 2007).

The role of the nonlinearity can also be examined from the perspective of dominant balance, without assuming either linearity or proximity to the bifurcation. Figure 10.4 shows the transient evolution of the cylinder wake, from the unstable steady-state through the linear growth region and ultimately to the vortex shedding limit cycle. We apply dominant balance analysis to the evolution equation for the base-subtracted snapshots, Eq. (10.8) in both the linear instability regime and the post-transient limit cycle. The method identifies four dominant balances: a far-field steady balance between base flow advection and the pressure gradient (green), a steady viscous region (orange), an inviscid linear balance (blue), and finally a balance including the nonlinear term (purple).

The nonlinear balance is largely absent from the transient growth snapshot. This is expected from the assumptions of linear stability analysis; the fluctuations at this stage are nearly identical to the linear instability mode (Barkley, 2006). On the other hand, the wake region of the post-transient snapshot is largely nonlinear, which confirms the picture of the role of nonlinearity in stabilizing the limit cycle.

Moreover, this analysis clearly demonstrates the instantaneous local balance of the terms, without relying on either assumptions of linearity or limiting parameter regimes. This distinguishes it from analyses such as weakly nonlinear expansions, self-consistent modeling, and energy balance arguments. Furthermore, this method is non-intrusive and could be applied to either experimental or numerical observations with more complex decomposition structures such as the harmonic balance expansion approach to nonlinear resolvent analysis (Rigas *et al.*, 2021).

10.3 A turbulent boundary layer

One of the major breakthroughs in the study of fluid mechanics in the 20th century was the development of boundary layer theory (Schlichting, 1955; Pope, 2000). In many practical applications fluids can be treated as inviscid, but close to solid boundaries strong velocity gradients lead to significant viscous forces. Prandtl showed in 1904 that careful scaling analysis applied to the governing Navier-Stokes equations reveals distinct regimes where the behavior of the fluid is essentially determined by a small subset of the full equations. In turn, these balance relations can be used to derive powerful scaling laws such as the so-called “law of the wall”.

Although such analyses can be intractable for general turbulent flows, one of the most important canonical configurations is zero pressure gradient flow over a flat plate parallel to the free stream velocity. The zero pressure gradient ensures that the free-stream velocity is constant in the streamwise direction at large distances from the wall. This flow is statistically two-dimensional; the configuration does not vary in the cross-stream direction so the mean flow only varies in the streamwise and wall-normal directions.

To study dominant balance physics in the turbulent boundary layer, we use the transitional DNS by Lee and Zaki (Zaki, 2013; Lee & Zaki, 2018; Wu *et al.*, 2019), openly available from the Johns Hopkins Turbulence Database (Perlman *et al.*, 2007; Li *et al.*, 2008)². The full computational domain consists of a long flat plate with an elliptical leading edge. The extent of the domain (in units defined by the plate half-thickness) is $(x, y, z) \in (1040, 40, 240)$ with periodic boundary conditions in the spanwise (z) direction, discretized to $(Nx, Ny, Nz) = (4097, 257, 2049)$. Since the configuration of interest is a zero pressure gradient flat plate boundary layer, the DNS results are only saved once the flow passes the elliptical leading edge ($x > 30.2185$). The inflow consists of small amplitude free-stream turbulence superimposed on a uniform streamwise velocity U_∞ incident on the plate. The interactions of these perturbations with the laminar boundary layer cause a downstream transition to turbulence (Zaki, 2013).

Since we are interested here in the mean momentum balance, we only use the 2D mean field (also available from JHTDB), which was computed from 4701 data snapshots once the flow reached a statistically stationary state. Without direct access to the gradients, we compute the constituent terms of the RANS equations with second-order accurate finite differences, as shown in Fig. 10.1b. Although some of these fields show small fluctuations, the overall smoothness suggests the statistics are approximately converged.

Governing equations and analytic scaling After performing the Reynolds decomposition of the variables into mean and fluctuating components, e.g. $u = \bar{u} + u'$, the mean flow is determined by the Reynolds-averaged Navier-Stokes (RANS) equations. For the

²<https://doi.org/10.7281/T17S7KX8>

streamwise mean velocity \bar{u} , the equation is

$$\bar{u} \frac{\partial \bar{u}}{\partial x} + \bar{v} \frac{\partial \bar{u}}{\partial y} = \rho^{-1} \frac{\partial \bar{p}}{\partial x} + \nu \nabla^2 \bar{u} - \frac{\partial}{\partial y} \overline{u'v'} - \frac{\partial}{\partial x} \overline{u'^2}. \quad (10.9)$$

The terms on the left represent mean flow advection, while those on the right are the pressure gradient, viscosity, wall-normal Reynolds stress, and streamwise Reynolds stress, respectively.

One of the challenges in studying this flow is that there are multiple length scales. Following (Holmes *et al.*, 1996), we may consider a streamwise length scale L , a wall-normal length scale ℓ , and a viscous length scale $\eta = \nu/u_\tau$, where u_τ is the “friction velocity” associated with the shear stress at the wall.

Beginning with the “outer” region of the boundary layer (where $y \gg \eta$), suppose the mean streamwise velocity \bar{u} scales with the free stream U_∞ , while the turbulent fluctuations u', v' scale with u_τ . As with the previous example, assume that the derivatives scale with the corresponding length scale, so that for instance $(\cdot)_y \sim 1/\ell$. For instance, the continuity equation $\bar{u}_x + \bar{v}_y = 0$ implies that $\bar{v} \sim U_\infty(\ell/L)$. By this reasoning typically we would expect the mean velocity gradient \bar{u}_y to scale with U_∞/ℓ , but as argued in (Holmes *et al.*, 1996), the gradients in the outer part of the layer are much weaker than near the wall, and empirically a better estimate is $\bar{u}_y \sim u_\tau/\ell$. Then for the streamwise momentum equation we find

$$\bar{u}\bar{u}_x \sim \frac{U_\infty^2}{L}, \quad \bar{v}\bar{u}_y \sim \frac{u_\tau U_\infty}{L}, \quad \nu \bar{u}_{xx} \sim \frac{\nu U_\infty}{L^2}, \quad \nu \bar{u}_{yy} \sim \frac{\nu u_\tau}{\ell^2}, \quad (\overline{u'v'})_y \sim \frac{u_\tau^2}{\ell}, \quad (\overline{u'^2})_x \sim \frac{u_\tau^2}{L},$$

and the pressure gradient is negligible by construction. Since $L \gg \ell$ we neglect the streamwise Reynolds stress compared to the wall-normal term. On the other hand, since $U_\infty \gg u_\tau$, we can assume the mean flow advection is dominated by the streamwise component $\bar{u}\bar{u}_x$. Finally, the viscous terms are smaller than the advection by a factor on the order of the Reynolds number $Re_L = U_\infty L/\nu \gg 1$. The outer part of the boundary layer is then determined by an inertial balance between streamwise mean flow advection and wall-normal Reynolds stress:

$$(\overline{u'v'})_y = -\bar{u}\bar{u}_x. \quad (10.10)$$

However, this relation cannot describe the near-wall regime, where viscosity is known to be important. In this region we expect the wall-normal derivatives to scale with $(\cdot)_y \sim 1/\eta = u_\tau/\nu$. As a consequence of the no-slip boundary conditions, in this region the free-stream velocity is not an appropriate scale for the streamwise component and we should instead use the friction velocity u_τ , so that

$$\bar{u}\bar{u}_x, \bar{v}\bar{u}_y \sim \frac{u_\tau^2}{L}, \quad \nu \bar{u}_{xx} \sim \left(\frac{\eta}{L}\right) \frac{u_\tau^2}{L}, \quad \nu \bar{u}_{yy} \sim \left(\frac{L}{\eta}\right) \frac{u_\tau^2}{L}, \quad (\overline{u'v'})_y \sim \left(\frac{L}{\eta}\right) \frac{u_\tau^2}{L}, \quad (\overline{u'^2})_x \sim \frac{u_\tau^2}{L}.$$

In this case the wall-normal Reynolds stress is larger than the mean flow advection by a factor of $L/\eta \gg 1$ and must instead be balanced by the viscosity. Therefore, in a thin viscous sublayer near the wall the dominant balance is

$$(\overline{u'v'})_y = \nu \bar{u}_{yy}. \quad (10.11)$$

The overall picture is then that the Reynolds stress must be balanced by mean flow advection in the inertial sublayer and by viscosity in the near-wall region. Outside of the turbulent boundary layer the Reynolds stresses and mean wall-normal velocity are negligible, so small variations, for instance due to incompletely converged statistics, should be described by the balance $\bar{u}\bar{u}_x = -\rho^{-1}\bar{p}_x$. In a true zero pressure gradient flow both of these would be zero in the free stream.

Identified dominant balance Figure 10.1 shows the equation space clusters and associated dominant balance models for the mean fields. As with the cylinder example, some sets of points have significantly reduced variance in certain directions of equation space, a strong signature of the dominant balance phenomenon.

The method identifies regions corresponding to the viscous sublayer (10.11), inertial sublayer (10.10), and slightly perturbed free stream. It also identifies a region near the inlet characterized by a lack of Reynolds stresses, suggesting the mean profile here should be consistent with the laminar solution. The boundaries between balance regimes need not be sharp, however, especially in a transitional flow. In this case a cluster containing all of the active terms in the zero-pressure-gradient flat plate turbulent boundary layer equation is identified between the laminar inflow region and fully developed turbulence downstream.

Equations (10.10) and (10.11) are a starting point for many of the results of boundary layer theory; from these a range of useful laws can be derived, such as the logarithmic mean velocity profile in the inertial sublayer. Although we ultimately hope that data-driven balance identification will open new avenues of analysis, we can also use established results to examine the validity of the proposed method.

For example, the dominant length scale ℓ in the inertial sublayer is expected to depend on the streamwise coordinate x via a power law $\ell \sim x^{4/5}$ (Schlichting, 1955). It is not usually obvious how to extract a specific value of ℓ for which this scaling can be checked. However, as a rough proxy we may consider the wall-normal coordinate at which the dominant balance changes from that of the inertial sublayer to the free-stream. Figure 10.1 shows the growth of the inertial sublayer thickness according to this definition along with a power law fit with exponent 0.81, showing close agreement with the expected value of $4/5$. Although this evidence is somewhat circumstantial, it is at least suggestive that the balance model identification procedure reflects the underlying physics.




Figure 10.5: Self-similarity in the transitional boundary layer. The spatial dominant balance regimes are consistent with known self-similarity in the boundary layer. The laminar balance (Fig. 10.1, purple and green) extends to approximately $x = 200$, a region over which the mean profile approximately matches the Blasius solution (left). Similarly, the turbulent viscous sublayer (Fig. 10.1, red) implies scaling with wall units (right). The mean profile with this scaling collapses until approximately the wall-normal extent of the identified viscous layer, as indicated by dashed lines.

Self-similarity Boundary layers are known to exhibit distinct self-similarity in both the laminar and turbulent regimes (Schlichting, 1955). In the laminar regime, Prandtl's boundary layer equations may be solved by introducing a similarity variable $\eta = y\sqrt{U_\infty/\nu x}$ and a dimensionless streamfunction $f(\eta)$ such that $u(x, y) = U_\infty f'(\eta)$. With this ansatz the momentum equations reduce to the Blasius equation

$$2f''' + f''f = 0 \quad (10.12)$$

with boundary conditions $f(0) = f'(0) = 0$ and $f'(\infty) = 1$. We expect that the mean field of the transitional boundary layer will be approximately given by the Blasius solution in the regions identified with negligible Reynolds stress (green and purple in Fig. 10.1), which extend to approximately $x = 200$.

The left panel of Fig. 10.5 confirms this by comparing the numerically computed Blasius solution to the mean DNS profile. The mean flow closely matches the Blasius profile for the region identified with a laminar dominant balance (solid lines), but significant discrepancies appear in the transitional region (dotted lines).

The scaling of the turbulent region is significantly more complicated, and aspects of it are still a topic of debate (Barrenblatt *et al.*, 1997; Zagarola & Smits, 1998; Morrison *et al.*, 2004; Nickels *et al.*, 2005; Marusic *et al.*, 2010; Smits *et al.*, 2011; Marusic *et al.*, 2013). In the inner layer, the balance between Reynolds stress and viscosity given by (10.11) suggests

the relevant length scale is determined by the friction velocity (Jiménez, 2013). In this case the appropriate scaling is in “wall-units”:

$$y^+ = \frac{yu_\tau}{\nu} \quad u^+ = \frac{u}{u_\tau}. \quad (10.13)$$

The specific variation of u^+ with y^+ is usually thought to transition from a linear dependence in the very near-wall region to a logarithmic “law of the wall” before the scaling changes to that of the outer inertial sublayer. Regardless of the specific functional form, the velocity profile should be self-similar in wall units throughout the viscous-dominated region (red in Fig. 10.1).

The right panel of Fig. 10.5 shows that the mean profile of the turbulent region does indeed collapse in wall units until approximately $y^+ \sim \mathcal{O}(10^2)$, as expected (Jiménez, 2013). The dashed lines indicate the wall-normal extent of the viscous layer; the self-similarity begins to noticeably deteriorate above this point for all streamwise locations.

The identified dominant balance regions are thus consistent with expected self-similarity in the boundary layer, in both laminar and turbulent regimes. This behavior is not built into the algorithm in any way, but the result suggests that the method does indeed identify the correct physical balance in space.

10.4 Discussion

In one guise or another, dominant balance analysis has played a major role in the development of our understanding of many complex systems. In this paper, we have proposed a method of identifying dominant balance regimes in an unsupervised manner directly from data. This approach leverages our understanding of the full physical complexity in the form of governing equations, but by using simple clustering and sparse approximation methods, we avoid any a priori assumptions about balance relations. Nevertheless, the method identifies dominant balance relationships that either recover classical scaling analysis (in the case of the boundary layer and Gulf of Mexico) or confirm arguments based on physical intuition (in the case of nonlinear optics, the Hodgkin-Huxley model, and the combustion analog).

The critical step in this process is the equation space perspective. By considering each term in the governing equation to describe a direction in this space, the dominant balance relations naturally manifest via restriction to subspaces, dramatically reducing variance in directions corresponding to negligible terms. This observation enables the Gaussian mixture model to identify clusters with variance in different directions, and the sparse principal components analysis to extract sparse subspaces by finding directions with significantly nonzero variance. These machine learning tools are applied in a targeted and clearly motivated context, and the equation space perspective necessarily ties the output to underlying physics.

This data-driven approach has the same goal as traditional methods such as scaling

analysis, but introduces several new features. It is a principled, objective approach that does not require the assumption of asymptotic parameter regimes, while providing an estimate of the locally active physical processes throughout domains with arbitrarily complex geometries. The proposed method retains the advantages of the classic approach, but generalizes to a range of disciplines to which traditional analysis cannot readily be applied.

Dominant balance analysis has historically been a critical tool for understanding local physical behavior in complex systems. Non-asymptotic data-driven methods could be used to better understand the behavior of more exotic dynamics such as non-Newtonian turbulence (Samanta *et al.*, 2013) or to study important transitional behavior in cases where the asymptotics are already well known (Hof *et al.*, 2004; Eckhardt *et al.*, 2007; Avila *et al.*, 2011). In the latter case, a clear understanding of the active mechanisms has proven crucial to successful control strategies (Du & Karniadakis, 2000; Hof *et al.*, 2010).

The existence of dominant balance limiting regimes even in complex nonlinear spatiotemporal systems is consistent with the observation that these systems can often be described with sparse representations in function space (Brunton *et al.*, 2016a; Rudy *et al.*, 2017). Building on this insight, we may even be able to identify local dominant balance behavior in spatiotemporal systems without clear governing equations, such as neuroscience, epidemiology, ecology, active fluids, and schooling. For example, the inclusion of spurious terms in the governing equation can be readily detected in the equation space representation (see Supplementary Information); in future work this feature might be leveraged to identify local balance relations in the absence of global conservation equations. This approach thus stands to shed light on more exotic physical processes that have remained elusive to traditional analysis.

However, as with all applications of machine learning and data science methods to physical systems, a critical step in application to any system will be careful validation that the balance identification procedure reproduces the expected results. The dominant balance modeling approach described here is designed to build on, rather than circumvent, physical expertise. The study of dominant balance regimes has been foundational to our understanding of many complex systems; we hope that data-driven methods can integrate with this legacy to enable even wider applicability.

Part VI

OUTLOOK

This thesis has proposed a variety of methods for modeling unsteady fluid flows, including applications to numerical simulations and experimental data. These methods are unified by their objective of characterizing and exploiting the multiscale nature of these flows through low-dimensional modeling and scaling analysis.

After reviewing foundational material in Part II, Part III began from the perspective of standard POD-Galerkin modeling. By examining the structure of these models for advection-dominated flows in Chapter 6, we saw fundamental shortcomings in the use of a linear modal decomposition for low-dimensional modeling. In a dynamical systems perspective this issue can be framed in terms of a latent nonlinear *slow manifold* embedded in the linear reduced-order subspace, leading to manifold modeling methods ranging from analytic power series expansions to deep autoencoders. These strategies accurately approximated the dynamics of shear-driven cavity flow in quasiperiodic and chaotic regimes. Chapter 7 further expanded on the concept of a slow manifold by proposing a multiscale averaging method that treats subscale variables as asymptotically fast relative to the resolved scales. The leading-order correction includes a cubic tensor, extending the Galerkin model to a generalized Stuart-Landau model. When applied to vortex pairing in a spatially developing mixing layer, this approach results in a model of the flow as a set of coupled nonlinear oscillators that preserves complicated phase relationships.

Part IV turned to system identification, proposing the *Langevin regression* extension to the SINDy model discovery framework for stochastic dynamics in Chapter 8. In the present context, this method was motivated by the multiscale view of turbulent flows, where we approximate the broadband turbulence as Gaussian white noise acting on a small number of global integral quantities, or order parameters. Chapter 9 explored an in-depth application of Langevin regression to experimental measurements of symmetry-breaking in a turbulent wake. Along with a mean-flow model inspired by the laminar cylinder wake in Chapter 5, the Langevin regression model reproduced the dominant features of the wake data. This combination of physically-motivated modeling with system identification is highly interpretable, suggesting a new explanation for axisymmetric fluctuations observed in the experimental data but inconsistent with linear and weakly nonlinear stability analyses.

Finally, Part V took a completely different perspective of model reduction in terms of scaling analysis. Rather than constructing a low-dimensional model of the flow with ordinary or stochastic differential equations, this approach aims to directly simplify the governing partial differential equations by identifying local regions where a subset of

the terms are active. Chapter 10 showed how this *dominant balance* analysis could be automated using simple machine learning algorithms and demonstrated the application to both the unsteady cylinder wake and the mean profile of a boundary layer in transition to turbulence. The dominant balance models capture key features of both flows, including nonlinear saturation and self-similarity.

Whether through dynamical systems analysis, stochastic approximations of turbulence, or dominant balance modeling, these approaches leverage the multiscale structure of the flow to derive simplified models. While many of these methods have long been applied analytically in limiting regimes (e.g. near the threshold of instability or for asymptotically large Reynolds number), the work in this thesis has sought to combine these classical approaches with targeted applications of data-driven modeling in order to extend the range of their validity. Throughout, a key focus has been on applying the methods to flows in these intermediate regimes where traditional modeling methods break down, including extensive simulations of laminar flows and experimental observations of high-Reynolds number turbulence.

The final section in each chapter has offered a discussion including potential for future work related to further development of these methods. Some of this work has already continued beyond what is presented in this thesis, including the application of Langevin regression to finite-size scaling in a system of Kuramoto oscillators (Snyder *et al.*, 2021), the development of automated scaling analysis for dimensionally consistent learning (Bakarji *et al.*, 2022), and ongoing work applying both Langevin regression and dominant balance modeling to unsteady shock-boundary layer interactions in wall-bounded hypersonic turbulence. Rather than restate proposed future extensions, these final paragraphs offer an opportunity for reflection on the broader context of this work.

A longstanding promise of reduced-order modeling is the possibility of using the low-dimensional surrogate models for design and control. While this has been accomplished to some extent in controlled academic settings (see Sec. 1.1, there are two critical factors that hinder its development as a robust engineering practice.

First, both design and control optimization depend crucially on the *sensitivity* of the flow to actuation or design changes (Jameson, 1988; Gunzburger, 2003; Noack *et al.*, 2011). This places the *adjoint* linear operators in a privileged role in optimization. The same is true in operator-based modeling methods, including weakly nonlinear analysis (Sipp & Lebedev, 2007) and nonmodal stability theory (Sujith *et al.*, 2016; Barthel *et al.*, 2022). This presents a key challenge for data-driven modeling methods, since it is difficult to empirically approximate adjoint sensitivities without substantial information about the response to perturbations (Juang & Pappa, 1985; Juang *et al.*, 1991; Ma *et al.*, 2011; Herrmann *et al.*, 2021). This is particularly true when the system is strongly nonlinear or non-normal.

Second, even in the case that the system is linear and all direct and adjoint operators are fully accessible, resolving parametric dependence in a model reduction framework is still not straightforward, particularly if there are a large number of parameters (Benner *et al.*, 2015). Dependence on physical parameters such as Reynolds, Mach, or Prandtl

numbers can typically be separated from the operators, as in Sec. 6.6, but the performance of low-order model-based controllers can still be highly sensitive to variable flow conditions (Rowley & Juttijudata, 2005). Moreover, any change to the geometry or control law will generally modify the operators themselves.

As a result of both of these challenges, one of the more exciting recent developments in scientific computing is the development of accessible tools for *differentiable programming*, or adjoint-enabled simulation. Modern differentiable programming tools include general-purpose packages originally designed for machine learning (e.g. PyTorch, Jax, and Flux.jl), but also adjoint-capable open-source multiphysics simulation frameworks such as OpenFOAM, SU2, and FEniCS/Firedrake. While the general mathematical framework for PDE-constrained optimization has been in place for more than 50 years in the context of optimal control theory (Lions, 1971), its application to practical problems has been limited by both computational capabilities and the relative difficulty of applying source-to-source automatic differentiation tools to legacy simulation codes.

As a result, the current landscape suggests that high-level differentiable programming tools in conjunction with advancements in matrix-free Krylov methods for large-scale linear algebra problems (Edwards *et al.*, 1994; Stewart, 2001; Knoll & Keyes, 2004; Antoulas, 2005) stand to deliver the advances in robust simulation-based engineering that have long been promised by reduced-order modeling. This approach combines elements of the computational architecture supporting deep learning with traditional physics and engineering, but does not require any “training” in the standard machine learning sense. For instance, recent work has explored application of differentiable physics models to control (Belbute-Peres *et al.*, 2018; Innes *et al.*, 2019), design optimization (Hoyer *et al.*, 2019), accelerated simulation (Bar-Sinai *et al.*, 2019; Kochkov *et al.*, 2021), and self-assembly (Goodrich *et al.*, 2021)³. When either the numerical code or the underlying physical model is not differentiable, graph-based interaction networks (Battaglia *et al.*, 2018; Thomas *et al.*, 2018; Pfaff *et al.*, 2020) could serve as a flexible, GPU-accelerated, differentiable surrogate models for optimization (Sanchez-Gonzalez *et al.*, 2018; Allen *et al.*, 2022).

This discussion might be summarized concisely with the prediction that even relatively coarse differentiable physics simulations will ultimately be a more robust approach to engineering optimization objectives than low-dimensional modeling, especially as computational and numerical analysis tools continue to advance. On the other hand, the analysis of unsteady fluid flow as a high-dimensional dissipative nonlinear dynamical system has already produced significant contributions to our understanding of basic phenomena such as hydrodynamic stability and the transition to turbulence. A key result of this thesis is that it is possible to capture the dominant features of relatively complex flows with very low-dimensional models; further development of these methods and analysis of the models themselves could yet provide insights into the fundamental physics and nonlinear dynamics of unsteady fluid flow.

³See also <https://su2code.github.io/publications.html> for an extensive reference list of aerodynamics applications.

BIBLIOGRAPHY

- ADRIAN, R. 1975 On the role of conditional averages in turbulence theory. *Proceedings of the 4th Biennial Symposium on Turbulence in Liquids* pp. 323–332.
- ADRIAN, RONALD J. 1979 Conditional eddies in isotropic turbulence. *Physics of Fluids* **22** (11), 2065.
- AHMED, M. A., BAE, H. J., THOMPSON, A. F. & MCKEON, B. J. 2021 Resolvent analysis of stratification effects on wall-bounded shear flows. *Physical Review Fluids* **6** (8), 084804.
- AÏT-SAHALIA, Y., MYKLAND, P. A. & ZHANG, L. 2005 How often to sample a continuous-time process in the presence of market microstructure noise. *Review of Financial Studies* **18**, 351–416.
- ALLEN, K. R., LOPEZ-GUEVARA, T., STACHENFELD, K., SANCHEZ-GONZALEZ, A., BATTAGLIA, P., HAMRICK, J. & PFAFF, T. 2022 Physical design using differentiable learned simulators. *arXiv:2202.00728*.
- ALLEN, L. J. S. 2008 *Mathematical Epidemiology*, vol. 1945, chap. An introduction to stochastic epidemic models. Springer-Verlag Berlin Heidelberg.
- AMSALLEM, D. & FARHAT, C. 2011 An online method for interpolating linear parametric reduced-order models. *SIAM Journal on Scientific Computing* **33** (5), 2169–2198.
- AMSALLEM, D. & FARHAT, C. 2012 Stabilization of projection-based reduced-order models. *International Journal for Numerical Methods in Engineering* **91**, 358–377.
- ANTOULAS, A. C. 2005 *Approximation of Large-Scale Dynamical Systems*. SIAM.
- ARBABI, H. & MEZIĆ, I. 2017 Study of dynamics in post-transient flows using Koopman mode decomposition. *Physical Review Fluids* **2**, 124402.
- ARBABI, H. & SAPSIS, T. 2019 Data-driven modeling of strongly nonlinear chaotic systems with non-Gaussian statistics. *arXiv:1908.08941*.
- AUBRY, N., HOLMES, P., LUMLEY, J. L. & STONE, E. 1988 The dynamics of coherent structures in the wall region of a turbulent boundary layer. *Journal of Fluid Mechanics* **192**, 115.
- AVILA, K., MOXEY, D., DE LOZAR, A., AVILA, M., BARKLEY, D. & HOF, B. 2011 The onset of turbulence in pipe flow. *Science* **333** (6039), 192–196.
- BAE, H. J., LOZANO-DURAN, A. & MCKEON, B. J. 2021 Nonlinear mechanism of the self-sustaining process in the buffer and logarithmic layer of wall-bounded flows. *Journal of Fluid Mechanics* **914**.
- BAGHERI, S., BRANDT, L. & HENNINGSON, D. S. 2009a Input-output analysis, model reduction and control of the flat-plate boundary layer. *Journal of Fluid Mechanics* **620**, 263–298.
- BAGHERI, S., HENNINGSON, D. S., HEPFFNER, J. & SCHMID, P. J. 2009b Input-output analysis and control design applied to a linear model of spatially developing flows. *Applied Mechanics Reviews* **62** (2), 020803.
- BAKARJI, J., CALLAHAM, J. L., BRUNTON, S. L. & KUTZ, J. N. 2022 Dimensionally consistent learning with Buckingham pi. *arXiv:2202.04643*.
- BAKARJI, J. & TARTAKOVSKY, D. 2020 Data-driven discovery of coarse-grained equations. *arXiv:2002.00790*.
- BALAJEWICZ, M., DOWELL, E. & NOACK, B. R. 2013 Low-dimensional modelling of high-Reynolds-number shear flows incorporating constraints from the Navier–Stokes equation. *Journal of Fluid Mechanics* **729**, 285–308.

- BAR-SINAI, Y., HOYER, S., HICKEY, J. & BRENNER, M. P. 2019 Learning data-driven discretizations for partial differential equations. *Proceedings of the National Academy of Sciences* **116** (31), 15344–15349.
- BARBAGALLO, A., SIPP, D. & SCHMID, P. J. 2009 Closed-loop control of an open cavity flow using reduced-order models. *Journal of Fluid Mechanics* **641**, 1.
- BARKLEY, D. 2006 Linear analysis of the cylinder wake mean flow. *Europhysics Letters* **75** (5).
- BARKLEY, D. 2016 Theoretical perspective on the route to turbulence in a pipe. *Journal of Fluid Mechanics* **803**, P1.
- BARKLEY, D. & HENDERSON, R. D. 1996 Three-dimensional Floquet stability analysis of the wake of a circular cylinder. *Journal of Fluid Mechanics* **322**, 215–241.
- BARKLEY, D., SONG, B., MUKUND, V., LEMOULT, G., AVILA, M. & HOF, B. 2015 The rise of fully turbulent flow. *Nature* **526** (7574), 550.
- BARREBLATT, G. I., CHORIN, A. J., HALD, O. H. & PROSTOKISHIN, V. M. 1997 Structure of the zero-pressure-gradient turbulent boundary layer. *Proceedings of the National Academy of Sciences* **94**, 7817–7819.
- BARROS, D., BORÈE, J., NOACK, B. R., SPOHN, A. & RUIZ, T. 2016 Bluff body drag manipulation using pused jets and Coanda effect. *Journal of Fluid Mechanics* **805**, 422–459.
- BARTHEL, B., GOMEZ, S. & MCKEON, B. J. 2022 Variational formulation of resolvent analysis. *Physical Review Fluids* **7**, 013905.
- BARTHEL, B., ZHU, X. & MCKEON, B. J. 2021 Closing the loop: nonlinear Taylor vortex flow through the lens of resolvent analysis. *Journal of Fluid Mechanics* **924**, A9.
- BATTAGLIA, P. W., HAMRICK, J. B., BAPST, V., SANCHEZ-GONZALEZ, A., ZAMBALDI, V., MALINOWSKI, M., TACCHETTI, A., RAPOSO, D., SANTORO, A., FAULKNER, R., GULCEHRE, C., SONG, F., BALLARD, A., GILMER, J., DAHL, G., VASWANI, A., ALLEN, K., NASH, C., LANGSTON, V., DYER, C., HEES, N., WIERSTRA, D., KOHLI, P., BOTVINICK, M., VINYALS, O., LI, Y. & PASCANU, R. 2018 Relational inductive biases, deep learning, and graph networks. *arXiv:1806.01261*.
- BAZILEVS, Y., CALO, V.M., COTTRELL, J.A., HUGHES, T.J.R., REALI, A. & SCOVAZZI, G. 2007 Variational multiscale residual-based turbulence modeling for large eddy simulation of incompressible flows. *Computer Methods in Applied Mechanics and Engineering* **197** (1), 173–201.
- BEAUDOIN, J.-F., CADOT, O., AIDER, J.-L. & WESFREID, J.-E. 2006 Drag reduction of a bluff body using adaptive control methods. *Physics of Fluids* **18** (8), 085107.
- BEETHAM, S. & CAPECELATRO, J. 2020 Formulating turbulence closures using sparse regression with embedded form invariance. *Physical Review Fluids* **5**, 084611.
- BELBUTE-PERES, F. A., SMITH, K., ALLEN, K., TENENBAUM, J. & KOLTER, J. Z. 2018 End-to-end differentiable physics for learning and control. *Advances in Neural Information Processing Systems*.
- BELSON, B. A., TU, J. H. & ROWLEY, C. W. 2014 Algorithm 945: modred - a parallelized model reduction library. *ACM Transactions on Mathematical Software* **40** (4), 1–23.
- BENDER, C. M. & ORSZAG, S. A. 1999 *Advanced Mathematical Methods for Scientists and Engineers I: Asymptotic Methods and Perturbation Theory*. New York: Springer.
- BENEDDINE, S., SIPP, D., ARNAULT, A., DANDOIS, J. & LESSHAFFT, L. 2016 Conditions for validity of mean flow stability analysis. *J. Fluid Mech.* **798**, 485–504.
- BENGANA, Y., LOISEAU, J.-C., ROBINET, J.-C. & TUCKERMAN, L. S. 2019 Bifurcation analysis and frequency prediction in shear-driven cavity flow. *Journal of Fluid Mechanics* **875**, 725–757.
- BENGIO, Y., COURVILLE, A. & VINCENT, P. 2013 Representation Learning: A Review and New Perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**, 1798–1828.

- BENNER, P., GUGERCIN, S. & WILLCOX, K. 2015 A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Review* **57** (4), 483–531.
- BERGER, E., SCHOLZ, D. & SCHUMM, M. 1990 Coherent vortex structures in the wake of a sphere and a circular disk at rest and under forced vibrations. *Journal of Fluids and Structures* **4**, 231–257.
- BERKOOZ, G., HOLMES, P. & LUMLEY, P. L. 1993 The proper orthogonal decomposition in the analysis of turbulent flows. *Annual Review of Fluid Mechanics* **25**, 539–575.
- BILLINGS, S. A. 2013 *Nonlinear system identification: NARMAX methods in the time, frequency, and spatio-temporal domains*. Wiley.
- BISHOP, C. 2006 *Pattern recognition and machine learning*. Springer New York.
- BLOW, K. J. & WOOD, D. 1989 Theoretical description of transient stimulated Raman scattering in optical fibers. *IEEE Journal of Quantum Electronics* **25**.
- BONINSEGNA, L., NÜSKE, F. & CLEMENTI, C. 2018 Sparse learning of stochastic dynamical equations. *Journal of Chemical Physics* **148**.
- BOUJO, E., BOURQUARD, C., XIONG, Y. & NOIRAY, N. 2020 Processing time-series of randomly forced self-oscillators: The example of beer bottle whistling. *Journal of Sound and Vibration* **464**, 114981.
- BOUJO, E. & CADOT, O. 2019 Stochastic modeling of a freely rotating disk facing a uniform flow. *Journal of Fluids and Structures* **86**, 34–43.
- BOUJO, E. & NOIRAY, N. 2017 Robust identification of harmonic oscillator parameters using the adjoint Fokker-Planck equation. *Proceedings of the Royal Society A* **473**.
- BRACKSTON, R. D., DE LA CRUZ, J. M. GARCIA, WYNN, A., RIGAS, G. & MORRISON, J. F. 2016 Stochastic modelling and feedback control of bistability in a turbulent bluff body wake. *Journal of Fluid Mechanics* **802**.
- BRADBURY, JAMES, FROSTIG, ROY, HAWKINS, PETER, JOHNSON, MATTHEW JAMES, LEARY, CHRIS, MACLAURIN, DOUGAL, NECULA, GEORGE, PASZKE, ADAM, VANDERPLAS, JAKE, WANDERMAN-MILNE, SKYE & ZHANG, QIAO 2018 JAX: composable transformations of Python+NumPy programs.
- BRADLEY, E. & KANTZ, H. 2015 Nonlinear time-series analysis revisited. *Chaos* **25**, 097610.
- BRANCHER, P. & CHOMAZ, J.-M. 1997 Absolute and convective instabilities in spatially periodic shear flows. *Physical Review Letters* **78** (4), 658–661.
- BRENNER, M. P., ELDREDGE, J. D. & FREUND, J. B. 2019 Perspective on machine learning for advancing fluid mechanics. *Physical Review Fluids* **4**, 100501.
- BROWN, G. L. & ROSHKO, A. 1974 On density effects and large structures in turbulent mixing layers. *Journal of Fluid Mechanics* **64** (4), 775–816.
- BRÜCKNER, D. B., RONCERAY, P. & BROEDERSZ, C. P. 2020 Inferring the dynamics of underdamped stochastic systems. *Physical Review Letters* **125**, 058103.
- BRUNTON, STEVEN L., BUDIŠIĆ, MARKO, KAISER, EURIKA & KUTZ, J NATHAN 2022 Modern Koopman theory for dynamical systems. *SIAM Review* .
- BRUNTON, S. L., DAWSON, S. & ROWLEY, C. W. 2014 State-space model identification and feedback control of unsteady aerodynamic forces. *Journal of Fluids and Structures* **50**.
- BRUNTON, S. L. & NOACK, B. R. 2015a Closed-loop turbulence control: Progress and challenges. *Applied Mechanics Reviews* **67**, 050801–1–050801–48.
- BRUNTON, S. L. & NOACK, B. R. 2015b Closed-loop turbulence control: Progress and challenges. *Applied Mechanics Review* **67** (5).
- BRUNTON, S. L., NOACK, B. R. & KOUMOUTSAKOS, P. 2020 Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics* **52** (1).

- BRUNTON, S. L., PROCTOR, J. L. & KUTZ, J. N. 2016a Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Nat. Acad. Sci.* **113** (15), 3932–3937.
- BRUNTON, S. L., PROCTOR, J. L. & KUTZ, J. N. 2016b Sparse identification of nonlinear dynamics with control (SINDYc). *IFAC-PapersOnLine* **49** (18), 710–715.
- BURGERS, J. M. 1948 A mathematical model illustrating the theory of turbulence. *Advances in Applied Mechanics* **1**, 171–199.
- CAIAZZO, A., ILIESCU, T., JOHN, V. & SCHYSCHLOWA, S. 2014 A numerical investigation of velocity-pressure reduced order models for incompressible flows. *Journal of Computational Physics* **259**, 598–616.
- CALLAHAM, J. L., BRUNTON, S. L. & LOISEAU, J.-C. 2022 On the role of nonlinear correlations in reduced-order modeling. *Journal of Fluid Mechanics* **938**, A1.
- CALLAHAM, J. L., KUTZ, J. N., BRUNTON, B. W. & BRUNTON, S. L. 2021a Learning dominant physical processes with data-driven balance models. *Nature Communications* **12** (1), 1–10.
- CALLAHAM, J. L., LOISEAU, J.-C., RIGAS, G. & BRUNTON, S. L. 2021b Nonlinear stochastic modeling with Langevin regression. *Proceedings of the Royal Society A* **477** (2250), 20210092.
- CALLAHAM, J. L., RIGAS, G., LOISEAU, J.-C. & BRUNTON, S. L. 2021c An empirical mean-field model of symmetry-breaking in a turbulent wake. *arXiv:2105.13990*.
- CARINI, M., AUTERI, F. & GIANNETTI, F. 2015 Centre-manifold reduction of bifurcating flows. *J. Fluid Mech.* **767**, 109–145.
- CARLBERG, KEVIN 2015 Adaptive h-refinement for reduced-order models. *International Journal for Numerical Methods in Engineering* **102** (5).
- CARLBERG, K., BARONE, M. & ANTIL, H. 2017 Galerkin v. least-squares Petrov–Galerkin projection in nonlinear model reduction. *J. Comput. Phys.* **330**, 693–734.
- CARLBERG, K., BOU-MOSLEH, C. & FARHAT, C. 2011 Efficient non-linear model reduction via a least-squares Petrov–Galerkin projection and compressive tensor approximations. *International Journal for Numerical Methods in Engineering* **86** (2).
- CARLBERG, KEVIN, FARHAT, CHARBEL, CORTIAL, JULIEN & AMSALLEM, DAVID 2013 The GNAT method for nonlinear model reduction: effective implementation and application to computational fluid dynamics and turbulent flows. *Journal of Computational Physics* **242**, 623–647.
- CAVALIERI, A. V. G. 2021 Structure interactions in a reduced-order model for wall-bounded turbulence. *Physical Review Fluids* **6**, 034610.
- CERDA, E. & MAHADEVAN, L. 2003 Geometry and physics of wrinkling. *Phys. Rev. Letters* **90** (7), 074302.
- CHAMPION, KATHLEEN, LUSCH, BETHANY, KUTZ, J. NATHAN & BRUNTON, STEVEN L. 2019a Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences* **116** (45), 22445–22451.
- CHAMPION, K., ZHENG, P., ARAVKIN, A., BRUNTON, S. L. & KUTZ, J. N. 2019b A unified sparse optimization framework to learn parsimonious physics-informed models from data. *IEEE Access* **8**, 169259–169271.
- CHARRU, F. 2011 *Hydrodynamic instabilities*. Cambridge University Press.
- CHIEN, W.-L., RISING, H. & OTTINO, J. M. 1986 Laminar mixing and chaotic mixing in several cavity flows. *J. Fluid Mech.* **170** (-1), 355.
- CHOMAZ, J.-M. 2005 Global instabilities in spatially developing flows: non-normality and non-linearity. *Annual Review of Fluid Mechanics* **37**, 357–392.
- COLONIUS, T., LELE, S. K. & MOIN, P. 1997 Sound generation in a mixing layer. *Journal of Fluid Mechanics* **330**, 375–409.

- CONSTANTINE, P. G., DOW, E. & WANG, Q. 2014 Active subspace methods in theory and practice: applications to Kriging surfaces. *SIAM Journal on Scientific Computing* **36** (4), A1500–1524.
- CORDIER, L., NOACK, B. R., TISSOT, G., LEHNASCH, G., DELVILLE, J., BALAJEWICZ, M., DAVILLER, G. & NIVEN, R. K. 2013 Identification strategies for model-based control. *Experiments in Fluids* **54**, 1580.
- COULLET, P. H. & SPIEGEL, E. A. 1983 Amplitude equations for systems with competing instabilities. *SIAM Journal of Applied Mathematics* **43** (4), 776–820.
- CRANMER, M., GREYDANUS, S., HOYER, S., BATTAGLIA, P., SPERGEL, D. & HO, S. 2020a Lagrangian neural networks. *arXiv:2003.04630*.
- CRANMER, M., SANCHEZ-GONZALEZ, A., BATTAGLIA, P., XU, R., CRANMER, K., SPERGEL, D. & HO, S. 2020b Discovering symbolic models from deep learning with inductive biases. *Advances in Neural Information Processing Systems* **33**, 17429–17442.
- CROSS, M. C. & HOHENBERG, P. C. 1993 Pattern formation outside of equilibrium. *Reviews of modern physics* **65** (3), 851.
- CRUTCHFIELD, J. P. & MCNAMARA, B. S. 1987 Equations of motion from a data series. *Complex Systems* **1**, 417–452.
- DEANE, A. E., KEVREKEDIS, I. G., KARNIADAKIS, G. E. & ORSZAG, S. A. 1991 Low-dimensional models for complex-geometry flows - application to grooved channels and circular-cylinders. *Physics of Fluids A* **3**.
- DELAHUNT, C. B. & KUTZ, J. N. 2022 A toolkit for data-driven discovery of governing equations in high-noise regimes. *IEEE Access* **10**, 31210–31234.
- DENG, N., NOACK, B. R., MORZYNSKI, M. & PASTUR, L. R. 2020 Low-order model for successive bifurcations of the fluidic pinball. *Journal of Fluid Mechanics* **884**, A37–1.
- DRAZIN, P. G. & REID, W. H. 2004 *Hydrodynamic stability*. Cambridge University Press.
- DU, Y. & KARNIADAKIS, G. E. 2000 Suppressing wall turbulence by means of a transverse traveling wave. *Science* **288** (5469), 1230–1234.
- DUCIMETIÈRE, Y.-M., BOUJO, E. & GALLAIRE, F. 2022 Weak nonlinearity for strong nonnormality. *arXiv:2110.08064*.
- DURASAMY, K., IACCARINO, G. & XIAO, H. 2019 Turbulence modeling in the age of data. *Annual Reviews of Fluid Mechanics* **51**, 357–377.
- E, W. 2011 *Principles of multiscale modeling*. Cambridge University Press.
- E, W. & ENGQUIST, B. 2003 The heterogeneous multiscale methods. *Communications in Mathematical Sciences* **1**, 87–132.
- E, W., ENGQUIST, B., LI, X., REN, W. & VANDEN-EIJNDEN, E. 2007 The heterogeneous multiscale method: a review. *Communications in Computational Physics* p. 2007.
- ECKHARDT, B., SCHNEIDER, T. M., HOF, B. & WESTERWEEL, J. 2007 Turbulence transition in pipe flow. *Annual Review of Fluid Mechanics* **39**, 447–468.
- EDWARDS, W. S., TUCKERMAN, L. S., FRIESNER, R. A. & SORENSEN, D. C. 1994 Krylov methods for the incompressible navier-stokes equations. *J. Comput. Phys.* **110** (1), 82–102.
- EINSTEIN, A. 1905 Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen. *Ann. Phys.* **322** (8), 549–560.
- EVERSON, R. & SIROVICH, L. 1995 Kauhunen-Loève procedure for gappy data. *Journal of the Optical Society of America* **12** (8), 1657–1664.
- FASEL, U., KAISER, E., KUTZ, J. N., BRUNTON, B. W. & BRUNTON, S. L. 2022 Ensemble-SINDy: Robust sparse model discovery in the low-data, high-noise limit, with active learning and control. *Proceedings of the Royal Society A* **478** (2260), 20210904.

- FISCHER, P. F., LOTTES, J. W. & KERKEMEIR, S. G. 2008 Nek5000 web pages. <http://nek5000.mcs.anl.gov>.
- FRIEDRICH, R. & PEINKE, J. 1997 Description of a turbulent cascade by a Fokker-Planck equation. *Physical Review Letters* **78** (5), 863–866.
- FRIEDRICH, R., PEINKE, J., SAHIMI, M. & TABAR, M. R. R. 2011 Approaching complexity by stochastic methods: from biological systems to turbulence. *Physics Reports* **506**, 87–162.
- FRIEDRICH, R., ZELLER, J. & PEINKE, J. 1998 A note on three-point statistics of velocity increments in turbulence. *Europhysics Letters* **41** (2), 153–158.
- FRISHMAN, A. & RONCERAY, P. 2020 Learning force fields from stochastic trajectories. *Physical Review X* **10**, 021009.
- GALLAIRE, F., BOUJO, E., MANTIČ-LUGO, V., ARRATIA, C., THIRA, B. & MELIGA, P. 2016 Pushing amplitude equations far from threshold: application to the supercritical Hopf bifurcation in the cylinder wake. *Fluid Dynamics Research* **48**, 061401.
- GAUTIER, N., AIDER, J. L., DURIEZ, T., NOACK, B. R., SEGOND, M. & ABEL, M. 2015 Closed-loop separation control using machine learning. *Journal of Fluid Mechanics* **770**, 442–457.
- GHASEMI, F., BAHRAMINASAB, A., MOVAHED, M. S., RAHVAR, S., SREENIVASAN, K. R. & TABAR, M. R. R. 2006 Characteristic angular scales in cosmic microwave background radiation. *Journal of Statistical Mechanics* **11**, P11008.
- GHIRARDO, G., JUNIPER, M. P. & MOECK, J. P. 2016 Weakly nonlinear analysis of thermoacoustic instabilities in annular combustors. *Journal of Fluid Mechanics* **805**, 52–87.
- GILL, A. 1982 *Atmosphere-Ocean Dynamics*. Academic Press.
- GIVON, D., KUPFERMAN, R. & STUART, A. M. 2004 Extracting macroscopic dynamics: model problems and algorithms. *Nonlinearity* **17** (6), R55–R127.
- GLAZ, B., MEZIĆ, I., FONOBEROVA, M. & LOIRE, S. 2017 Quasi-periodic intermittency in oscillating cylinder flow. *Journal of Fluid Mechanics* **828**, 680–707.
- GOLLUB, J. P. & SWINNEY, H. L. 1975 Onset of turbulence in a rotating fluid. *Physical Review Letters* **35** (14), 927–930.
- GOLUBITSKY, M. & LANGFORD, W. 1988 Pattern formation and bistability in flow between counterrotating cylinders. *Physica D* **32**, 362–392.
- GOLUBITSKY, M. & STEWART, I. 1985 Hopf bifurcation in the presence of symmetry. *Archive for Rational Mechanics and Analysis* **87**, 107–165.
- GOODFELLOW, I., BENGIO, Y. & COURVILLE, A. 2016 *Deep learning*. MIT Press.
- GOODRICH, C. P., KING, E. M., SCHOENHOLZ, S. S., CUBUK, E. D. & BRENNER, M. P. 2021 Designing self-assembling kinetics with differentiable statistical physics models. *Proceedings of the National Academy of Sciences* **118** (10).
- GRANDEMANGE, M., GOHLKE, M. & CADOT, O. 2012 Reflectional symmetry breaking of the separated flow over three-dimensional bluff bodies. *Physical Review E* **86** (3), 035302.
- GRANDEMANGE, M., GOHLKE, M. & CADOT, O. 2013 Turbulent wake past a three-dimensional blunt body. Part 1. Global modes and bi-stability. *Journal of Fluid Mechanics* **722**, 51–84.
- GRIMBERG, S., FARHAT, C. & YOUKILIS, N. 2020 On the stability of projection-based model order reduction for convection-dominated laminar and turbulent flows. *Journal of Computational Physics* **419**, 109681.
- GRZYBOWSKI, B., STONE, H. A. & WHITESIDES, G. M. 2000 Dynamic self-assembly of magnetized, millimetre-sized objects rotating at a liquid-air interface. *Nature* **405**, 1033–1036.
- GUAN, Y., BRUNTON, S. L. & NOVOSSELOV, I. 2021 Sparse nonlinear models of chaotic electroconvection. *Royal Society Open Science* **8** (8), 202367.

- GUCKENHEIMER, J. 1986 Strange attractors in fluids: another view. *Annual Review of Fluid Mechanics* **18**, 15–31.
- GUCKENHEIMER, J. & HOLMES, P. 1983 *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*. Springer.
- GUNZBURGER, M. D. 2003 *Perspectives in flow control and optimization*. Society for Industrial & Applied Mathematics (SIAM).
- HAKEN, H. 1983 *Synergetics: an Introduction: Nonequilibrium Phase Transitions and Self-Organization in Physics, Chemistry and Biology*. Springer New York.
- HARRISON, L. M., DAVID, O. & FRISTON, K. J. 2005 Stochastic models of neuronal dynamics. *Phil. Trans. R. Soc. B* **360**, 1075–1091.
- HEMATI, M. S., DAWSON, S. T. M. & ROWLEY, C. W. 2017 Parameter-varying aerodynamics models for aggressive pitching-response prediction. *AIAA Journal* **55** (3).
- HENNIGAN, TOM, CAI, TREVOR, NORMAN, TAMARA & BABUSCHKIN, IGOR 2020 Haiku: Sonnet for JAX.
- HERRMANN, B., BADDOO, P. J., SEMAAN, R., BRUNTON, S. L. & MCKEON, B. J. 2021 Data-driven resolvent analysis. *Journal of Fluid Mechanics* **918**, A10.
- HERRMANN, B., OSWALD, P., SEMAAN, R. & BRUNTON, S. L. 2020 Modeling synchronization in forced turbulent oscillator flows. *arxiv:2001.12000* .
- HESSSEL, MATTEO, BUDDEN, DAVID, VIOLA, FABIO, ROSCA, MIHAELA, SEZENER, EREN & HENNIGAN, TOM 2020 Optax: composable gradient transformation and optimisation, in jax!
- HESTHAVEN, J. S. & UBBIALI, S. 2018 Non-intrusive reduced order modeling of nonlinear problems using neural networks. *Journal of Computational Physics* **363**, 55–78.
- HOF, B., VAN DOORNE, C. W. H., WESTERWEEL, J., NIEUWSTADT, F. T. M., FAISST, H., ECKHARDT, B., WEDIN, H., KERSWELL, R. R. & WALEFFE, F. 2004 Experimental observation of nonlinear traveling waves in turbulent pipe flow. *Science* **305** (5960), 1594–1598.
- HOF, B., JUEL, A. & MULLIN, T. 2003 Scaling of the turbulence transition threshold in a pipe. *Physical Review Letters* **91** (24), 244502.
- HOF, B., DE LOZAR, A., AVILA, M., TU, X. & SCHNEIDER, T. M. 2010 Eliminating turbulence in spatially intermittent flows. *Science* **327** (5972), 1491–1494.
- HOLMES, P., LUMLEY, J. L. & BERKOOZ, G. 1996 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge Monographs on Mechanics.
- HONISCH, C. & FRIEDRICH, R. 2011 Estimation of Kramers-Moyal coefficients at low sampling rates. *Physical Review E* **83**.
- HOPF, E. 1948 A mathematical example displaying features of turbulence. *Communications on Pure and Applied Mathematics* **1** (4), 302–322.
- HOYER, S., SOHL-DICKSTEIN, J. & GREYDANUS, S. 2019 Neural reparameterization improves structural optimization. *arXiv:1909.04240* .
- HUERRE, P. & MONKEWITZ, P. A. 1985 Absolute and convective instabilities in free shear layers. *Journal of Fluid Mechanics* **159**, 151–168.
- HUERRE, P. & MONKEWITZ, P. A. 1990 Local and global instabilities in spatially developing shear flows. *Annual Review of Fluid Mechanics* **22**, 473–537.
- HUSSAIN, A. K. M. F. 1981 Role of coherent structures in turbulent shear flows. *Proceedings of the Indian Academy of Sciences Section C: Engineering Sciences* **4** (129).
- INNES, M., EDELMAN, A., FISCHER, K., RACKAUCKAS, C., SABA, E., SHAH, V. B. & TEBBUTT, W. 2019 ∂P a differentiable programming system to bridge machine learning and scientific computing. *arXiv:1907.07587* .

- JACKSON, C. P. 1987 A finite-element study of the onset of vortex shedding in flow past variously shaped bodies. *Journal of Fluid Mechanics* **182**, 23–45.
- JACOBI, I., CHUNG, D., DUVVURI, S. & MCKEON, B. J. 2021 Interactions between scales in wall turbulence: phase relationships, amplitude modulation and the importance of critical layers. *Journal of Fluid Mechanics* **914**.
- JAMESON, A. 1988 Aerodynamic design via control theory. *Journal of Scientific Computing* **3** (3), 233–260.
- JIMÉNEZ, J. 2013 Near-wall turbulence. *Physics of Fluids* **25**, 101302.
- JIMÉNEZ, J. 2018 Coherent structures in wall-bounded turbulence. *Journal of Fluid Mechanics* **842**, P1.
- JUANG, J. N. & PAPPAS, R. S. 1985 An eigensystem realization algorithm for modal parameter identification and model reduction. *Journal of Guidance, Control, and Dynamics* **8** (5), 620–627.
- JUANG, J. N., PHAN, M., HORTA, L. G. & LONGMAN, R. W. 1991 Identification of observer/Kalman filter Markov parameters: theory and experiments. Technical Memorandum 104069. NASA.
- KAHEMAN, K., KUTZ, J. N. & BRUNTON, S. L. 2020 SINDy-PI: a robust algorithm for parallel implicit sparse identification of nonlinear dynamics. *Proceedings of the Royal Society A* **476** (2242), 20200279.
- KAISER, EURIKA, KUTZ, J NATHAN & BRUNTON, STEVEN L 2018 Sparse identification of nonlinear dynamics for model predictive control in the low-data limit. *Proc. Roy. Soc. London A* **474** (2219).
- KANTZ, H. & SCHREIBER, T. 2003 *Nonlinear time series analysis*, 2nd edn. Cambridge University Press.
- KAPTANOGLU, A. A., CALLAHAM, J. L., HANSEN, C. J., ARAVKIN, A. & BRUNTON, S. L. 2021a Promoting global stability in data-driven models of quadratic nonlinear dynamics. *Physical Review Fluids* **6** (9), 094401.
- 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. *Physical Review E* **104** (1), 015206.
- KAPTANOGLU, A. A., DE SILVA, B. M., FASEL, U., KAHEMAN, K., GOLDSCHMIDT, A. J., CALLAHAM, J., DELAHUNT, C. B., NICOLAU, Z. G., CHAMPION, K., LOISEAU, J.-C., KUTZ, J. N. & BRUNTON, S. L. 2022 PySINDy: A comprehensive Python package for robust sparse system identification. *Journal of Open Source Software* **7** (69), 3994.
- KINGMA, D. P. & WELLING, M. 2013 Auto-encoding variational Bayes. *arXiv:1312.6114* .
- KLUS, S., NÜSKE, F., PEITZ, S., NIEMANN, J.-H., CLEMENTI, C. & SCHÜTTE, C. 2020 Data-driven approximation of the Koopman generator: Model reduction, system identification, and control. *Physica D: Nonlinear Phenomena* **406**, 132416.
- KNOLL, D. A. & KEYES, D. E. 2004 Jacobian-free Newton-Krylov methods: a survey of approaches and applications. *Journal of Computational Physics* **193** (2), 357–397.
- KOCHKOV, D., SMITH, J. A., ALIEVA, A., WANG, Q., BRENNER, M. P. & HOYER, S. 2021 Machine learning-accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences* **118** (21), e2101784118.
- KOLMOGOROV, A. N. 1941 The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. *Doklady Akademiia Nauk SSSR* **30**, 301–305.
- KRAICHNAN, R. H. & CHEN, S. 1989 Is there a statistical mechanics of turbulence? *Physica D: Nonlinear Phenomena* **37**, 160–172.
- KURTZ, T. G. 1973 A limit theorem for perturbed operator semigroups with applications to random evolutions. *Journal of Functional Analysis* **12**, 55–67.

- KUTZ, J. N., BRUNTON, S. L., BRUNTON, B. W. & PROCTOR, J. L. 2016 *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM.
- LADE, S. J. 2009 Finite sampling interval effects in Kramers-Moyal analysis. *Physics Letters A* **373** (41).
- LAING, C. R., FREWEN, T. & KEVREKEDIS, I. G. 2010 Reduced models for binocular rivalry. *Journal of Computational Neuroscience* **28**, 459–476.
- LANDAU, L. D. 1944 On the problem of turbulence. *Doklady Akademiia Nauk SSSR* **44**, 339–342.
- LANDAU, L. D. & LIFSHITZ, E. M. 1959 *Fluid Mechanics*. Elsevier.
- LANFORD, O. E. III 1982 The strange attractor theory of turbulence. *Annual Review of Fluid Mechanics* **14**, 347–364.
- LANGEVIN, P. 1908 Sur la théorie du mouvement brownien. *C. R. Acad. Sci. Paris* **146**, 530–533.
- LECLERCQ, C., DEMOURANT, F., POUSSOT-VASSAL, C. & SIPP, D. 2019 Linear iterative method for closed-loop control of quasiperiodic flows. *J. Fluid Mech.* **868**, 26–65.
- LEE, J. & ZAKI, T. A. 2018 Detection algorithm for turbulent interfaces and large-scale structures in intermittent flows. *Computers & Fluids* **175**, 142–158.
- LEE, K. & CARLBERG, K. T. 2020 Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *Journal of Computational Physics* **404**, 108973.
- LEGOLL, F. & LELIÈVRE, T. 2010 Effective dynamics using conditional expectations. *Nonlinearity* **23** (9).
- LEVIN, S. A. 1992 The problem of pattern and scale in ecology. *Ecology* **73** (6), 1943–1967.
- LGUENSAT, R., TANDEO, P., AILLIOT, P., PULIDO, M. & FABLET, R. 2017 The analog data assimilation. *Monthly Weather Review* **145** (10), 4093–4107.
- LI, Y., PERLMAN, E., WAN, M., YANG, Y., BURNS, R., MENEVEAU, C., BURNS, R., CHEN, S., SZALAY, A. & EYINK, G. 2008 A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence. *Journal of Turbulence* **9** (31).
- LIGHTHILL, M. J. 1952 On sound generated aerodynamically. Part I. *Proc. R. Soc. A* **211**, 564–587.
- LIGHTHILL, M. J. 1966 Dynamics of rotating fluids: a survey. *Journal of Fluid Mechanics* **26**, 411–431.
- LING, J., KURZAWSKI, A. & TEMPLETON, J. 2016 Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. *Journal of Fluid Mechanics* **807**, 155–166.
- LIONS, J.-L. 1971 *Optimal control of systems governed by partial differential equations*. Springer Berlin Heidelberg.
- LIU, Q., GÓMEZ, F. & THEOFILIS, V. 2016 Linear instability analysis of low-re incompressible flow over a long rectangular finite-span open cavity. *J. Fluid Mech.* **799**, R2 (16 pages).
- LIU, Q., SUN, Y., YEH, C.-A., UKEILEY, L. S., CATTAFESTA, L. N. & TAIRA, K. 2021 Unsteady control of supersonic turbulent cavity flow based on resolvent analysis. *Journal of Fluid Mechanics* **925**, A5.
- LOISEAU, J.-C. 2020 Data-driven modeling of the chaotic thermal convection in an annular thermosyphon. *Theoretical and Computational Fluid Dynamics* **34** (4), 339–365.
- LOISEAU, J.-C. & BRUNTON, S. L. 2018 Constrained sparse Galerkin regression. *Journal of Fluid Mechanics* **838**, 42–67.
- LOISEAU, J.-C., BRUNTON, S. L. & NOACK, B. R. 2018a *Handbook on Model Order Reduction*, chap. From the POD-Galerkin method to sparse manifold models. De Gruyter GmbH.
- LOISEAU, J.-C., BUCCI, M. A., CHERUBINI, S. & ROBINET, J.-C. 2019 *Computational Modelling of Bifurcations and Instabilities in Fluid Dynamics*, chap. Time-stepping and Krylov methods for large-scale instability problems, pp. 33–73. Springer.

- LOISEAU, J.-C., NOACK, B. R. & BRUNTON, S. L. 2018b Sparse reduced-order modelling: sensor-based dynamics to full-state estimation. *Journal of Fluid Mechanics* **844**, 454–490.
- LOPEZ-PAZ, D., HENNIG, P. & SCHÖLKOPF, B. 2013 The randomized dependence coefficient. In *Advances in Neural Information Processing Systems 26* (ed. C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani & K. Q. Weinberger), pp. 1–9. Curran Associates, Inc.
- LORENZ, E. N. 1963 Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences* **20** (2), 130–141.
- LUCHTENBURG, D. M., GÜNTHER, B., NOACK, B. R., KING, R. & TADMOR, G. 2009 A generalized mean-field model of the natural and high-frequency actuated flow around a high-lift configuration. *Journal of Fluid Mechanics* **623**, 283.
- LUHAR, M., SHARMA, A. S. & MCKEON, B. J. 2014 Opposition control within the resolvent analysis framework. *Journal of Fluid Mechanics* **749**, 597–626.
- LUMLEY, JOHN L. 1967 *Atmospheric Turbulence and Wave Propagation*, chap. The structure of inhomogeneous turbulence.
- MA, XIA & KARNIADAKIS, GEORGE EM 2002 A low-dimensional model for simulating three-dimensional cylinder flow. *Journal of Fluid Mechanics* **458**, 181–190.
- MA, Z., AHUJA, S. & ROWLEY, C. W. 2011 Reduced order models for control of fluids using the eigensystem realization algorithm. *Theoretical and Computational Fluid Dynamics* **25**, 233–247.
- MAJDA, A. & WANG, X. 2006 *Nonlinear dynamics and statistical theories for basic geophysical flows*. Cambridge University Press.
- MAJDA, ANDREW J & TIMOFEYEV, ILYA 2000 Remarkable statistical behavior for truncated Burgers–Hopf dynamics. *Proceedings of the National Academy of Sciences* **97** (23), 12413–12417.
- MAJDA, A. J., TIMOFEYEV, I. & VANDEN-EIJNDEN, E. 2001 A mathematical framework for stochastic climate models. *Communications on Pure and Applied Mathematics* **54** (8), 891–974.
- MAJDA, A. J., TIMOFEYEV, I. & VANDEN-EIJNDEN, E. 2006 Stochastic models for selected slow variables in large deterministic systems. *Nonlinearity* **19**, 769–794.
- MALKUS, W. V. R. 1956 Outline of a theory of turbulent shear flow. *Journal of Fluid Mechanics* **1** (5), 521–539.
- MANGAN, N. M., BRUNTON, S. L., PROCTOR, J. L. & KUTZ, J. N. 2016 Inferring biological networks by sparse identification of nonlinear dynamics. *IEEE Transactions on Molecular, Biological, and Multi-Scale Communications* **2** (1), 52–63.
- MANGAN, N. M., KUTZ, J. N., BRUNTON, S. L. & PROCTOR, J. L. 2017 Model selection for dynamical systems via sparse regression and information criteria. *Proceedings of the Royal Society A* **473** (2204).
- MANTIČ-LUGO, V., ARRATIA, C. & GALLAIRE, F. 2014 Self-consistent mean flow description of the nonlinear saturation of the vortex shedding in the cylinder wake. *Physical Review Letters* **113**, 084501.
- MANTIČ-LUGO, V. & GALLAIRE, F. 2016 Self-consistent model for the saturation mechanism of the response to harmonic forcing in the backward-facing step flow. *Journal of Fluid Mechanics* **793**, 777–797.
- MARUSIC, I., MATHIS, R. & HUTCHINS, N. 2010 Predictive model for wall-bounded turbulent flow. *Science* **329** (5988), 193–196.
- MARUSIC, I., MONTY, J. P., HULTMARK, M. & SMITS, A. J. 2013 On the logarithmic region in wall turbulence. *Journal of Fluid Mechanics* **716**, R3.
- MAULIK, R., SAN, O., RASHEED, A. & VEDULA, P. 2019 Subgrid modelling for two-dimensional turbulence using neural networks. *Journal of Fluid Mechanics* **858**, 122–144.

- MCKEON, B. J. 2017 The engine behind (wall) turbulence: perspectives on scale interactions. *Journal of Fluid Mechanics* **817**.
- MCKEON, B. J. & SHARMA, A. S. 2010 A critical-layer framework for turbulent pipe flow. *J. Fluid Mech.* **658**, 336–382.
- MCKEON, BEVERLY J., SHARMA, ATI S. & JACOBI, IAN 2013 Experimental manipulation of wall turbulence: A systems approach. *Physics of Fluids* **25**.
- MELIGA, P. 2017 Harmonics generation and the mechanics of saturation in flow over an open cavity: a second-order self-consistent description. *Journal of Fluid Mechanics* **826**, 503–521.
- MELIGA, P. & CHOMAZ, J.-M. 2011 An asymptotic expansion for the vortex-induced vibrations of a circular cylinder. *Journal of Fluid Mechanics* **671**, 137–167.
- MELIGA, P., CHOMAZ, J.-M. & SIPP, D. 2009 Global mode interaction and pattern selection in the wake of a disk: a weakly nonlinear expansion. *Journal of Fluid Mechanics* **633**, 159–189.
- MENDIBLE, A., BRUNTON, S. L., ARAVKIN, A. Y., LOWRIE, W. & KUTZ, J. N. 2020 Dimensionality reduction and reduced-order modeling for traveling wave physics. *Theoretical and Computational Fluid Dynamics* **34** (4), 385–400.
- MENIER, E., BUCCI, M. A., YAGOUBI, M., MATHELIN, L. & SCHCOENAUER, M. 2022 CD-DOM: Complementary deep - reduced order model. *arXiv:2202.10746* .
- MESENTER, D. A. & BORTZ, D. M. 2021 Weak SINDy for partial differential equations. *Journal of Computational Physics* p. 110525.
- MEZIĆ, I. 2005 Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics* **41** (1-3), 309–325.
- MEZIĆ, I. 2013 Analysis of fluid flows via spectral properties of the Koopman operator. *Ann. Rev. Fluid Mech.* **45** (1), 357–378.
- MICHALKE, A. 1964 On the inviscid instability of the hyperbolic-tangent velocity profile. *Journal of Fluid Mechanics* **19** (4), 543–556.
- MOEHLIS, J., FAISST, H. & ECKHARDT, B. 2004 A low-dimensional model for turbulent shear flows. *New Journal of Physics* **6**, 56.
- MOHEBUJJAMAN, M., REBHOLZ, L. G. & ILIESCU, T. 2018 Physically constrained data-driven correction for reduced-order modeling of fluid flows. *International Journal for Numerical Methods in Fluids* **89** (3), 103–122.
- MOHEBUJJAMAN, M., REBHOLZ, L. G., XIE, X. & ILIESCU, T. 2017 Energy balance and mass conservation in reduced order models of fluid flows. *Journal of Computational Physics* **346** (1), 262–277.
- MONIN, A. S. & YAGLOM, A. M. 1971 *Statistical fluid mechanics, volume I: mechanics of turbulence*. MIT Press.
- MONKEWITZ, P. A & HUERRE, P. 1982 Influence of the velocity ratio on the spatial instability of mixing layers. *Physics of Fluids* **25** (7), 1137.
- MOORE, B. C. 1981 Principle component analysis in linear systems: controllability, observability, and model reduction. *IEEE Transactions on Automatic Control* **26**, 17–32.
- MORRIS, S. W., BODENSCHATZ, E., CANNELL, D. S. & AHLERS, G. 1993 Spiral defect chaos in large aspect ratio Rayleigh-Bénard convection. *Physical Review Letters* **71** (13), 2026.
- MORRISON, J. F., MCKEON, B. J., JIANG, W. & SMITS, A. J. 2004 Scaling of the streamwise velocity component in turbulent pipe flow. *Journal of Fluid Mechanics* **508**, 99–131.
- NICKELS, T. B., MARUSIC, I., HAFEZ, S. & CHONG, M. S. 2005 Evidence of the k_1^{-1} law in a high-Reynolds-number turbulent boundary layer. *Physical Review Letters* **95**, 074501.

- NOACK, B. R., AFANASIEV, K., MORZYNSKI, M., TADMOR, G. & THIELE, F. 2003 A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.* **497**, 335–363.
- NOACK, B. R., MORZYNSKI, M. & TADMOR, G. 2011 *Reduced-order modelling for flow control*, , vol. 528. Springer Science & Business Media.
- NOACK, B. R., PAPAS, P. & MONKEWITZ, P. A. 2005 The need for a pressure-term representation in empirical Galerkin models of incompressible shear flows. *Journal of Fluid Mechanics* **523**, 339–365.
- NOACK, B. R., SCHLEGEL, M., AHLBORN, B., MUTSCHKE, G., MORZYNSKI, M., COMTE, P. & TADMOR, G. 2008 A finite-time thermodynamics of unsteady fluid flows. *Journal of Non-Equilibrium Thermodynamics* **33**, 103–148.
- NOÉ, F., OLSSON, S., KÖHLER, J. & WU, H. 2019 Boltzmann generators: sampling equilibrium states of many-body systems with deep learning. *Science* **365** (6457), eaaw1147.
- NOÉ, F., TKATCHENKO, A., MÜLLER, K. R. & CLEMENTI, C. 2020 Machine learning for molecular simulation. *Annual review of physical chemistry* **71**, 361–390.
- NOVATI, G., DE LAROUSILHE, H. L. & KOUMOUTSAKOS, P. 2021 Automating turbulence modeling by multi-agent reinforcement learning. *Nature Machine Intelligence* **3** (1), 87–96.
- ORCHINI, ALESSANDRO, RIGAS, GEORGIOS & JUNIPER, MATTHEW P 2016 Weakly nonlinear analysis of thermoacoustic bifurcations in the rijke tube. *Journal of Fluid Mechanics* **805**, 523–550.
- ÖSTH, J., NOACK, B. R., KRAJNOVIĆ, S., BARROS, D. & BORÈE, J. 2014 On the need for a nonlinear subscale turbulence term in POD models as exemplified for a high-Reynolds-number flow over an Ahmed body. *Journal of Fluid Mechanics* **747**, 518–544.
- PADOVAN, A., OTTO, S. E. & ROWLEY, C. W. 2020 Analysis of amplification mechanisms and cross-frequency interactions in nonlinear flows via the harmonic resolvent. *Journal of Fluid Mechanics* **900**, A14.
- PAN, S. & DURAISAMY, K. 2018 Data-driven discovery of closure models. *SIAM Journal of Applied Dynamical Systems* **17** (4), 2381–2413.
- PAPANICOLAOU, G. 1976 Some probabilistic problems and methods in singular perturbations. *Rocky Mountain Journal of Mathematics* **6** (4), 653–673.
- PASTOOR, M., HENNING, L., NOACK, B. R., KING, R. & TADMOR, G. 2008 Feedback shear layer control for bluff body drag reduction. *Journal of Fluid Mechanics* **608**, 161–196.
- PATHAK, J., HUNT, B., GIRVAN, M., LU, Z. & OTT, E. 2018 Model-free prediction of large spatiotemporally chaotic systems from data: A reservoir computing approach. *Physical Review Letters* **120** (2).
- PAVLIOTIS, G. A. & STUART, A. M. 2007 Parameter estimation for multiscale diffusions. *Journal of Statistical Physics* **127** (4), 741–781.
- PAVLIOTIS, G. A. & STUART, A. M. 2012 *Multiscale methods: averaging and homogenization*. Springer, New York.
- PEHERSTORFER, B. & WILLCOX, K. 2016 Data-driven operator inference for nonintrusive projection-based model reduction. *Computer Methods in Applied Mechanics and Engineering* **306**, 196–215.
- PEINKE, J., TABAR, M. R. R. & WÄCHTER, M. 2019 The Fokker-Planck approach to complex spatiotemporal disordered systems. *Annual Review of Condensed Matter Physics* **10**, 107–132.
- PERLMAN, E., BURNS, R., LI, Y. & MENEVEAU, C. 2007 Data exploration of turbulence simulations using a database cluster. In *Supercomputing SC07*. IEEE.
- PFAFF, T., FORTUNATO, M., SANCHEZ-GONZALEZ, A. & BATTAGLIA, P. W. 2020 Learning mesh-based simulation with graph networks. *arXiv:2010.03409* .

- PICELLA, F., LOISEAU, J.-CH., LUSSEYRAN, F., ROBINET, J.-CH., CHERUBINI, S. & PASTUR, L. 2018 Successive bifurcations in a fully three-dimensional open cavity flow. *J. Fluid Mech.* **844**, 855–877.
- PICKERING, E., RIGAS, G., SCHMIDT, O. T., SIPP, D. & COLONIUS, T. 2021 Optimal eddy viscosity for resolvent-based models of coherent structures in turbulent jets. *Journal of Fluid Mechanics* **917**, A29.
- POPE, S. B. 2000 *Turbulent Flows*. Cambridge University Press.
- PORTWOOD, G. D., DE BRUYN KOPS, S. M., TAYLOR, J. R., SALEHIPOUR, H. & CAULFIELD, C. P. 2016 Robust identification of dynamically distinct regions in stratified turbulence. *J. Fluid Mech.* **807** (R2).
- PRADAS, M., PAVLIOTIS, G. A., KALLIADASIS, S., PAPAGEORGIU, D. T. & TSELUIKO, D. 2012 Additive noise effects in active nonlinear spatially extended systems. *European Journal of Applied Mathematics* **23**, 563–591.
- PRINZ, J.-H., WU, H., SARICH, M., KELLER, B., SENNE, M., HELD, M., CHODERA, J. D., SCHÜTTE, C. & NOÉ, F. 2011 Markov models of molecular kinetics: Generation and validation. *Journal of Chemical Physics* **134**, 174105.
- PROVANSAL, M., MATHIS, C. & BOYER, L. 1987 Bénard-von Kármán instability: Transient and forced regime. *Journal of Fluid Mechanics* **182**, 1–22.
- QIAN, E., KRAMER, B., PEHERSTORFER, B. & WILLCOX, K. 2020 Lift & learn: Physics-informed machine learning for large-scale nonlinear dynamical systems. *Physica D: Nonlinear Phenomena* **406**, 132401.
- RACKAUCKAS, C. & QING, N. 2017 DifferentialEquations.jl—a performant and feature-rich ecosystem for solving differential equations in Julia. *Journal of Open Research Software* **5** (1).
- RAGWITZ, M. & KANTZ, H. 2001 Indispensable finite time corrections for Fokker-Planck equations from time series data. *Physical Review Letters* **87** (25), 254501.
- RAISSI, MAZIAR & KARNIADAKIS, GEORGE EM 2018 Hidden physics models: Machine learning of nonlinear partial differential equations. *Journal of Computational Physics* **357**, 125–141.
- RAISSI, MAZIAR, YAZDANI, ALIREZA & KARNIADAKIS, GEORGE EM 2020 Hidden fluid mechanics: A Navier-Stokes informed deep learning framework for assimilating flow visualization data. *Science* p. eaaw4741.
- ÅKERVIK, ESPEN, BRANDT, LUCA, HENNINGSON, DAN S., HÖPFNER, JÉRÔME, MARXEN, OLAF & SCHLATTER, PHILIPP 2006 Steady solutions of the Navier-Stokes equations by selective frequency damping. *Physics of Fluids* **18**.
- REDDY, SATISH C., SCHMID, PETER J. & HENNINGSON, DAN S. 1993 Pseudospectra of the orr-sommerfeld operator. *SIAM Journal of Applied Mathematics* **53** (1), 15–47.
- REINBOLD, P. A., GUREVICH, D. R. & GRIGORIEV, R. O. 2020 Using noisy or incomplete data to discover models of spatiotemporal dynamics. *Physical Review E* **101** (1), 010203.
- REISS, J., SCHULZE, P., SESTERHENN, J. & MEHRMANN, V. 2018 The shifted proper orthogonal decomposition: A mode decomposition for multiple transport phenomena. *SIAM Journal on Scientific Computing* **40** (3), A1322–A1344.
- REMPFER, D. & FASEL, H. F. 1994 Dynamics of three-dimensional coherent structures in a flat-plate boundary layer. *Journal of Fluid Mechanics* **275**, 257–282.
- RÉNYI, A. 1959 On measures of dependence. *Acta Mathematica Academiae Scientiarum Hungaricae* **10**, 441–451.
- RICHARDSON, LEWIS F. 1922 *Weather Prediction by Numerical Process*. Cambridge University Press.
- RIGAS, G., ESCLAPEZ, L. & MAGRI, L. 2016 Symmetry breaking in a 3d bluff-body wake. Center for Turbulence Research, Stanford University.

- RIGAS, G., MORGANS, A. S., BRACKSTON, R. D. & MORRISON, J. F. 2015 Diffusive dynamics and stochastic models of turbulent axisymmetric wakes. *Journal of Fluid Mechanics* **778**.
- RIGAS, G., MORGANS, A. S. & MORRISON, J. F. 2017 Weakly nonlinear modelling of a forced turbulent axisymmetric wake. *Journal of Fluid Mechanics* **814**, 570–591.
- RIGAS, G., OXLADE, A. R., MORGANS, A. S. & MORRISON, J. F. 2014 Low-dimensional dynamics of a turbulent axisymmetric wake. *Journal of Fluid Mechanics* **755**.
- RIGAS, GEORGIOS, SIPP, DENIS & COLONIUS, TIM 2021 Nonlinear input/output analysis: application to boundary layer transition. *Journal of Fluid Mechanics* **911**.
- RIM, D., MOE, S. & LEVEQUE, R. J. 2018 Transport reversal for model reduction of hyperbolic partial differential equations. *SIAM/ASA J. Uncert. Quant.* **6** (1), 118–150.
- RISKEN, H. 1996 *The Fokker-Planck equation: methods of solution and applications*. Springer-Verlag Berlin Heidelberg.
- ROSENBERG, K. & MCKEON, B. J. 2019 Efficient representation of exact coherent states of the navier–stokes equations using resolvent analysis. *Fluid Dynamics Research* **51** (1), 011401.
- ROSENBERG, K., SYMON, S. & MCKEON, B. J. 2019 Role of parasitic modes in nonlinear closure via the resolvent feedback loop. *Physical Review Fluids* **4**, 052601(R).
- ROSSITER, J. E. 1964 Wind tunnel experiments on the flow over rectangular cavities at subsonic and transonic speeds. *Tech. Rep.*. Ministry of Aviation; Royal Aircraft Establishment; RAE Farnborough.
- RÖSSLER, A. 2010 Runge-Kutta methods for the strong approximation of solutions of stochastic differential equations. *SIAM Journal of Applied Numerical Analysis* **48** (3), 922–952.
- ROWLEY, C. 2005 Model reduction for fluids using balanced proper orthogonal decomposition. *International Journal of Bifurcation and Chaos* **15** (03), 997–1013.
- ROWLEY, C., MEZIĆ, I., BAGHERI, S., SCHLATTER, P. & HENNINGSON, D. S. 2009 Spectral analysis of nonlinear flows. *Journal of Fluid Mechanics* **641**, 115–127.
- ROWLEY, CLARENCE W., COLONIUS, TIM & BASU, AMIT J. 2002 On self-sustained oscillations in two-dimensional compressible flow over rectangular cavities. *Journal of Fluid Mechanics* **455**, 315–346.
- ROWLEY, CLARENCE W., COLONIUS, TIM & MURRAY, RICHARD M. 2004 Model reduction for compressible flows using POD and Galerkin projection. *Physica D* **189**, 115–129.
- ROWLEY, C. W. & DAWSON, S. T. M. 2017 Model reduction for flow analysis and control. *Ann. Rev. Fluid Mech.* **49**, 387–417.
- ROWLEY, C. W. & JUTTIJUDATA, V. 2005 Model-based control and estimation of cavity flow oscillations. In *Proceedings of the 44th IEEE conference on decision and control*.
- ROWLEY, C. W. & MARSDEN, J. E. 2000 Reconstruction equations and the Karhunen–Loève expansion for systems with symmetry. *Physica D: Nonlinear Phenomena* **142** (1-2), 1–19.
- ROWLEY, C. W. & WILLIAMS, D. R. 2006 Dynamics and control of high-Reynolds-number flow over open cavities. *Annual Review of Fluid Mechanics* **38**, 251–276.
- ROWLEY, C. W., WILLIAMS, D. R., COLONIUS, T. & MURRAY, R. M. 2006 Linear models for control of cavity flow oscillations. *Journal of Fluid Mechanics* **547**, 317–330.
- RUBINI, R., LASAGNA, D. & RONCH, A. DA 2020 The ℓ_1 -based sparsification of energy interaction in unsteady lid-driven cavity flow. *Journal of Fluid Mechanics* **905**, A15.
- RUDY, S. H., BRUNTON, S. L., PROCTOR, J. L. & KUTZ, J. N. 2017 Data-driven discovery of partial differential equations. *Science Advances* **3** (4).
- RUELLE, D. & TAKENS, F. 1971 On the nature of turbulence. *Communications in Mathematical Physics* **20** (3), 167–192.

- SAMANTA, D., DUBIEF, Y., HOLZNER, M., SCHÄFER, C., MOROZOV, A. N., WAGNER, C. & HOF, B. 2013 Elasto-inertial turbulence. *Proceedings of the National Academy of Sciences* **110** (26), 10557–10562.
- SAN, O. & MAULIK, R. 2018a Extreme learning machine for reduced order modeling of turbulent geophysical flows. *Physical Review E* **97**, 042322.
- SAN, O. & MAULIK, R. 2018b Neural network closures for nonlinear model order reduction. *Advances in Computational Mathematics* **44** (6), 1717–1750.
- SANCHEZ-GONZALEZ, A., HEES, N., SPRINGENBERG, J. T., MEREL, J., RIEDMILLER, M., HADSELL, R. & BATTAGLIA, P. 2018 Graph networks as learnable physics engines for inference and control. *arXiv:1806.01242*.
- SAXTON-FOX, T. & MCKEON, B. J. 2017 Coherent structures, uniform momentum zones and the streamwise energy spectrum in wall-bounded turbulent flows. *Journal of Fluid Mechanics* **826**.
- SCHAEFFER, H., TRAN, G. & WARD, R. 2017 Extracting sparse high-dimensional dynamics from limited data. *arXiv:1707.08528*.
- SCHLEGEL, M. & NOACK, B. R. 2015 On long-term boundedness of galerkin models. *Journal of Fluid Mechanics* **765**, 325–352.
- SCHLICHTING, H. 1955 *Boundary-Layer Theory*. McGraw-Hill.
- SCHMID, P. J. 2007 Nonmodal stability theory. *Annual Review of Fluid Mechanics* **39**, 129–162.
- SCHMID, P. J. 2010 Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics* **656**, 5–28.
- SCHMID, P. J. & HENNINGSON, D. S. 2001 *Stability and transition in shear flows*. Springer.
- SCHMIDT, M. & LIPSON, H. 2009 Distilling free-form natural laws from experimental data. *Science* **324** (5923), 81–85.
- SCHMIDT, O. T. 2020 Bispectral mode decomposition of nonlinear flows. *Nonlinear Dynamics* **102**, 2479–2501.
- SCHNEIDER, T., STUART, A. M. & WU, J.-L. 2020 Learning stochastic closures using ensemble Kalman inversion. *arXiv:2004.08376*.
- SCHÜTT, K. T., SAUCEDA, H. E., KINDERMANS, P.-J., TKATCHENKO, A. & MULLER, K.-R. 2018 SchNet - a deep learning architecture for molecules and materials. *Journal of Chemical Physics* **148**, 241722.
- SEMINARA, A., ANGELINI, T. E., WILKING, J. N., VLAMAKIS, H., EBRAHIM, S., KOLTER, R., WEITZ, D. A. & BRENNER, M. P. 2012 Osmotic spreading of bacillus subtilis biofilms driven by an extracellular matrix. *Proceedings of the National Academy of Sciences* **109** (4), 1116–1121.
- SHARMA, ATI S. & MCKEON, BEVERLY J. 2013 On coherent structure in wall turbulence. *Journal of Fluid Mechanics* **728**, 196–238.
- SHI, X. D., BRENNER, M. P. & NAGEL, S. R. 1994 A cascade of structure in a drop falling from a faucet. *Science* **265** (5169), 219–222.
- SIEBER, M., PASCHEREIT, C. O. & OBERLEITHNER, K. 2020 Stochastic modelling of a noise driven global instability in a turbulent swirling jet. *arXiv:2007.07555*.
- DE SILVA, B. M., CALLAHAM, J. L., JONKER, J., GOEBEL, N., KLEMISCH, J., McDONALD, D., HICKS, N., KUTZ, J. N., BRUNTON, S. L. & ARAVKIN, A. Y. 2021 Hybrid learning approach to sensor fault detection with flight test data. *AIAA Journal* (1-14).
- DE SILVA, B. M., CHAMPION, K., QUADE, M., LOISEAU, J.-C., KUTZ, J. N. & BRUNTON, S. L. 2020 PySINDY: A Python package for the sparse identification of nonlinear dynamical systems from data. *Journal of Open Source Software* **5** (49), 1–4.
- SIPP, D. & LEBEDEV, A. 2007 Global stability of base and mean flows: a general approach and its applications to cylinder and open cavity flows. *Journal of Fluid Mechanics* **593**, 333–358.

- SIPP, D., MARQUET, O., MELIGA, P. & BARBAGALLO, A. 2010 Dynamics and control of global instabilities in open-flows: a linearized approach. *Applied Mechanics Reviews* **63**.
- SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures, parts i-iii. *Quar. J. Appl. Math* **45** (3), 561–90.
- SMITS, A. J., MCKEON, B. J. & MARUSIC, I. 2011 High-Reynolds number wall turbulence. *Annual Review of Fluid Mechanics* **43**.
- SNYDER, J., CALLAHAM, J. L., BRUNTON, S. L. & KUTZ, J. N. 2021 Data-driven stochastic modeling of coarse-grained dynamics with finite-size effects using langevin regression. *Physica D: Nonlinear Phenomena* **427**, 133004.
- SOIZE, C. 1988 Steady-state solution of Fokker-Planck equation in higher dimension. *Probabilistic Engineering Mechanics* **3** (4), 196–206.
- SONNEWALD, M., WUNSCH, C. & HEIMBACH, P. 2019 Unsupervised learning reveals geography of global ocean regimes. *Earth and Space Science* **6**, 784–794.
- STANLEY, S. & SARKAR, S. 1997 Simulations of spatially developing two-dimensional shear layers and jets. *Theoretical and Computational Fluid Dynamics* **9** (2), 121–147.
- STEWART, G. W. 2001 A Krylov-Schur algorithm for large eigenproblems. *SIAM J. Matrix Anal. Appl.* **23**, 601–614.
- STEWARTSON, K. & STUART, J. T. 1971 A non-linear instability theory for a wave system in plane Poiseuille flow. *Journal of Fluid Mechanics* **48** (3), 529–545.
- STUART, J. T. 1958 On the non-linear mechanics of hydrodynamic stability. *J. Fluid Mech.* **4** (1), 1–21.
- STUART, J. T. 1960 On the non-linear mechanics of wave disturbances in stable and unstable parallel flows Part 1. The basic behavior in plane Poiseuille flow. *Journal of Fluid Mechanics* **9** (3), 353–370.
- STUART, J. T. 1967 On finite amplitude oscillations in laminar mixing layers. *Journal of Fluid Mechanics* **29** (3), 417–440.
- STUART, J. T. 1971 Nonlinear stability theory. *Annual Review of Fluid Mechanics* **3** (347-370).
- SUJITH, R. I., JUNIPER, M. P. & SCHMID, P. J. 2016 Non-normality and nonlinearity in thermoacoustic instabilities. *International Journal of Spray and Combustion Dynamics* **8** (2), 119–146.
- SWIFT, J. & HOHENBERG, P. C. 1977 Hydrodynamic fluctuations at the convective instability. *Physical Review A* **15**, 319–328.
- SWINNEY, H. L. & GOLLUB, J. P. 1981 *Hydrodynamic instabilities and the transition to turbulence*. Springer-Verlag Berlin Heidelberg.
- SYMON, S., ROSENBERG, K., DAWSON, S. T. M. & MCKEON, B. J. 2018 Non-normality and classification of amplification mechanisms in stability and resolvent analysis. *Physical Review Fluids* **3**, 053902.
- SYMON, S., SIPP, D. & MCKEON, B. J. 2019 A tale of two airfoils: resolvent-based modelling of an oscillator versus an amplifier from an experimental mean. *Journal of Fluid Mechanics* **881**, 51–83.
- TABAR, M. REZA RAHMI 2019 *Analysis and data-based reconstruction of complex nonlinear dynamical systems*. Springer International.
- TADMOR, G., LEHMANN, O., NOACK, B. R., CORDIER, L., DELVILLE, J., BONNET, J.-P. & MORZYNSKI, M. 2011 Reduced-order models for closed loop wake control. *Philosophical Transactions of the Royal Society of London A* **369**, 1513–1524.
- TADMOR, G., NOACK, B. R., MORZYNSKI, M. & SIEGEL, S. 2004 Low-dimensional models for feedback flow control. part II: control design and dynamic estimation. *AIAA Flow Control Conference* .

- TAIRA, K., BRUNTON, S. L., DAWSON, S. T. M., ROWLEY, C. W., COLONIUS, T., MCKEON, B. J., SCHMIDT, O. T., GORDEYEV, S., THEOFILIS, V. & UKEILEY, L. S. 2017 Modal analysis of fluid flows: an overview. *AIAA Journal* **55** (12), 4013–4041.
- TAIRA, K. & COLONIUS, T. 2007 The immersed boundary method: A projection approach. *Journal of Computational Physics* **225**, 2118–2137.
- TAIRA, K., HEMATI, M. S., BRUNTON, S. L., SUN, Y., DURAISAMY, K., BAGHERI, SHERVIN, DAWSON, SCOTT & YEH, CHI-AN 2020 Modal analysis of fluid flows: Applications and outlook. *AIAA Journal* **58** (3), 998–1022.
- TAKENS, F. 1981 Detecting strange attractors in turbulence. *Lecture Notes in Mathematics* **898**, 366–381.
- TAYLOR, G. I. 1923 Stability of a viscous liquid contained between two rotating cylinders. *Philosophical Transactions of the Royal Society of London A* **233**, 289–343.
- TAYLOR, GEOFFREY I. 1935 Statistical theory of turbulence: Parts i-iii. *Proceedings of the Royal Society A* **151** (873), 421–464.
- TENNEKES, H. & LUMLEY, J. L. 1972 *A First Course in Turbulence*. MIT Press.
- THEOFILIS, V. 2011 Global linear instability. *Annual Review of Fluid Mechanics* **43**, 319–352.
- THOMAS, N., SMIDT, T., KEARNES, S., YANG, L., LI, L., KOHLHOFF, K. & RILEY, P. 2018 Tensor field networks: rotation- and translation-equivariant neural networks for 3D point clouds. *arXiv:1802.08219*.
- TOWNE, A. 2021 *Space-time Galerkin projection via spectral proper orthogonal decomposition and resolvent modes*. AIAA Scitech 2021 Forum.
- TOWNE, A., SCHMIDT, O. T. & COLONIUS, T. 2018 Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *Journal of Fluid Mechanics* **847**, 821–867.
- TREFETHEN, L. N., TREFETHEN, A. N., REDDY, S. C. & DRISCOLL, T. A. 1993 Hydrodynamic stability without eigenvalues. *Science* **261** (5121), 578–584.
- TU, J. H., ROWLEY, C. W., KUTZ, J. N. & SHANG, J. K. 2014 Spectral analysis of fluid flows using sub-Nyquist-rate PIV data. *Experiments in Fluids* **55**.
- TUNG, T. C. & ADRIAN, R. J. 1980 Higher-order estimates of conditional eddies in isotropic turbulence. *Physics of Fluids* **23** (7).
- UKEILEY, L., CORDIER, L., MANCEAU, R., DELVILLE, J., GLAUSER, M. & BONNET, J. P. 2001 Examination of large-scale structures in a turbulent plane mixing layer. Part 2. Dynamical systems model. *Journal of Fluid Mechanics* **441**, 67–108.
- VAN OVERSCHEE, P. & DE MOOR, B. 1994 N4SID: Subspace algorithms for the identification of combined deterministic-stochastic systems. *Automatica* **30** (1), 75–93.
- VERMA, S., NOVATI, G. & KOUMOUTSAKOS, P. 2018 Efficient collective swimming by harnessing vortices through deep reinforcement learning. *Proc. Nat. Acad. Sci.* **115** (23), 5849–5854.
- VLACHAS, P. R., B., W., WAN, Z. Y., SAPSIS, T. P. & KOUMOUTSAKOS, P. 2018 Data-driven forecasting of high-dimensional chaotic systems with long short-term memory networks. *Proceedings of the Royal Society A* **474** (2213).
- WAN, Z. Y. & SAPSIS, T. P. 2018 Machine learning the kinematics of bubbles in fluid flows. (*Submitted*).
- WAN, Z. Y., VLACHAS, P., KOUMOUTSAKOS, P. & SAPSIS, T. 2018 Data-assisted reduced-order modeling of extreme events in complex dynamical systems. *PloS one* **13** (5), e0197704.
- WANG, Z., AKHTAR, I., BORGGARD, J. & ILIESCU, T. 2012 Proper orthogonal decomposition closure models for turbulent flows: a numerical comparison. *Computer Methods in Applied Mechanics and Engineering* **237**, 10–26.

- WANT, J., OLSSON, S., WEHMEYER, C., PÉREZ, A., CHARRON, N. E., DE FABRITIIS, G., NOÉ, F. & CLEMENTI, C. 2019 Machine learning of coarse-grained molecular dynamics force fields. *ACS Central Science* **5** (5), 755–767.
- WEIDEMANN, J. A. C. & REDDY, S. C. 2000 A MATLAB differentiation matrix suite. *ACM Transactions on Mathematical Software* **26**, 465–519.
- WIGGINS, S. 1990 *Introduction to Applied Nonlinear Dynamical Systems and Chaos*. Springer.
- WINANT, C. D. & BROWAND, F. K. 1974 Vortex pairing: the mechanism of turbulent mixing-layer growth at moderate Reynolds number. *Journal of Fluid Mechanics* **63** (2).
- WU, Z., LEE, J., MENEVEAU, C. & ZAKI, T. 2019 Application of a self-organizing map to identify the turbulent-boundary-layer interface in a transitional flow. *Physical Review Fluids* **4**, 023902.
- XIE, X., MOHEBUJJAMAN, M., REBHOLZ, L. G. & ILIESCU, T. 2018 Data-driven filtered reduced order modeling of fluid flows. *SIAM Journal on Scientific Computing* **40** (3), B834–b857.
- YAMOUNI, S., SIPP, D. & JACQUIN, L. 2013 Interaction between feedback aeroacoustic and acoustic resonance mechanisms in a cavity flow: a global stability analysis. *J. Fluid Mech.* **717**, 134–165.
- YEH, C.-A. & TAIRA, K. 2019 Resolvent-analysis-based design of airfoil separation control. *Journal of Fluid Mechanics* **867**, 572–610.
- YU, Y. H. 1977 Measurements of sound radiation from cavities at subsonic speeds. *Journal of Aircraft* **14** (9), 838–843.
- ZAGAROLA, M. V. & SMITS, A. J. 1998 Mean-flow scaling of turbulent pipe flow. *J. Fluid Mech.* **373**, 33–79.
- ZAKI, T. A. 2013 From streaks and spots and on to turbulence: exploring the dynamics of boundary layer transition. *Flow, turbulence, and combustion* **91** (3), 451–473.
- ZARE, A., JOVANOVIĆ, M. R. & GEORGIU, T. T. 2017 Colour of turbulence. *Journal of Fluid Mechanics* **812**, 636–680.
- ZEBIB, Z. 1987 Stability of viscous flow past a circular cylinder. *Journal of Engineering Mathematics* (21), 155–165.
- ZHANG, L., MYKLAND, P. A. & AÏT-SAHALIA, Y. 2005 A tale of two time scales: determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association* **100** (472), 1394–1411.
- ZOU, H., HASTIE, T. & TIBSHIRANI, R. 2012 Sparse principal component analysis. *Journal of Computational and Graphical Statistics* **15** (2).
- ZWANZIG, R. 2001 *Nonequilibrium Statistical Mechanics*. Oxford University Press.