

THE UNIVERSITY OF CHICAGO

OPTIMIZATION AND CONTROL FOR SCALABLE MODELING AND INFERENCE
IN DYNAMICAL SYSTEMS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

COMMITTEE ON COMPUTATIONAL AND APPLIED MATHEMATICS

BY
HONGLI ZHAO

CHICAGO, ILLINOIS
MARCH 2026

Copyright © 2026 by Hongli Zhao
All Rights Reserved

In dedication to my family

TABLE OF CONTENTS

LIST OF FIGURES	vii
LIST OF TABLES	xii
LIST OF ALGORITHMS	xiii
ACKNOWLEDGMENTS	xiv
ABSTRACT	xvi
1 INTRODUCTION	1
1.1 Literature Review	2
1.2 Overview of Chapter 2	9
1.3 Overview of Chapter 3	10
1.4 Overview of Chapter 4	11
1.5 Overview of Chapter 5	12
2 DISCOVERY OF LINEAR REPRESENTATIONS FOR NONAUTONOMOUS TRANS-	
LATION INVARIANT PROBLEMS	13
2.1 Introduction	13
2.2 Problem formulation	15
2.3 Brief introduction to DMD algorithms	17
2.3.1 Standard DMD	18
2.3.2 Lagrangian DMD	19
2.3.3 Local DMD	22
2.4 Proposed methodology	23
2.5 Error analysis	25
2.5.1 Prediction errors of standard and local DMDs	25
2.5.2 Impact of truncation on DMD error	29
2.5.3 System-specific upper bounds	34
2.6 Numerical experiments	39
2.6.1 Incompressible Navier-Stokes equations	39
2.6.2 Advection-dominated transport	43
2.7 Summary and conclusions	46
3 DATA-DRIVEN CLOSURES AND ASSIMILATION FOR STIFF MULTISCALE	
RANDOM DYNAMICS	49
3.1 Introduction	49
3.2 Reduced-order Method of Distributions	52
3.3 Data Assimilation	58
3.3.1 Nudging	60
3.3.2 Deep Neural Networks	61

3.4	Implementation and Analysis	64
3.4.1	Assimilation Training	65
3.4.2	Computational Complexity	66
3.5	Numerical Experiments	68
3.5.1	Stiff Linear System	68
3.5.2	Power System Cascade Outages	74
3.6	Summary and Conclusions	85
3.7	Appendix	86
3.7.1	Nudging Convergence	86
4	DATA-DRIVEN ESTIMATION OF FAILURE PROBABILITIES IN CORRELATED STRUCTURE-PRESERVING STOCHASTIC POWER SYSTEM MODELS	89
4.1	Introduction	89
4.2	Stochastic Power System Model	91
4.3	Proposed Method	92
4.3.1	Reduced-Order PDF Equations	93
4.3.2	Joint Reduced-Order PDF Equations	95
4.3.3	Approximation of Closure Terms	96
4.3.4	Implementation of the Reduced-Order PDF Method	97
4.4	IEEE Cases	98
4.4.1	Density of Line Energy	100
4.4.2	Joint Dynamics of Line Pairs	101
4.5	Numerical Experiments	103
4.5.1	Experimental Setup	103
4.5.2	Marginal Probability Densities	105
4.5.3	Joint Failure Probabilities	108
4.5.4	Empirical Sample Complexity	110
4.5.5	Mutual Information	112
4.6	Conclusions and Future Work	113
4.7	Appendix	114
4.7.1	Derivations of Reduced-Order PDF Coefficients	114
5	OVERLAPPING SCHWARZ SCHEME FOR LINEAR-QUADRATIC PROGRAMS IN CONTINUOUS-TIME	117
5.1	Introduction	117
5.1.1	Organization	121
5.1.2	Notation	121
5.2	Problem formulation	122
5.3	Continuous-Time Overlapping Schwarz Decomposition	128
5.4	Exponential Decay of Sensitivity in Linear-Quadratic Control	133
5.5	Convergence of Continuous-Time OSD	155
5.6	Numerical Experiment	158
5.7	Conclusion and Future Work	163
5.8	Appendix	165

5.8.1	Gradient Method for Subproblems	165
5.8.2	Appendix: Proof of Theorem 6	169
5.8.3	Proof of Lemma 6	173
6	FUTURE WORK	182
7	REFERENCES	184

LIST OF FIGURES

2.1	Left: Operator 2-norm upper bound (3) for time-shifted data matrix $\mathbf{Y}$, compared with the actual operator 2-norms, for the forced linear system (2.36). Right: Perturbation upper bounds (2.24) as function of the training-dataset size. From largest to smallest: the generic upper bound of Proposition 2, the physics-based upper bound (3), and the actual DMD perturbation.	38
2.2	Spatial distributions of the speed of fluid flow past a cylinder, predicted via the numerical solution of Navier-Stokes equations (top row) and the standard (middle row) and local (bottom row) DMD algorithms, at times $t = 0.15$ (left column), 0.25 (middle column) and 0.5 (right column). The standard DMD (middle row) yielded an unstable solution. All DMD interpolations are obtained non-intrusively from an existing numerical solver.	40
2.3	Left: Relative interpolation errors of the standard and local DMD ($N_T = 100$) with respect to the reference solution. The standard-DMD error grows exponentially due to the inherent model inaccuracy. Right: Relative errors (zoomed on the interval $10^{-8} < \epsilon < 10^2$) of the local DMD with different numbers of sub-intervals, N_T . The error tends to improve as N_T increases.	42
2.4	Spatial distribution of solute concentration $c(\mathbf{x}, t)$ predicted with the numerical (reference) solution of (2.39) and the three DMD variants (standard, Lagrangian, and local Lagrangian), at time $t = 4$	44

2.5	Left: Estimated advection-velocity along the trajectory of center-of-mass (2.40). Right: Temporal evolution of the relative prediction errors $\epsilon(t)$, of the four DMD variants (standard DMD, local DMD, Lagrangian DMD, and local Lagrangian DMD). The interpolation errors are small at early times due to small advection velocity, and remain bounded except for the standard and local DMD variants due to violation of their assumptions (i.e., time-inhomogeneity and Lagrangian dynamics).	45
3.1	(Left) Learned $\hat{\mathcal{R}}$ from $N_{\text{MC}}^{\text{tr}} = 5 \times 10^2$ MC realizations of Equation (3.20) at sparse times $\mathbf{T}_\nu$ with $\nu = 2 \times 10^2$. (Middle) The learned advection coefficient of Equation (3.23). (Right) Evolution of the homogeneous solution $f_{x_1}^{\text{h}}$ to Equation (3.23).	69
3.2	(Left) Temporal evolution of $f_{x_1}^{\text{h}}$ L_1 error against the yardstick f_{MC} for $\nu \in \{1, 2, 5\} \times 10^2$ and $N_{\text{MC}}^{\text{tr}} = 5 \times 10^2$. (Middle) Snapshot of f_{MC} , $f_{x_1}^{\text{h}}$, $\hat{f}_{x_1}^{\text{NR}}$, and $\hat{f}_{x_1}^{\text{DNN}}$ for $\nu = 2 \times 10^2$ at time $t = 7.2$. (Right) Defects $\hat{f}_{x_1}^{\text{d}}$ corresponding to the middle plot for DNN and NR observers. The latter is computed <i>ex post facto</i> as $\hat{f}_{x_1}^{\text{NR}} - f_{x_1}^{\text{h}}$.	71
3.3	Evolution of L_1 error for $\nu \in \{1, 2, 5\} \times 10^2$ and $N_{\text{MC}}^{\text{tr}} = 5 \times 10^2$. (Left) $\hat{f}_{x_1}^{\text{NR}}$ against the assimilated $f_{\text{MC}}^{\text{tr}, \nu}$. Green ticks on the t -axis represent short assimilation periods, i.e., when $\lambda_\nu(t) > 0$. (Middle) $\hat{f}_{x_1}^{\text{NR}}$ against the yardstick f_{MC} . (Right) $\hat{f}_{x_1}^{\text{DNN}}$ against f_{MC} .	73
3.4	Convergence rates, on a log-log scale, for the spatiotemporal L_1 error of the observations $f_{\text{MC}}^{\text{tr}, \nu}$ (left), $\hat{f}_{x_1}^{\text{NR}}$ (middle), and $\hat{f}_{x_1}^{\text{DNN}}$ (right) for $\nu = \{1, 2, 5\} \times 10^2$ as $N_{\text{MC}}^{\text{tr}}$ increases.	74

- 3.5 $N_{\text{MC}}^{\text{tr}} = 2 \times 10^3$ MC realizations (blue dots) of $(\gamma_{12}, \tilde{h}_{12})$ at time $t = 0.1$. (Left) GLLR estimate of $\hat{\mathcal{R}}$ using 10-fold CV for bandwidth selection. (Middle) Data transformed to standard Gaussian variates and corresponding GLLR fit with simple plug-in bandwidth. (Right) Fit from the middle plot transformed back to the original scale to obtain a more accurate estimate of $\hat{\mathcal{R}}$ 79
- 3.6 (Left) Learned regression function $\hat{\mathcal{R}}(Z_{12}, t)$ estimated from $N_{\text{MC}}^{\text{tr}} = 2 \times 10^3$ MC realizations of equation (3.25) at sparse observation times $\mathbf{T}_\nu$ with $\nu = 10^3$. The behavior of $\hat{\mathcal{R}}$ is difficult to distinguish using a single scale (bottom). Hence, we partition the domain about the (red dashed) line $Z_{12} \approx 0$ and plot the left and right sides on difference scales (top). (Middle) The partially separable advection coefficient based on the learned $\hat{\mathcal{R}}$. Likewise, the domain is split about the (gray dotted) line $Z_{12} \approx 0.95$. The left side is omitted on top as it closely resembles the bottom plot. (Right) Evolution of the homogeneous RoPDF $f_{\gamma_{12}}^h$ for $\gamma_{12}(t)$. The evolution is displayed only near the phase space boundaries since the middle portion of the domain always corresponds to relatively low probability states. . . 81
- 3.7 Comparison of the yardstick f_{MC} (solid black), the homogeneous solution $f_{\gamma_{12}}^h$ (dot-dashed red), the nudged observer $\hat{f}_{\gamma_{12}}^{\text{NR}}$ (dotted blue), and the DNN observer $\hat{f}_{\gamma_{12}}^{\text{DNN}}$ (dashed purple) for $\nu = 10^3$ and $N_{\text{MC}}^{\text{tr}} = 2 \times 10^3$ at times $t = 0.005$ (left), 0.1 (middle), and 0.4, (right). Similar to the PDFs in Figure 3.6, the phase space is restricted near the boundaries to emphasize the bimodal behavior at the latter two times. Moreover, at $t = 0.4$ (right), the modes are given on two scales, revealing that the right mode shrinks, but does not vanish. 83

3.8	(Left) The defect solutions $\hat{f}_{\gamma_{12}}^d$ for the DNN and nudged observers at time $t = 0.1$ with sparsity factor $\nu = 10^3$ and $N_{\text{MC}}^{\text{tr}} = 2 \times 10^3$ MC samples. The former is the (prediction of the) DNN in Equation (3.18) while the latter is computed <i>ex post facto</i> as $\hat{f}_{\gamma_{12}}^{\text{NR}} - f_{\gamma_{12}}^h$. (Middle) Temporal evolution of L_1 errors for the homogeneous solution, $f_{\gamma_{12}}^h$, the nudged observer $\hat{f}_{\gamma_{12}}^{\text{NR}}$, and the DNN observer $\hat{f}_{\gamma_{12}}^{\text{DNN}}$. (Right) Convergence rates, on a log-log scale, for the normalized spatiotemporal L_1 errors of the nudged and DNN observers as $N_{\text{MC}}^{\text{tr}}$ increases.	84
4.1	(Top to bottom) Predicted PDF of line energies for WSCC case 9, IEEE case 30, and IEEE case 57, plotted at respective peak times, with exceedance thresholds zoomed in and marked in red. The 1D reduced-order marginal PDF equations (4.11) were solved with $m_{\text{R}} = 5 \times 10^3$ samples, compared with a $m_{\text{KDE}} = 10^4$ KDE benchmark.	106
4.2	Predicted reduced-order joint PDFs (4.26) with $m_{\text{R}} = 5 \times 10^3$ samples, at selected PDE solution times. Particularly, IEEE Case 30 and 57 show non-trivial evolution and correlation structure over time.	107
4.3	Ratio $\hat{\mathbb{P}}/\mathbb{P}$ between predicted and benchmark probabilities for the three cases (left to right). Grey markers correspond to the independence assumption; see Section 4.5.1.	110
4.4	(Top) Convergence of L^1 -norm error from benchmark (4.29), averaged by the number of lines for each case. (Bottom) Aggregate number of samples required for KDE or reduced-order PDF method to reach within $\gamma = 1\%$ L^1 -norm error from the benchmark.	112
4.5	Estimated mutual information (4.34) of WSCC case 9 and IEEE cases 30 and 57 over PDE simulation time up to $T = 10$	114

5.1	(Left) Convergence of the overlapping Schwarz method under varying overlap sizes ($\{1\%, 5\%, 10\%, 20\%, 30\%, 60\%\}$ of the subdomain length). Larger overlaps yield faster convergence. (Right) Observed convergence rate and its exponential fit. For each overlap size, we estimate the convergence rate by averaging the ratios of two consecutive errors, and fit an exponential curve of the theoretical form $c \cdot \exp(-\rho_Z \tau)$. The fitted values are $\hat{c} = 0.98$ and $\hat{\rho}_Z = 0.05$; and we have $\hat{\rho}_Z < 0.02$ at the significance level of 5%. We observe that the convergence rate improves exponentially in the overlap size, which is consistent with Theorem 9. .	161
5.2	(Left) Comparison of numerical solvers in resolving the subproblem Hamiltonian dynamics. The Runge-Kutta method (RK45) yields the highest accuracy per iteration. (Right) Comparison of first 30 Schwarz iterations for different solvers and time step sizes Δt . Adaptive solvers maintain low errors even on coarse grids, while FE accuracy degrades significantly in the stiff regime.	163

LIST OF TABLES

4.1	Tail probabilities at energy-peak times for individual exceedances (4.30) and joint exceedances (4.31). Values from reduced-order PDF equations use $m_R = 5 \times 10^3$ and are compared against a KDE benchmark with $m_{\text{KDE}} = 10^4$. Grey numbers correspond to the independence assumption (4.32).	110
-----	--	-----

LIST OF ALGORITHMS

1	Standard DMD algorithm [33, Ch. 3]	20
2	Local Lagrangian DMD	24
3	Data-driven finite volume scheme with adaptive time-stepping	97
4	Continuous-Time Overlapping Schwarz Decomposition	132
5	Gradient Descent for Subproblem $\mathcal{P}_j([t_j^0, t_j^1]; p_j, q_j)$	165

ACKNOWLEDGMENTS

I would like to begin by expressing my deepest gratitude to my advisor and life mentor, Prof. Mihai Anitescu, for his unwavering guidance, intellectual generosity, and continual support throughout my PhD journey. His high standards for rigor and clarity have profoundly shaped the way I approach both research and life. I am indebted to him for his trust and patience he has granted me during my PhD studies.

I am also deeply grateful to Prof. Guillaume Bal and Prof. Daniel Sanz-Alonso for serving on my dissertation committee. Prof. Bal's courses in partial differential equations and functional analysis during my first year provided a solid analytical foundation for much of my later work. I also benefited greatly from stimulating discussions with Prof. Sanz-Alonso on Bayesian inference. Their thoughtful feedback on my thesis and steadfast encouragement in my career development have been invaluable.

My sincere thanks go to Dr. Daniel Adrian Maldonado and Dr. Emil Constantinescu for hosting and mentoring me at Argonne National Laboratory during 2023–2024 as a Givens Associate. Working in such a world-class research environment has broadened my perspective on scientific inquiry and offered me a glimpse into the life of an academic researcher. The experience and training I received there will continue to guide my future professional pursuits.

I am also grateful to my collaborators—Prof. Daniel Tartakovsky and Dr. Tyler Maltba—whose insight and dedication to research continue to serve as role models for me. I have also received valuable suggestions and constructive feedback from Prof. Sen Na, whose dedication and work ethic have been a constant source of inspiration. I would also like to thank Prof. Yuehaw Khoo and Prof. Michael Lindsey for advising me during my first year and providing rigorous training that helped me navigate the critical transition from undergraduate to doctoral study.

I am thankful as well to my cohort of fellow CCAM PhD students for their friendship, collaboration, and countless conversations that have enriched both my academic and personal

growth at the University of Chicago.

Finally, I owe my deepest gratitude to my family for their unconditional love and support. This dissertation represents the culmination of their faith in me, and I dedicate it to them.

ABSTRACT

In this thesis, we investigate methodologies for discovery, assimilation, and control of (stochastic) dynamical systems, and introduce algorithms that exploit structural properties for efficiency. Data-driven modeling problems in physics and engineering are governed by high-dimensional dynamical systems, where direct simulation and control can become prohibitively expensive.

Chapter 2 begins with the discovery of mathematical descriptors directly from data, focusing on nonautonomous and translation-invariant systems. By combining the Lagrangian frame of reference with locally time-invariant Koopman approximations, we obtain low-rank linear time-varying surrogate models with a priori predictive error estimators. We validate the approach on canonical flow and 2D advective transport problems. In Chapter 3, we shift focus to stochastic systems, and address uncertainty quantification in multiscale models where stiffness impedes statistical convergence. We propose reduced-order probability density equations, which employ closure terms defined as state-space conditional expectations and learned directly through data assimilation strategies, such as nudging. Chapter 4 generalizes this methodology to high-dimensional observables. The methodology is then applied to stochastic power system models to design estimators for extreme-event probabilities, such as cascading line failures. Numerical results demonstrate both stability and efficiency, outperforming naive Monte Carlo simulation and kernel density estimations.

In Chapter 5, we address the problem of control, which is a central problem in the characterization of extreme events. We present an optimize-then-discretize framework for linear-quadratic control governed by time-inhomogeneous dynamics. Our method employs a modified overlapping Schwarz decomposition in continuous-time, which ensures convergence through the exponential decay of sensitivity property in Hamiltonian systems. Unlike discretize-then-optimize approaches, the framework allows flexible numerical integration schemes while preserving stability. This contribution paves way for distributed nonlinear

control as well as deep learning applications, whose analysis is left for future work.

CHAPTER 1

INTRODUCTION

Many phenomena in physics and engineering are governed by dynamical systems, which are described by ordinary (ODE) or partial differential equations (PDE) that evolve in time according to underlying physical or statistical laws. These systems often exhibit nonlinear interactions, are subject to uncertain boundary conditions and forcing, and evolve over long time horizons, making them both expressive yet computationally demanding.

Modern engineering and scientific applications call for advanced computational frameworks capable of addressing two fundamental challenges. The first concerns *model reduction*: how to approximate complex dynamics in low-dimensional, interpretable representations that accelerate model-based simulations while maintaining a controlled (and efficient) trade-off between physical fidelity and computational cost [10, 80]. The second pertains to the task of *control*: how to optimally synthesize, predict, and plan system behavior under physical constraints so as to achieve desired outcomes in interactive, feedback-driven settings [18, 36, 98]. These two aspects form the conceptual backbone of modern computational science, and recur across virtually every field in which high-dimensional dynamical systems must be simulated and optimized: from fluid dynamics [101, 159], materials science [198, 1, 146], power network analysis [56, 221], energy [7], to biological networks [44], financial modeling [177, 29], and, more recently, deep learning [59, 90, 106]. Despite their disciplinary diversity, these applications share a common goal: to model, approximate, and control dynamical behavior in ways that preserve underlying structure while enabling scalable computation.

The aim of the present work is to address the challenges of efficient computations in this joint perspective of model reduction and control, through the lens of four chapters. Chapter 2 and Chapter 3 address deterministic and stochastic aspects of model reduction. Extending the framework developed in the stochastic case, Chapter 4 shifts from model reduction to the task of estimation; and we discuss an important case of control in Chapter 5. Although

historically developed as distinct disciplines—one arising from numerical analysis and the other from nonlinear optimization theory—reduction and control are intrinsically coupled. Effective control depends on reduced representations that capture essential system behavior, while models must maintain the stability, sensitivity, and controllability properties necessary for feedback synthesis. This interplay has become especially prominent in modern data-driven settings, where models are inferred directly from observations and deployed for both prediction and decision making under uncertainty [34, 212, 166].

1.1 Literature Review

In Chapter 2, we begin with the problem of model discovery and reduction for deterministic dynamical systems, focusing on the identification of linear representations of complex processes directly from data. This question is motivated by systems dominated by transport, diffusion, and external forcing, including advection-diffusion equations in fluid mechanics [169], reactive flows [208], and probability diffusion in stochastic dynamical systems [178]. In such settings, spatial and temporal nonautonomous problems—arising from variable advection fields, time-dependent source terms, or externally driven boundary conditions—renders direct numerical simulation or analytical closure intractable [167]. Therefore, the objective is to construct reduced-order models that remain both stable and time-adaptive across evolving regimes, capturing essential dynamics while mitigating the computational cost inherent in such systems. Classical approaches such as proper orthogonal decomposition [101, 17, 194] and Galerkin projection [119] project high-dimensional systems onto low-dimensional subspaces defined by dominant modes. While these techniques yield optimal low-rank approximations in their parametric forms, they depend on linearization around an equilibria and full knowledge of governing equations. These assumptions are difficult to satisfy in the numerical studies of complex and multiscale systems. As a result, their stability and accuracy often degrade when applied to time-dependent source terms.

Over the past decade, a more flexible alternative is offered by Koopman operator theory [35], which represents nonlinear dynamics as linear evolution in an infinite-dimensional space of observables. The dynamic mode decomposition (DMD) [184] provides a finite-dimensional, numerical realization of this construction, which enables modal analysis directly from data without explicit equations. DMD has achieved broad success in uncovering coherent structures in turbulence [76] and plasma physics [108]. At the same time, its global linear assumption makes it sensitive to nonstationarity. When advection, nonlinear forcing, or parameter variation dominate stationary diffusive behavior, standard DMD produces modes that fail to separate temporal scales and cause errors to accumulate rapidly. Variants such as extended DMD [133, 64], kernel DMD [14, 160], and time-delay embeddings [123] greatly improve expressiveness but fails under time-varying forcing. On the other hand, temporal variations of forcing or variable advection further complicate the task of identifying a low-rank operator due to the inherent constant (in time) assumption of the conventional Koopman operator [79]. While online and streaming DMD formulations [226, 92] mitigate the unstable behavior, they do not account well for advection-dominant regimes.

The limitations of existing reduced-order modeling approaches in handling nonautonomous and advection-dominated systems motivate the approach developed in Chapter 2. To address these issues, we propose to enhance the standard approaches by incorporating temporal locality and draw on the insight that advective transport can be captured more faithfully when observables are represented in a Lagrangian frame that moves with the flow. This perspective enables the Koopman operator to be approximated in a form that remains linear in the feature space but adaptive to evolving dynamics.

Chapter 2 discusses the reduced-order modeling of a *deterministic* dynamical system, the next major question underlying practical decision-making is the propagation of uncertainty. Uncertainty may arise from noisy observations, model discrepancy, or structured randomness in boundary or forcing terms. In realistic multiscale systems—such as climate modeling or

biological networks—the number of uncertain degrees of freedom can be enormous, and their interactions are frequently nonlinear [140, 163, 73]. Traditional uncertainty quantification (UQ) techniques, while well established, often struggle in this regime.

While deterministic model reduction focuses on the simulation of state trajectories, many systems of practical interest are driven by stochastic excitations that necessitate a probabilistic description of dynamics, or at least the ensemble moments thereof. In such settings, the goal is to approximate the evolution—or suitable low-order statistics—of the probability density function (PDF) governed by the Fokker-Planck or Liouville equation, rather than tracking individual sample paths. Classical approaches to UQ, including Monte Carlo simulations and stochastic collocation [215, 70], are known to converge slowly and become prohibitively expensive when the underlying dynamics are high-dimensional. To mitigate this cost, PDF-based methods have been developed to propagate the full distribution of system states. Early work on moment-closure approximations [43, 48] introduced tractable surrogate models by truncating infinite-order statistical moments, while spectral representations such as polynomial chaos (PCE) [214] and Karhunen-Loève (KL) expansions [194] provided orthogonal bases in measure-induced Hilbert spaces. Despite their elegance, these global expansions become inappropriate for correlated noise processes and suffer from poor stability when applied to strongly nonlinear or non-Gaussian systems, where the support of the underlying distribution evolves in a highly nontrivial manner. For low-dimensional systems exhibiting these traits, the method of distributions (MoD), comprised of PDF and cumulative distribution function (CDF) methods, has been highly successful via the derivation and learning of closed-form deterministic partial differential equations (PDEs) for joint PDFs/CDFs of system states [142, 143]. The approach has also been adapted to quantify parametric uncertainty in hyperbolic [199, and the references therein] and parabolic [25] PDEs. Its major strength relies on random input fields being treated exactly, which is in contrast to implementations based on KL expansions [31], meaning that, unlike PCEs, the

MoD is well-suited for systems with short-range colored noise but fail to efficiently characterize high-dimensional systems.

In Chapter 3, we develop a principled model-reduction framework for stochastic dynamics that operates directly on the underlying system properties. We study a low-dimensional quantity of interest (QoI) that depends on a high-dimensional, multiscale stochastic system, and aim to describe its probabilistic evolution within the MoD framework. When projected onto the QoI, the governing PDF equation induces conditional expectation-based closure terms that represent the influence of unresolved variables on the reduced dynamics. These conditional closures are analytically intractable in complex nonlinear systems and must therefore be inferred from data. To this end, we incorporate a data assimilation step that learns these closure terms directly from trajectory observations that are consistent between the reduced PDF dynamics and the original high-dimensional system. The resulting hybrid approach provides a semi-analytic and physically interpretable framework for propagating uncertainty in multiscale stochastic systems.

While high-dimensional stochastic systems are inherently difficult to fully characterize, in Chapter 4 we present a practical application that extends the framework of Chapter 3 to vector-valued QoIs arising in large-scale power system models. This chapter marks a conceptual transition from model reduction to statistical estimation of probabilities associated with rare but high-impact events. Large power networks provide a natural and societally critical testbed for stochastic modeling and uncertainty propagation. Their dynamics are intrinsically multiscale, combining fast electromagnetic transients with slower mechanical and economic control loops, and are continuously perturbed by stochastic renewable generation and fluctuating demand. From a reliability perspective, the central challenge lies in quantifying the probability of cascading failures, and Chapter 4 aims to provide a perspective towards this direction.

Traditional reliability assessment techniques in power networks rely on chained condi-

tional simulations, such as $N - k$ contingency [220] and kinetic Monte Carlo (kMC) simulation [182], are conceptually straightforward but have constraints on the noise processes (i.e. convergence is only to hold in the large-deviation limit) and becomes computationally intractable for high-dimensional systems. They fail to capture the nonlinear coupling between uncertain injections and network constraints and scale poorly with system size.

To address the above challenges, the extension in Chapter 4 highlights the observation from Chapter 3 that the probability of each transmission-line failure depends on a scalar function that can be efficiently evaluated (albeit highly nonlinear). We then propagate its PDF directly by expressing it as the solution of a low-dimensional partial differential equation, which we solve numerically. This reduced-order formulation allows us to consider the joint density of multiple transmission lines simultaneously, and capture their conditional dependence structures that are otherwise difficult to simulate using kMC methods. Furthermore, the approach is inherently data-driven: measurements of system states can be assimilated to estimate unknown closure terms in the PDE. Finally, unlike the large-deviation or low-temperature approximations of kMC, the proposed method provides a minimal-assumption and distributionally consistent methodology for tail probability estimation.

While the previous chapters focus on constructing and simulating data-driven reduced-order models, the techniques were ultimately motivated by a broader goal: to predict and plan for rare but critical events in complex dynamical systems under uncertainty. Beyond describing statistical properties, an effective model should also reveal the most likely pathways by which extreme events occur. Recent advances in computational methods have strengthened the link between large deviations theory and optimal control, showing that the so-called instanton path—the least unlikely trajectory leading to a rare excursion in a stochastic system—can be characterized as the solution of an optimal control problem constrained by differential equations [67]. This viewpoint reformulates reliability analysis as a control problem: rather than sampling stochastic dynamics directly or characterizing

their entire probability distribution, one identifies specific noise realizations that drive the system toward failure and employs asymptotic approximations to estimate the corresponding rare-event probability [75]. In recent applications, this approach has proven effective in dramatically reducing the number of required simulations for probability estimation, as well as in related tasks such as importance sampling [201, 202].

Motivated by this perspective, we close in Chapter 5 by considering the solution of long-horizon optimal control problems (OCPs) constrained by linear, time-varying differential equations, and by developing the theoretical foundation for this task. General-purpose computational methods for control remain formulated in discrete time, where long horizons lead to rapidly growing system dimensions and severe scalability bottlenecks (see [95], and references therein). These challenges have motivated the search for structure-preserving optimization methods that reduce the dimensionality of the Karush-Kuhn-Tucker (KKT) systems, including reduced-space Newton-Krylov approaches [171] and matrix-free interior-point methods [74, 47]. Despite these advances, optimization with discretized control variables becomes increasingly expensive as the temporal resolution grows, motivating a formulation of the problem in continuous-time.

In contrast, another promising class of methods is the indirect approach, which derives first-order optimality conditions directly in the continuous setting before introducing any discretization. This formulation expresses the problem through the Hamiltonian system associated with Pontryagin’s Minimum Principle (PMP), and establishes a coupled relationship between control, state, and adjoint variables [45]. By exploiting the adjoint equation to compute gradients efficiently, indirect methods avoid the large matrix numerical differentiation required in discretized formulations and retain the structure of the continuous problem [61]. Over the years, numerical techniques such as multiple-shooting and adjoint-based schemes [19, 85] have been developed to solve these systems alternatively. Because they adhere to an optimize-then-discretize philosophy, indirect methods offer greater flexibility

across temporal and spatial resolutions, making them well-suited for high-fidelity dynamical optimization.

Nonetheless, the foundational challenge in differential equation-constrained control still lies in scalability. Chapter 5 aims to address this issue by developing a framework that exploits structural properties of the underlying dynamical system to enable distributed computation. While the large number of coupled variables and constraints in long-horizon or high-dimensional systems renders centralized solvers computationally prohibitive [51], we develop a continuous-time framework that leverages domain decomposition to achieve scalable and distributed computation. In particular, Schwarz-type methods—originally introduced for elliptic and parabolic partial differential equations [137, 57, 170]—serve as the conceptual foundation of our approach. These methods achieve convergence by iteratively solving local subproblems and exchanging information across overlapping subdomains, a principle we adapt here to the setting of continuous-time optimal control.

In recent years, this idea has been extended to OCPs, where overlapping subdomains correspond to locally coupled control intervals. Allowing overlap between design variables has been shown to enhance convergence and robustness [192, 152]. A key theoretical insight underlying these methods is the *exponential decay of sensitivity*: perturbations introduced at subdomain interfaces diminish rapidly as they propagate inward, ensuring that local corrections synchronize toward the global optimum. This principle forms the foundation of iterative Schwarz-based optimization schemes. Recent developments, such as the fast overlapping temporal decomposition (FOTD) algorithm [153], have demonstrated notable acceleration for discretized OCPs.

While these advances have established the effectiveness of Schwarz-based decomposition in discrete settings, they remain limited by the lack of a continuous-time theoretical foundation. In this chapter, we extend this line of research to continuous-time control, developing an overlapping Schwarz framework that provides the first explicit convergence constants in

this setting.

This thesis contains material from two published papers [228, 145] by the author, along with two submitted manuscripts [229, 230]. In particular, Chapter 2 is based on [228], co-authored with Prof. Daniel Tartakovsky. Chapter 3 is based on [145], co-authored with Dr. Tyler Maltba and Dr. Daniel Adrian Maldonado; and Chapter 4 is based on [229], co-authored with Dr. Tyler Maltba, Dr. Daniel Adrian Maldonado, Dr. Emil Constantinescu, and Prof. Mihai Anitescu. Both Chapter 3 and Chapter 4 were completed during the author’s appointment at Argonne National Laboratory (MCS Division). Chapter 5 is based on [230], co-authored with Prof. Sen Na and Prof. Mihai Anitescu.

For the remainder of this introduction, we provide detailed overview of each chapter. We discuss future directions and extensions in Chapter 6.

1.2 Overview of Chapter 2

In Chapter 2, we begin with the discovery of (linear) mathematical descriptors for physical phenomena directly from observational and simulated data. This question is central to many areas of applied mathematics and physics, yet it is made difficult by two recurring obstacles: the inherent time-inhomogeneity of inputs and translation-invariance of underlying dynamical systems. Motivated by these challenges, we propose to develop a framework that connects Lagrangian dynamic mode decomposition (DMD), proposed in [130], with a locally time-invariant approximation of the Koopman operator. The Lagrangian formulation allows us to construct low-rank linear surrogate models that optimally capture system dynamics in the least-squares sense, while the Koopman approximation provides a principled way to address nonlinearity and non-autonomous behavior. The downstream tasks of constructing a time-varying model relies heavily on linear algebra operations, including the Woodbury matrix identity [84] to update the low-rank least squares-fit linear operator.

In addition to the algorithmic formulation, we establish theoretical bounds on both inter-

polation accuracy and perturbation error, which quantify their dependence on the numerical rank and sampling frequency of the observation data. Furthermore, an improved upper bound is provided for the case of linear time-varying systems. We provide a verification of the derived error bounds using a 2D time-varying linear system (see [226]) and highlight its practical improvement in tightening the *a priori* interpolation error estimates. To demonstrate the practicality of our approach in identifying low-rank surrogate models, we provide detailed numerical experiments on two complex systems: fluid flow past a solid cylinder, governed by the incompressible Navier–Stokes equations; and the transport of a scalar field in 2D, described by the advection-diffusion equation. In both cases, our approach significantly outperforms standard [33] and Lagrangian DMD methods [130], demonstrating superior numerical accuracy and stability under time evolutions.

1.3 Overview of Chapter 3

In this chapter, we consider reduced-order models in stochastically-forced systems and extend a first principles approach towards addressing the challenge of sampling observables that are derived from such systems. Instead of a deterministic governing system with fully predictable behavior, we discuss the approximation of time-varying probability diffusion of the system, from which the observables also exhibit structured, albeit random, behavior.

We introduce a data-driven and physics-informed framework for propagating uncertainty in stiff, multiscale random ordinary differential equations (RODEs) driven by correlated noise. Unlike systems subjected to Gaussian white noise, such problems do not admit a closed-form deterministic equation for the joint probability density function (PDF) of state variables, and any formal PDF representation would be prohibitively high-dimensional. We derive exact but unclosed reduced-order PDF (RoPDF) equations for low-dimensional observables or quantities of interest, where the missing closure terms appear as state-dependent conditional expectations estimated directly from data at sparse observation times.

Furthermore, in stiff, multiscale regimes, sparse grid leads to regression discrepancies that accumulate during RoPDF evolution. To mitigate this issue, we propose to enrich the RoPDF equations with a kinetic-like defect term that is adaptively learned by assimilating sparse, low-fidelity PDF estimates through plug-in density estimators (e.g. kernel density estimation). Within this framework, we investigate two assimilation strategies—nudging and deep neural networks—and demonstrate that both can effectively stabilize and correct the RoPDF evolution. Through numerical experiments, we show that the proposed approach achieves accurate long-range uncertainty propagation, benchmarked against high-fidelity Monte Carlo simulations, while requiring significantly fewer computational resources. We note, however, that while the methodology demonstrates empirical success, the precise characterization of sample complexity remains a challenging, albeit promising, direction.

1.4 Overview of Chapter 4

Chapter 4 extends our focus from low-dimensional observables to vector-valued observables, where correlations among variables play a central role. Such extensions are essential for applications in large-scale engineering systems, where uncertainties propagate across coupled components.

We propose a data-driven approach for propagating uncertainty in stochastic power grid simulations and apply it to the estimation of transmission line failure probabilities. Starting from the Fokker-Planck equation of the full-order continuous Markov model, we derive a reduced-order equation governing the evolution of joint line energy probability density function. Our approach relies on numerically integrating this reduced-order model, which yields sample-efficient estimates of line failure probabilities without resorting to costly full-scale Monte Carlo simulations. We conduct numerical experiments for both scalar- and vector-valued energy functions using the classical multimachine model under spatiotemporally correlated noise perturbations. The results show that our method not only achieves

greater efficiency in computing tail-event probabilities compared with kernel density estimation, but also produces meaningfully richer estimates of joint event capture (measured through Kullback-Liebr divergence) than independent models.

1.5 Overview of Chapter 5

In the final chapter of this thesis, we present an optimize-then-discretize framework for solving linear-quadratic optimal control problems governed by time-inhomogeneous differential equations. Our method employs a novel variant of overlapping Schwarz decomposition, formulated in continuous-time through the Pontryagin Minimum Principle (PMP), which partitions the temporal domain into overlapping intervals and solves associated Hamiltonian systems in parallel. Convergence is ensured by systematically updating boundary conditions across subproblems, a mechanism whose theoretical foundation rests on the exponential decay of sensitivity (EDS), previously established in discrete-time OCPs, extends to the continuous-time setting. Unlike discretize-then-optimize approaches, our framework accommodates flexible numerical integration schemes, including adaptive-time integrators, while preserving stability and convergence guarantees. We illustrate the practicality of our method with a linear-quadratic control example, which paves way for its use in long-range and non-linear control applications.

CHAPTER 2

DISCOVERY OF LINEAR REPRESENTATIONS FOR NONAUTONOMOUS TRANSLATION INVARIANT PROBLEMS

2.1 Introduction

Most, if not all, efforts in mathematical modeling involve the task of identifying constitutive relations or closures from experimental and/or simulated data. Traditionally, this task has been the province of human intelligence, as exemplified in experimental determination of density-pressure and permeability-saturation relations or theoretical derivation of coefficients in the Fokker-Planck and Boltzmann equations. It has now become a domain of machine learning and artificial intelligence [116].

While some of the machine-learning techniques are tailored to specifically “discover” a closure term in, e.g., a partial-differential equation (PDE) [38, 164, etc.], others aim to approximate the whole differential operator or flow map. Among the latter are deep neural networks (DNN) and dynamic mode decomposition (DMD). DNN architectures, e.g., DeepONet, possess high expressiveness and can serve as nonlinear surrogates of PDE-based models, provided sufficient amount of training samples are available [41, 172]. DMD provides an optimal linear approximation of the model and relies on the Koopman operator to account for nonlinearity [33, and references therein]. Lacking precise error estimators, DNN’s performance on any given problem cannot be guaranteed *a priori*. In contrast, DMD is equipped with theoretical estimates of prediction accuracy, e.g., [131]. Its various generalizations are designed to handle advection-dominated phenomena [130], shock waves and discontinuous solutions [132], inhomogeneity of differential operators [134] and a problem’s parametric dependence [135].

Physical features of a PDE model, such as translation invariance and time-dependent coefficients, pose challenges for both DNNs and DMD. For instance, direct-solution DNNs using

soft regularization to enforce advection and mass conservation may lead to ill-conditioning during training [117]. Operator DNNs have also been observed to yield poor performance when the finite data samples are not representative of global transport phenomena [233, 216]. Likewise, standard DMD is also not devoid of shortcomings and fails to cope with sharp gradients [16]. Furthermore, its construction is based on the assumption of time homogeneity (e.g., parameters and/or source terms do not vary in time), which is not suitable for nonautonomous problems.

A prime example of the twin challenges to model discovery is advection-diffusion problems which encapsulate conservation of momentum [191, 195], thermal energy [20], and probability mass [179]. In the diffusion-dominated and intermediary regimes, these problems have been successfully treated via standard reduced-order basis methods including DMD [131, 155] and POD [82]. The advection-dominated regime, characterized by, e.g., high Péclet and Reynolds numbers, complicates not only numerical solution of advection-diffusion equations but also discovery of these equations (or corresponding equation-free models) from observations. Although its convergence properties has been well-studied [114], standard DMD yields quantitatively and qualitatively incorrect solutions, spurring the development of Lagrangian DMD [130].

Nonautonomous dynamical systems pose an additional challenge for surrogate modeling. Typically, they necessitate the deployment of either a global spatiotemporal basis or a time-dependent parameterization, e.g., via Fourier spectral expansions [149, 147]. Examples of such extensions of standard DMD include streaming DMD [92], online time-varying DMD [226], and nonautonomous Koopman operators designed for periodic or quasi-periodic forcings [148].

We reformulate the dynamics in the Lagrangian frame of reference without requiring any knowledge of the underlying PDE or its parameters. The transformed observations are then used to construct a reduced-order model via a piecewise-constant approximation of

the nonautonomous Koopman operator. Our proposed method, local Lagrangian DMD, extends the original Lagrangian DMD by introducing temporal localization, enabling accurate modeling of translation-invariant, time-dependent systems directly from data.

In Section 2.2, we formulate a class of parabolic translation-invariant PDEs with time-dependent inputs and, upon spatial discretization, express them as a nonautonomous dynamical system. Section 2.3 contains a brief description of the Koopman operator theory and the DMD framework for construction of reduced-order representations of PDE-based models. In Section 2.4, we present a local Lagrangian DMD, whose implementation shares relevant features of the time-varying DMD [226] and the Lagrangian DMD [130] to effectively represent translation-invariant PDEs with time-dependent inputs. Upper bounds of both the prediction error of our method and the operator norm error are derived in Section 2.5, as functions of reduction of rank and number of collected snapshots. This theoretical analysis demonstrates that the local Lagrangian DMD is more accurate than either time-varying DMD or Lagrangian DMD alone. A series of numerical experiments, reported in Section 2.6, serve to demonstrate our approach and to compare these error bounds. Main conclusions drawn from our study are summarized in Section 2.7, accompanied by a discussion of the method’s limitations and future research directions.

2.2 Problem formulation

The spatiotemporal evolution of N_u state variables $u(\mathbf{x}, t) = \{u_1, \dots, u_{N_u}\} \in \mathbb{R}^{N_u}$ in the d -dimensional domain $\Omega \subset \mathbb{R}^d$ over the time horizon $[0, T]$ is described by (nonlinear, parabolic) PDEs

$$\frac{\partial u}{\partial t} = -\nabla_{\mathbf{x}} \cdot \mathbf{J}, \quad \mathbf{J} \equiv -D(\mathbf{x}, t, u) \nabla_{\mathbf{x}} u + \mathbf{G}(\mathbf{x}, t, u), \quad (\mathbf{x}, t) \in \Omega \times (0, T]. \quad (2.1a)$$

These equations are subject to appropriate boundary conditions on the bounding surface of Ω and to the initial condition

$$u(\mathbf{x}, 0) = u_0(\mathbf{x}), \quad \mathbf{x} \in \Omega.$$

The functions $D(\mathbf{x}, t, u) : \Omega \times (0, T] \times \mathbb{R}^{N_u} \mapsto \mathbb{R}^+$ and $\mathbf{G}(\mathbf{x}, t, u) : \Omega \times (0, T] \times \mathbb{R}^{N_u} \mapsto \mathbb{R}^d$ are sufficiently smooth with respect to all variables to ensure the existence of a smooth solution $u(\mathbf{x}, t)$. For (2.1) to be well-posed in the classical sense, the scalar diffusion coefficient D is assumed to satisfy a uniform coercivity condition: there exists a constant $\lambda > 0$ such that $D(\mathbf{x}, t, u) \geq \lambda|u|^2$ for all $(\mathbf{x}, t) \in \Omega \times [0, T]$. This ensures that the system remains strictly parabolic and prevents degenerate or hyperbolic behavior (e.g., shock formation in inviscid Burgers' equation). Moreover, the nonlinear advection term $\mathbf{G}$ is assumed to satisfy a polynomial growth condition in u , such as $|\mathbf{G}(\mathbf{x}, t, u)| \leq C(1 + |u|^k)$, $k \geq 1$. Under these assumptions, the system (2.1) admits a unique smooth solution for smooth initial and boundary data [62, Chapter 7].

To simplify the presentation, and without loss of generality, we discretize the computational domain Ω with N_{gp} grid points in each of the d dimensions, yielding a mesh with a total of N_{gp}^d grid points. Let $\mathbf{u}(t) \in \mathcal{M}$ be a solution vector comprising values of the N_u state variables $u(\mathbf{x}, t)$ at the N_{gp}^d discretization points at time t ; and let $\mathcal{M}$ denote a bounded subset of $\mathbb{R}^N$, with $N = N_u N_{\text{gp}}^d$. A suitable discretization of the differential operator in (2.1) leads to a nonautonomous dynamical system of the following general form:

$$\frac{d\mathbf{u}}{dt} = \mathbf{N}(t, \mathbf{u}), \quad \mathbf{u}(0) = \mathbf{u}_0. \quad (2.2)$$

Here, $\mathbf{u}_0 \in \mathcal{M}$ is the solution vector at time $t = 0$; and the discretized differential operator in (2.1) is arranged in a vector of length N , i.e., $\mathbf{N}(t, \mathbf{u}) : \mathbb{R}^+ \times \mathcal{M} \mapsto \mathbb{R}^N$.

For any two times $t \leq t'$, a continuous flow map $\Phi_{t'}(\cdot; t) : \mathbb{R}^N \mapsto \mathbb{R}^N$ is defined via the

integration of (2.2), such that

$$\mathbf{u}(t') = \mathbf{u}(t) + \int_t^{t'} \mathbf{N}(\tau, \mathbf{u}(\tau)) d\tau \equiv \Phi_{t'}(\mathbf{u}(t); t). \quad (2.3)$$

Let us discretize the time horizon $[0, T]$ with $(N_{\text{sn}} + 1)$ points $0 = t_1 < t_2 < \dots < t_{N_{\text{sn}}+1} = T$; values of the state vector $\mathbf{u}(t)$ at these points (i.e., snapshots) are denoted by $\mathbf{u}_i = \mathbf{u}(t_i)$. To simplify the presentation, we use a uniform time step Δt , so that $t_i = i\Delta t$ for $i = 1, \dots, N_{\text{sn}} + 1$. Then, a discrete flow map $\Phi_{\Delta t}(\cdot; t_i)$ is defined as

$$\mathbf{u}_{i+1} = \mathbf{u}_i + \int_{t_i}^{t_{i+1}} \mathbf{N}(\tau, \mathbf{u}(\tau)) d\tau \equiv \Phi_{\Delta t}(\mathbf{u}_i; t_i). \quad (2.4)$$

A goal of DMD is to construct $\Phi_{\Delta t}(\mathbf{u}_i; t_i)$ from the temporal snapshots $\mathbf{u}_i$ ($i = 1, \dots, N_{\text{sn}} + 1$) and to obtain its low-rank approximation, yielding a reduced-order surrogate of (2.1). Alternatively, one can think of $\mathbf{u}(t)$ as a quantity of interest computed via post-processing the numerical solution to (2.1). Since a direct model for the temporal evolution of this quantity of interest is seldom available, DMD can be used to learn this model from the snapshots $\mathbf{u}_i$ ($i = 1, \dots, N_{\text{sn}}$).

2.3 Brief introduction to DMD algorithms

Consider an observable function $g(\mathbf{u}(t)) : \mathcal{M} \mapsto \mathbb{R}^N$. Families of the continuous, $\mathcal{K}_t^{t'}$, and discrete, $\mathcal{K}_i^{\Delta t}$, Koopman operators for the dynamical system (2.4) are defined by the relations

$$\mathcal{K}_t^{t'}\{g(\mathbf{u}(t))\} = g(\mathbf{u}(t')) \quad \text{and} \quad \mathcal{K}_i^{\Delta t}\{g(\mathbf{u}_i)\} = g(\mathbf{u}_{i+1}), \quad (2.5)$$

respectively. Both $\mathcal{K}_t^{t'}$ and $\mathcal{K}_i^{\Delta t}$ are infinite-dimensional operators on the Hilbert space of all observable functions g . They evolve a set of observables along their flow and amount to linear maps even though the dynamical system (2.4) is nonlinear.

DMD approximates the eigenmodes of the Koopman operator in order to identify dominant frequencies and reconstruct the underlying dynamics from temporal snapshots. Let $\mathcal{S}$ denote a training dataset containing N_{sn} snapshots of the observables, $\mathbf{g}_i = g(\mathbf{u}_i)$ with $i = 1, \dots, N_{\text{sn}}$. Informed by (2.5), DMD aims to construct a best-fit linear operator $\mathbf{K}$ such that

$$\mathbf{g}_{i+1} \approx \mathbf{K}\mathbf{g}_i, \quad \forall i = 1, \dots, N_{\text{sn}}.$$

2.3.1 Standard DMD

The standard DMD algorithm constructs $\mathbf{K}$ directly in the solution space, i.e., from $g(\mathbf{u}_i) = \mathbf{u}_i$, by minimizing the mean squared error (MSE),

$$L_{\mathcal{S}}(\mathbf{K}) = \frac{1}{N_{\text{sn}}} \sum_{i=1}^{N_{\text{sn}}} \|\mathbf{u}_{i+1} - \mathbf{K}\mathbf{u}_i\|_2^2. \quad (2.6)$$

If the snapshots $\mathbf{u}_i$ are arranged in the form of two data matrices of size $N \times N_{\text{sn}}$,

$$\mathbf{X} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_{N_{\text{sn}}} \\ | & | & \cdots & | \end{bmatrix} \quad \text{and} \quad \mathbf{Y} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{u}_2 & \mathbf{u}_3 & \cdots & \mathbf{u}_{N_{\text{sn}}+1} \\ | & | & \cdots & | \end{bmatrix},$$

then an exact solution of this minimization problem is

$$\mathbf{K} = \mathbf{Y}\mathbf{X}^\dagger, \quad (2.7)$$

where $\dagger$ denotes the Moore-Penrose pseudoinverse. To compute the pseudoinverse stably and tractably, a truncated SVD is often applied to the data matrix $\mathbf{X}$, namely,

$$\mathbf{X} \approx \mathbf{U}_\rho \mathbf{\Sigma}_\rho \mathbf{V}_\rho^\top.$$

The subscript ρ denotes a prescribed rank, typically determined on a Frobenius-norm error threshold; $\mathbf{V}_\rho \in \mathbb{R}^{N_{\text{sn}} \times \rho}$; and $\mathbf{\Sigma}_\rho \in \mathbb{R}^{\rho \times \rho}$ is a diagonal matrix containing the singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_\rho$ of $\mathbf{X} \in \mathbb{R}^{N \times N_{\text{sn}}}$. The columns of matrix $\mathbf{U}_\rho \in \mathbb{R}^{N \times \rho}$ span a ρ -dimensional subspace, making it computationally efficient first to project the observations on that lower-dimensional subspace, then to compute predictions on this subspace, and finally to transform back to the original state-space [33, Ch. 1]. This procedure is summarized in Algorithm 1.

2.3.2 Lagrangian DMD

Standard DMD algorithms, or other SVD-based methods for order reduction such as principal orthogonal decomposition (POD), fail to honor fundamental physical constraints for translation invariant problems, e.g., solution's non-negativity and mass conservation for advection-dominated transport phenomena [130]. Recasting such (autonomous) problems in a Lagrangian frame of reference and treating the moving Lagrangian grid as an additional observable resolves this challenge, giving rise to Lagrangian DMD.

Characteristics $\mathcal{X}(t; \boldsymbol{\xi})$ for (2.1) are defined elementwise by

$$\frac{d\mathcal{X}(t; \boldsymbol{\xi})}{dt} = \left[\nabla_{\mathbf{x}} D(\mathbf{x}, t, u(\mathbf{x}, t)) + \frac{\partial \mathbf{G}(\mathbf{x}, t, u(\mathbf{x}, t))}{\partial u} \right] \Big|_{\mathbf{x}=\mathcal{X}(t; \boldsymbol{\xi})} \quad (2.9a)$$

where $\boldsymbol{\xi} \in \Omega$ is a characteristic's label, which identifies its starting point at $t = 0$. For the domain Ω discretized with N_{gp}^d grid points $\mathbf{x}_i$ ($i = 1, \dots, N_{\text{gp}}^d$), there are N_{gp}^d characteristics, $\mathcal{X}_i$ ($i = 1, \dots, N_{\text{gp}}^d$), with the i th characteristic originating from (being labeled by) $\boldsymbol{\xi} = \mathbf{x}_i$. Along these characteristics, $u(\mathcal{X}, t)$ satisfies

$$\frac{du}{dt} = \left[\nabla_{\mathbf{x}} \cdot (D(\mathbf{x}, t, u(\mathbf{x}, t)) \nabla_{\mathbf{x}} u(\mathbf{x}, t)) \right]_{\mathbf{x}=\mathcal{X}} \quad u(\mathcal{X}(0), 0) = u_0(\boldsymbol{\xi}).$$

Under the regularity assumptions stated for equation (2.1), the trajectories of (2.9) are

Algorithm 1 Standard DMD algorithm [33, Ch. 3]

Inputs: SVD truncation-error threshold $\epsilon > 0$; and data matrices $\mathbf{X}$ and $\mathbf{Y}$, containing N_{sn} snapshots of the vector of state variables (or observables) of length N :

$$\mathbf{X} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_{N_{\text{sn}}} \\ | & | & \cdots & | \end{bmatrix} \quad \text{and} \quad \mathbf{Y} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{u}_2 & \mathbf{u}_3 & \cdots & \mathbf{u}_{N_{\text{sn}}+1} \\ | & | & \cdots & | \end{bmatrix}$$

1. Compute SVD of $\mathbf{X}$ with truncation rank ρ , $\mathbf{X} \approx \mathbf{U}_\rho \mathbf{\Sigma}_\rho \mathbf{V}_\rho^\top$, such that

$$\left\| \mathbf{X} - \mathbf{U}_\rho \mathbf{\Sigma}_\rho \mathbf{V}_\rho^\top \right\|_F^2 = \frac{\sum_{k=\rho+1}^N \sigma_k^2}{\sum_{k=1}^N \sigma_k^2} < \epsilon$$

2. Compute a low-rank approximation of the Koopman operator:

$$\hat{\mathbf{K}} = \mathbf{U}_\rho^\top \mathbf{Y} \mathbf{V}_\rho \mathbf{\Sigma}_\rho^{-1} \quad (2.8)$$

3. Compute eigendecomposition of $\hat{\mathbf{K}}$:

$$\hat{\mathbf{K}} = \mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^{-1},$$

where $\mathbf{\Lambda}$ is an $\rho \times \rho$ diagonal matrix containing the eigenvalues of $\hat{\mathbf{K}}$

4. Obtain the projected DMD modes:

$$\mathbf{\Phi} = \mathbf{U}_\rho \mathbf{Q}$$

associated with eigenvalues $\lambda_k = \mathbf{\Lambda}_{kk}$.

5. Compute DMD predictions:

$$\mathbf{u}_{\text{DMD}}(t) = \mathbf{\Phi} e^{t\mathbf{\Omega}} \mathbf{\Phi}^\dagger \mathbf{u}_0$$

where $\mathbf{\Omega}$ is a diagonal matrix of DMD frequencies with entries $\omega_k = \ln \lambda_k / \Delta t$. In particular, $\mathbf{u}_{\text{DMD}}(t_k) = \mathbf{\Phi} \mathbf{\Lambda}^k \mathbf{\Phi}^\dagger \mathbf{u}_0$ at time point t_k .

well-defined and the solution $u(\mathbf{x}, t)$ remains smooth for $t \in [0, T]$. While the advective component aligns naturally with Lagrangian trajectories, the diffusive term generally cannot be decoupled and must still be evaluated in the Eulerian frame and interpolated. However, through constructing a data-driven linear best-fit model for the full dynamics, this does not hinder our formulation in advection-dominated regimes.

Except at the initial time $t_1 = 0$, at any time t_k ($k > 1$), the Lagrangian mesh comprising nodes $\{\boldsymbol{\mathcal{X}}_i^k \equiv \boldsymbol{\mathcal{X}}_i(t_k)\}$ ($i = 1, \dots, N_{\text{gp}}^d$) differs from its Eulerian counterpart consisting of nodes $\{\mathbf{x}_i\}$ ($i = 1, \dots, N_{\text{gp}}^d$). Hence, the mapping between the discretized Eulerian solution $\mathbf{u}_k \equiv \mathbf{u}(t_k)$ and the Lagrangian solution $\tilde{\mathbf{u}}_k \equiv \tilde{\mathbf{u}}(t_k)$, i.e., the respective solutions to (2.1) and (2.9) requires interpolation between these two meshes at each time step t_k .

In this work, we use the terms *Lagrangian frame* and *characteristic curves* interchangeably. Specifically, the Lagrangian coordinates are defined via the characteristic flow and are evaluated at each t_k by integrating the ODE system (2.9). In this sense, the Lagrangian frame corresponds to the initial reference frame, defined at $\mathbf{x}_i$ ($i = 1, \dots, N_{\text{gp}}^d$), that moves with the flow.

An observable in Lagrangian DMD is the vector $\mathbf{w}_k := g(\mathbf{u}_k)$ constructed by concatenating the Lagrangian mesh $\boldsymbol{\mathcal{X}}(t_k)$ and the Eulerian solution $\mathbf{u}_k$ (interpolated from the Lagrangian solution $\tilde{\mathbf{u}}_k$), namely

$$\mathbf{w}(t) := \begin{bmatrix} \boldsymbol{\mathcal{X}}(t) \\ \mathbf{u}(t) \end{bmatrix} \in \mathbb{R}^{N_e}, \quad N_e = N_{\text{gp}}^d + N_u N_{\text{gp}}^d. \quad (2.10)$$

The length N_e is the effective state dimension of the resulting dynamical system. Temporal

snapshots $\mathbf{w}_k$ ($k = 1, \dots, N_{\text{sn}}$) are arranged into data matrices

$$\mathbf{X} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{w}_1 & \mathbf{w}_2 & \cdots & \mathbf{w}_{N_{\text{sn}}} \\ | & | & \cdots & | \end{bmatrix}, \quad \mathbf{Y} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{w}_2 & \mathbf{w}_3 & \cdots & \mathbf{w}_{N_{\text{sn}}+1} \\ | & | & \cdots & | \end{bmatrix}. \quad (2.11)$$

Finally, Algorithm 1 is used to build the best linear model for $\mathbf{w}(t)$.

The formulation of state vector $\mathbf{w}(t)$ in (2.10) suffers from the “curse of dimensionality,” as the PDE solution is defined on a d -dimensional spatial grid. Consequently, Lagrangian DMD for advection-dominated phenomena works best for low-dimensional problems. The state dimension is of $\mathcal{O}(N_{\text{gp}}^d)$, i.e., grows exponentially with d . The use of the one-dimensional grid points in each dimension, instead of their tensor product, would reduce the storage complexity to $dN_{\text{gp}} + N_u N_{\text{gp}}^d$. Alternatively, one can adopt model order-reduction techniques, e.g., tensor-based methods [21]. Further discussion of high-dimensional PDE solutions lies outside the scope of this investigation.

2.3.3 Local DMD

Local DMD is designed to deal with time-dependent inputs (of diffusion-dominated problems) by partitioning the simulation horizon $[0, T]$ into N_T non-overlapping intervals $[\tau_{p-1}, \tau_p]$, with $p = 1, \dots, N_T$ and $\tau_0 = 0$, $\tau_{N_T} = T$. For simplicity, each interval contains r snapshots and $N_{\text{sn}} = rN_T$. Local DMD approximates the nonautonomous Koopman operator (2.5) by introducing time-dependence in the linear operator,

$$\mathbf{g}_{i+1} \approx \mathbf{K}(t_i)\mathbf{g}_i \quad (2.12)$$

It is common to treat $\mathbf{K}(t)$ as piecewise constant on each interval $[\tau_{p-1}, \tau_p]$,

$$\mathbf{K}(t) = \sum_{p=1}^{N_T} \mathbf{K}_p \mathcal{I}_{[\tau_{p-1}, \tau_p]} \quad (2.13a)$$

where $\mathcal{I}_{[\tau_{p-1}, \tau_p]}$ is the indicator function for time interval $[\tau_{p-1}, \tau_p]$. The constants $\mathbf{K}_p$ ($p = 1, \dots, N_T$) are determined by solving N_T minimization problems,

$$\min_{\mathbf{K}_1, \dots, \mathbf{K}_{N_T}} L_{\mathcal{S}}(\mathbf{K}(t)) = \min_{\mathbf{K}_1, \dots, \mathbf{K}_{N_T}} \sum_{p=1}^{N_T} L_{\mathcal{S}_p}(\mathbf{K}_p).$$

Here, $\mathcal{S}_i$ is a collection of snapshots available during $[\tau_{p-1}, \tau_p]$, and $\mathcal{S} = \bigcup_{p=1}^{N_T} \mathcal{S}_p$. The resulting operator $\mathbf{K}(t)$ is a local best-fit given by standard DMD on interval $[\tau_{p-1}, \tau_p]$.

Instead of the piece-wise constant approximation of $\mathbf{K}(t)$, one can deploy other parameterizations of $\mathbf{K}(t)$, e.g., basis polynomials or a universal function approximator [173].

2.4 Proposed methodology

Both standard DMD and Lagrangian DMD are designed for either autonomous or periodic dynamical systems (2.2), for which the Koopman operator in (2.5) is constructed on a time-invariant set of solution trajectories. At the same time, local DMD is built to deal with nonautonomous systems, but fails to preserve the translation invariance that characterizes advection-dominated problems. Our transient Lagrangian DMD combines the strengths, and ameliorates the weaknesses, of these two DMD variants. Algorithm 2 implements our method.

Algorithm 2 relies on the piecewise-constant Koopman operator $\mathbf{K}(t)$ in (2.13). Modifications of $\mathbf{K}(t)$, which allow superpositions of DMD frequencies in each time interval, would give rise to other forms of DMD, such as the multi-resolution DMD [120] and the windowed DMD [5].

Algorithm 2 Local Lagrangian DMD

Inputs: Data matrices $\mathbf{X}, \mathbf{Y}$, containing N_{sn} snapshots of enlarged states of dimension N_e .
SVD truncation error threshold $\epsilon > 0$:

$$\mathbf{X} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{g}_1 & \mathbf{g}_2 & \cdots & \mathbf{g}_{N_{\text{sn}}} \\ | & | & \cdots & | \end{bmatrix}, \quad \mathbf{Y} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{g}_2 & \mathbf{g}_3 & \cdots & \mathbf{g}_{N_{\text{sn}}+1} \\ | & | & \cdots & | \end{bmatrix}$$

where $\mathbf{g}_i = g(\mathbf{u}_i) = \mathbf{w}_i$ are the concatenation of PDE solution and Lagrangian grid, as defined in (2.10).

1. For $k = 1, \dots, N_T$, compute a low-rank (rank r_k) approximation of the Koopman operator on interval $[\tau_{k-1}, \tau_k]$ by applying Algorithm 1 to the data matrices

$$\mathbf{X}_k = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{g}_{(k-1)r+1} & \mathbf{g}_{(k-1)r+2} & \cdots & \mathbf{g}_{kr} \\ | & | & \cdots & | \end{bmatrix},$$

$$\mathbf{Y}_k = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{g}_{(k-1)r+2} & \mathbf{g}_{(k-1)r+3} & \cdots & \mathbf{g}_{kr+1} \\ | & | & \cdots & | \end{bmatrix}.$$

2. Obtain DMD modes, compute

$$\Phi^{(k)} = \mathbf{U}_r^{(k)} \mathbf{Q}^{(k)},$$

where $\mathbf{U}^{(k)}$ contains the left singular vectors of data matrix $\mathbf{X}_i$, and $\mathbf{Q}^{(k)}$ contains the eigenvectors of operator $\hat{\mathbf{K}}^{(i)}$, analogous to (2.8) of Algorithm 1, along with eigenvalues: $\Lambda^{(1)}, \dots, \Lambda^{(N_T)}$.

3. Obtain a DMD prediction for $t \in [\tau_k, \tau_{k+1}]$. First, evaluate

$$\mathbf{w}_{\text{DMD}}(\tau_k) = \left(\prod_{i=1}^{k-1} \Phi^{(i)} e^{(\tau_i - \tau_{i-1}) \Omega^{(i)}} [\Phi^{(i)}]^\dagger \right) \mathbf{w}_0.$$

Then, continue to evolve for time $(t - \tau_k)$ and obtain the final prediction:

$$\mathbf{w}_{\text{DMD}}(t) = \Phi^{(k)} e^{(t - \tau_k) \Omega^{(k)}} [\Phi^{(k)}]^\dagger \mathbf{w}_{\text{DMD}}(\tau_k).$$

4. Transform back to Eulerian solution:

$$\mathbf{u}_{\text{DMD}}(t) = g^{-1}(\mathbf{w}_{\text{DMD}}(t)).$$

One could use the incremental SVD updates with adaptive rank truncation to directly update each local operator $\mathbf{K}_p$ to $\mathbf{K}_{p+1}$ in low-rank format [30]. However, the inclusion of the Lagrangian grid in the observable (2.10) translates into the data matrices of dimensions $N_e \gg N_{\text{sn}}$ and with full column rank. The size constraint manifests itself for high-dimensional PDEs. Our numerical experiments did not reveal a significant computational advantage of applying incremental SVD update to computed operators $\mathbf{K}^{(1)}, \dots, \mathbf{K}^{(i)}$. In particular, a direct pseudo-inverse computation in standard DMD involves $\mathcal{O}(N_{\text{sn}}^2 N_e)$ runtime complexity, which is asymptotically more expensive than N_T separate SVD computations, yielding $\mathcal{O}(N_T \rho^2 N_e) = \mathcal{O}(N_{\text{sn}} \rho N_e)$. Incremental SVD updates might yield a minor computational saving if the highest rank of the data matrices in each time interval is bounded by some $\rho' < \rho$, in which case the runtime complexity is $\mathcal{O}(N_T \rho \rho' N_e) = \mathcal{O}(N_{\text{sn}} \rho' N_e)$.

2.5 Error analysis

The prediction accuracy of local DMD depends crucially on a judicious choice of the partition of the simulation time horizon $[0, T]$ (see Section 2.3.3). At present, this choice is often informed by numerical experimentation.

We put the partition selection on a more solid foundation by deriving upper bounds for the pointwise and domain-averaged errors of local DMD (see Section 2.5.1). We also establish upper bounds for perturbations to the learned operator due to truncation of frequencies and deletion of training data (see Section 2.5.2). Finally, for a class of linear dynamical systems, we improve these general bounds by analyzing the norms of time-shifted training data $\mathbf{Y}$ and $\mathbf{X}$ (Section 2.5.3).

2.5.1 Prediction errors of standard and local DMDs

The following theorems establish the relative accuracy of the standard and local DMDs.

Theorem 1 (Pointwise error of local DMD). *Let the operator $\mathbf{N}(t, \cdot) : \mathcal{M} \rightarrow \mathbb{R}^N$ in the system of ODEs (2.2) be uniformly Lipschitz (in time) with constant $L > 0$. For the piecewise constant Koopman operator $\mathbf{K}(t)$ in (2.13) and the flow map $\Phi_{\Delta t}(\cdot; t)$ in (2.4), let local DMD in (2.12) satisfy the following conditions:*

1. *Reconstructed solutions $\mathbf{u}_{DMD}(t)$ belong to the set of trajectories that are generated by the flow map in time Δt , i.e., $\mathcal{M}_{\Delta t} = \{\mathbf{u} \in \mathcal{M} : \Phi_{\Delta t}(\mathbf{u}; t_i) \in \mathcal{M}\}$.*
2. *$\sup_{s \in [0, T]} \|\Phi_{\Delta t}(\cdot; s) - \mathbf{K}(s)\|_{L^\infty(\mathcal{M}_{\Delta t})} < +\infty$, where*

$$\|\Phi_{\Delta t}(\cdot; s) - \mathbf{K}(s)\|_{L^\infty(\mathcal{M}_{\Delta t})} := \sup_{\mathbf{u} \in \mathcal{M}_{\Delta t}} \|\Phi_{\Delta t}(\mathbf{u}; s) - \mathbf{K}(s)\mathbf{u}\|_\infty.$$

Let $\mathbf{u}_n \equiv \mathbf{u}(\tau_n)$ be an exact solution at time τ_n , and $\hat{\mathbf{u}}_n \equiv \mathbf{u}_{DMD}(\tau_n)$ be its local-DMD approximation. Then the pointwise (Euclidean norm) error of local DMD,

$$\mathcal{E}_n = \|\mathbf{u}_n - \hat{\mathbf{u}}_n\|_2^2, \quad n \geq 1, \quad (2.15)$$

is bounded from above by

$$\mathcal{E}_n \leq (1 + e^{\Delta t L})^{N_{sn}} \mathcal{E}_0 + \sum_{j=1}^{N_T} \sum_{l=0}^{r-1} (1 + e^{\Delta t L})^l \|\Phi_{\Delta t}^{(n-jr)-l} - \mathbf{A}_{(n-jr)-(l+1)}\|_\infty^2. \quad (2.16)$$

Here, $\mathbf{A}_k \equiv \mathbf{A}(t_k)$ with $\mathbf{A}(t) \equiv \mathbf{K}(t) - \mathbf{I}_N$ and $\mathbf{I}_N$ denoting the identity matrix of size N . Finally, $\mathcal{E}_0$ denotes the initial error at time τ_0 , such that $\mathcal{E}_0 \equiv 0$ if an analytic form of the initial condition for (2.2) is known.

Proof. The matrix $\mathbf{A}(t)$ in (2.16) arises from rewriting the DMD model (2.12) in residual form, i.e., $\hat{\mathbf{u}}_{k+1} = \mathbf{K}(t_k)\hat{\mathbf{u}}_k$, as $\hat{\mathbf{u}}_{k+1} = \hat{\mathbf{u}}_k + \mathbf{A}_k\hat{\mathbf{u}}_k$.

Accounting for (2.4), the error $\mathcal{E}_n$ in (2.15) is rewritten as

$$\mathcal{E}_n = \|\mathbf{u}_{n-1} + \Phi_{\Delta t}(\mathbf{u}_{n-1}; \tau_{n-1}) - (\hat{\mathbf{u}}_{n-1} + \mathbf{A}_{n-1}\hat{\mathbf{u}}_{n-1})\|_2^2. \quad (2.17)$$

In what follows, we simplify the presentation by writing $\Phi_{\Delta t}(\mathbf{u}(t))$ instead of $\Phi_{\Delta t}(\mathbf{u}(t); t)$, and $\|\cdot\|_\infty$ instead of $\|\cdot\|_{L^\infty(\mathcal{M}_{\Delta t})}$ to denote operator norms when the context is clear. Then, by the triangle inequality,

$$\begin{aligned} \mathcal{E}_n &\leq \|\mathbf{u}_{n-1} - \hat{\mathbf{u}}_{n-1}\|_2^2 + \|\Phi_{\Delta t}(\mathbf{u}_{n-1}) - \mathbf{A}_{n-1}\hat{\mathbf{u}}_{n-1}\|_2^2 \\ &= \mathcal{E}_{n-1} + \|\Phi_{\Delta t}(\mathbf{u}_{n-1}) - \Phi_{\Delta t}(\hat{\mathbf{u}}_{n-1}) + \Phi_{\Delta t}(\hat{\mathbf{u}}_{n-1}) - \mathbf{A}_{n-1}\hat{\mathbf{u}}_{n-1}\|_2^2 \\ &\leq \mathcal{E}_{n-1} + \|\Phi_{\Delta t}(\mathbf{u}_{n-1}) - \Phi_{\Delta t}(\hat{\mathbf{u}}_{n-1})\|_2^2 + \|\Phi_{\Delta t}(\hat{\mathbf{u}}_{n-1}) - \mathbf{A}_{n-1}\hat{\mathbf{u}}_{n-1}\|_2^2. \end{aligned} \quad (2.18)$$

Applying Grönwall's inequality to (2.4), and accounting for L -Lipschitz assumption of the operator $\mathbf{N}$, solutions $\mathbf{u}(t), \hat{\mathbf{u}}(t) \in \mathcal{M}_{\Delta t}$ at any time t satisfy

$$\|\Phi_{\Delta t}(\mathbf{u}) - \Phi_{\Delta t}(\hat{\mathbf{u}})\|_2 \leq e^{\tau L} \|\mathbf{u} - \hat{\mathbf{u}}\|_2, \quad \tau \in [0, \Delta t].$$

Let $\Phi(\cdot; \tau)$ denote the (possibly nonlinear) map $\mathbf{x} \mapsto \Phi(\mathbf{x}; \tau)$. By the uniform boundedness assumption,

$$\mathcal{E}_n \leq (1 + e^{\Delta t L})\mathcal{E}_{n-1} + \|\Phi_{\Delta t}^n - \mathbf{A}_{n-1}\|_\infty^2,$$

where $\Phi_{\Delta t}^k := \Phi_{\Delta t}(\cdot; \tau_k)$, for $k = 0, \dots, N_T$, denotes the discrete operators at temporal

grid points. Repeated application of these inequalities yields:

$$\begin{aligned}
\mathcal{E}_n &\leq (1 + e^{\Delta t L})^2 \mathcal{E}_{n-2} + (1 + e^{\Delta t L}) \|\Phi_{\Delta t}^{n-1} - \mathbf{A}_{n-2}\|_\infty^2 + \|\Phi_{\Delta t}^n - \mathbf{A}_{n-1}\|_\infty^2 \\
&\leq \dots \leq (1 + e^{\Delta t L})^r \mathcal{E}_{n-r} + \sum_{l=0}^{r-1} (1 + e^{\Delta t L})^l \|\Phi_{\Delta t}^{n-l} - \mathbf{A}_{n-(l+1)}\|_\infty^2 \\
&\leq \dots \leq (1 + e^{\Delta t L})^{N_{\text{sn}}} \mathcal{E}_0 + \sum_{j=1}^{N_T} \sum_{l=0}^{r-1} (1 + e^{\Delta t L})^l \|\Phi_{\Delta t}^{(n-jr)-l} - \mathbf{A}_{(n-jr)-(l+1)}\|_\infty^2.
\end{aligned}$$

This completes the proof. $\square$

If the Koopman operator were time-invariant, $\mathbf{K}(t) \equiv \mathbf{K}$, then the upper bound (2.16) reduces to that in [172, Theorem 4.3] and [133, Eq. 3.11].

Since the local DMD incorporates more variations in each prediction interval $[\tau_{p-1}, \tau_p]$ than the standard DMD does, it exhibits better average performance in the least squares sense.

Theorem 2. *For any prescribed sub-interval size r , local DMD in (2.13) is at least as accurate in the least squares sense as standard DMD in (2.7).*

Proof. Let $\mathbf{K}^*$ and $\mathbf{K}^*(t)$ denote the constant-in-time (standard DMD) and stepwise constant (local DMD) solutions of the minimization problem (2.6), respectively. It follows from (2.6) that

$$\begin{aligned}
L_{\mathcal{S}}(\mathbf{K}^*) &= \min_{\mathbf{K}} \frac{1}{N_{\text{sn}}} \sum_{i=1}^{N_{\text{sn}}} \|\mathbf{u}_{i+1} - \mathbf{K} \mathbf{u}_i\|_2^2 \\
&= \min_{\mathbf{K}} \frac{1}{N_T} \sum_{i=1}^{N_T} \frac{1}{r} \sum_{j=1}^r \left\| \mathbf{u}_{n-(ir-j)+1} - \mathbf{K} \mathbf{u}_{n-(ir-j)} \right\|_2^2.
\end{aligned} \tag{2.19}$$

By construction of the minimizer $\mathbf{K}^*$,

$$L_{\mathcal{S}}(\mathbf{K}^*) \geq \frac{1}{N_T} \sum_{i=1}^{N_T} \min_{\mathbf{K}_i} \frac{1}{r} \sum_{j=1}^r \left\| \mathbf{u}_{n-(ir-j)+1} - \mathbf{K}_i \mathbf{u}_{n-(ir-j)} \right\|_2^2 = L_{\mathcal{S}}(\mathbf{K}^*(t)).$$

This completes the proof. $\square$

2.5.2 Impact of truncation on DMD error

Dimensionality reduction in DMD (or POD) typically projects observations onto a low-dimensional space, wherein predictions are computationally cheaper. The projection is done by truncating the singular values of the data matrices $\mathbf{X}$ and $\mathbf{Y}$, thus discarding the contribution from their eigenmodes/directions. The following proposition quantifies the effect of such a projection on the learned DMD operator.

Proposition 1 (Operator 2-norm error of standard DMD). *Consider the data matrices $\mathbf{X}$ and $\mathbf{Y}$ in (2.11) and SVD of the former, $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$. The matrix $\mathbf{\Sigma}$ contains the singular values of $\mathbf{X}$ arranged in non-increasing order, $\sigma_1 \equiv \sigma_{\max} \geq \sigma_2 \geq \dots \geq \sigma_{\text{rank}(\mathbf{X})} \equiv \sigma_{\min}$. Let $\mathbf{X}_\rho = \mathbf{U}_\rho\mathbf{\Sigma}_\rho\mathbf{V}_\rho^\top$ denote a truncated SVD with rank $\rho \leq \text{rank}(\mathbf{X})$, i.e., only the first ρ columns are retained in $\mathbf{U}_\rho, \mathbf{V}_\rho$, and the first ρ singular values are retained in $\mathbf{\Sigma}_\rho$. The operator norm error is bounded from above by*

$$\|\mathbf{K} - \mathbf{K}_\rho\|_2 \leq \frac{\sigma_{\max}(\mathbf{Y})}{\sigma_{\min}(\mathbf{X})}. \quad (2.20)$$

Proof. It follows from (2.7) that $\|\mathbf{K} - \mathbf{K}_\rho\|_2 = \|\mathbf{Y}\mathbf{X}^\dagger - \mathbf{Y}\mathbf{X}_\rho^\dagger\|_2 \leq \|\mathbf{Y}\|_2 \cdot \|\mathbf{X}^\dagger - \mathbf{X}_\rho^\dagger\|_2 = \sigma_{\max}(\mathbf{Y})/\sigma_{\min}(\mathbf{X})$. $\square$

The operator norm bound in Proposition 1 is independent of the rank ρ due to pseudoinverse operation. For a given an arbitrary input $\mathbf{q} \in \mathbb{R}^N$, the pointwise prediction error offers

a more granular bound, e.g.,

$$\begin{aligned} \|\mathbf{K}\mathbf{q} - \mathbf{K}_\rho\mathbf{q}\|_2^2 &\leq \sigma_{\max}^2(\mathbf{Y}) \left\| - \sum_{k=\rho+1}^{\text{rank}(\mathbf{X})} \frac{(\mathbf{u}_k^\top \mathbf{q}) \mathbf{v}_k}{\sigma_k(\mathbf{X})} \right\|_2^2 \\ &= \sigma_{\max}^2(\mathbf{Y}) \sum_{k=\rho+1}^{\text{rank}(\mathbf{X})} \frac{(\mathbf{u}_k^\top \mathbf{q})^2}{\sigma_k^2(\mathbf{X})}. \end{aligned} \quad (2.21)$$

where $\mathbf{u}_k, \mathbf{v}_k$ denote the k th left and right singular vectors of the data matrix $\mathbf{X}$.

To investigate mathematical properties of our local DMD strategy, we view the individual solutions $\mathbf{K}_i$ on time interval $[\tau_{i-1}, \tau_i]$ as a standard DMD solution with fewer observations. This necessitates an analysis of the effect of adding/deleting observations in the training data. We do so by focusing on the operator norm perturbation due to the deletion of the most recent observation; the general case of deleting several most recent observations can be treated via the Sherman-Morrison-Woodbury update formula [88]. Our first intermediate result involves the perturbation of pseudoinverse of a matrix.

Lemma 1 (Update of pseudoinverse). *Consider a data matrix with $N_{sn} \leq N$ columns, $\mathbf{X}_{N_{sn}} \in \mathbb{R}^{N \times N_{sn}}$ with full column rank. Let $\mathbf{X} = [\mathbf{X}_{N_{sn}}, \mathbf{x}] \in \mathbb{R}^{N \times (N_{sn}+1)}$ denote a data matrix with an additional column representing the next snapshot, $\mathbf{x} \in \mathbb{R}^N$. Pseudoinverse of $\mathbf{X}$ is related to pseudoinverse of $\mathbf{X}_{N_{sn}}$ by*

$$\mathbf{X}^\dagger = \begin{bmatrix} \mathbf{X}_{N_{sn}}^\dagger \\ \mathbf{0}_{1 \times N} \end{bmatrix} - c \begin{bmatrix} \mathbf{X}_{N_{sn}}^\dagger \mathbf{x} \\ \mathbf{1}_{1 \times N} \end{bmatrix} ((\mathbf{I} - \mathbf{X}_{N_{sn}} \mathbf{X}_{N_{sn}}^\dagger) \mathbf{x})^\top,$$

where $\mathbf{0}$ and $\mathbf{1}$, respectively, denote vectors in $\mathbb{R}^N$ of all 0's and all 1's, and $\mathbf{0}_{1 \times N} = \mathbf{0}^\top, \mathbf{1}_{1 \times N} = \mathbf{1}^\top$. And the constant is defined:

$$c = \frac{1}{\|\mathbf{x}\|_2^2 - \mathbf{x}^\top \mathbf{X}_{N_{sn}} (\mathbf{X}_{N_{sn}}^\top \mathbf{X}_{N_{sn}})^{-1} \mathbf{X}_{N_{sn}}^\top \mathbf{x}} \geq \frac{1}{\|\mathbf{x}\|_2^2}. \quad (2.22a)$$

The lower bound is attained if $\mathbf{x}$ is orthogonal to the range of $\mathbf{X}_{N_{sn}}$.

Proof. For $\mathbf{X} = [\mathbf{X}_{N_{sn}}, \mathbf{x}]$,

$$(\mathbf{X}^\top \mathbf{X})^{-1} = \begin{bmatrix} \mathbf{X}_{N_{sn}}^\top \mathbf{X}_{N_{sn}} & \mathbf{X}_{N_{sn}}^\top \mathbf{x} \\ \mathbf{x}^\top \mathbf{X}_{N_{sn}} & \|\mathbf{x}\|_2^2 \end{bmatrix}^{-1}$$

and the block matrix inverse formula [54] yields

$$(\mathbf{X}^\top \mathbf{X})^{-1} = \begin{bmatrix} (\mathbf{X}_{N_{sn}}^\top \mathbf{X}_{N_{sn}})^{-1} + c \mathbf{X}_{N_{sn}}^\dagger \mathbf{x} \mathbf{x}^\top (\mathbf{X}_{N_{sn}}^\dagger)^\top & -c \mathbf{X}_{N_{sn}}^\dagger \mathbf{x} \\ -c \mathbf{x}^\top (\mathbf{X}_{N_{sn}}^\dagger)^\top & c \end{bmatrix}, \quad (2.23)$$

where c is defined in (2.22a). By the lemma's assumption, all N_{sn} columns of $\mathbf{X}_{N_{sn}}$ are linearly independent. Hence, $\mathbf{X}_{N_{sn}}^\dagger = (\mathbf{X}_{N_{sn}}^\top \mathbf{X}_{N_{sn}})^{-1} \mathbf{X}_{N_{sn}}^\top$. Post-multiplying (2.23) with $\mathbf{X}^\top$, we obtain

$$\mathbf{X}^\dagger = \begin{bmatrix} \mathbf{X}_{N_{sn}}^\dagger + c \mathbf{X}_{N_{sn}}^\dagger \mathbf{x} \mathbf{x}^\top (\mathbf{X}_{N_{sn}} \mathbf{X}_{N_{sn}}^\dagger)^\top - c \mathbf{X}_{N_{sn}}^\dagger \mathbf{x} \mathbf{x}^\top \\ -c \mathbf{x}^\top (\mathbf{X}_{N_{sn}} \mathbf{X}_{N_{sn}}^\dagger)^\top + c \mathbf{x}^\top \end{bmatrix},$$

which leads directly to (2.22). The inequality in (2.22a) stems from the property

$$\mathbf{x}^\top \mathbf{X}_{N_{sn}} (\mathbf{X}_{N_{sn}}^\top \mathbf{X}_{N_{sn}})^{-1} \mathbf{X}_{N_{sn}}^\top \mathbf{x} \geq 0,$$

which is equal to zero if and only if $\mathbf{X}_{N_{sn}}^\top \mathbf{x} = 0$. □

Proposition 2 (DMD operator error due to column deletion). *Let $\mathbf{K}_{N_{sn}}$ and $\mathbf{K}$ denote DMD operators trained on N_{sn} and $(N_{sn} + 1)$ data snapshots, respectively. Under the same conditions as those of Lemma 1, and for the shifted data matrices $\mathbf{Y}_{N_{sn}} \in \mathbb{R}^{N \times N_{sn}}$ and $\mathbf{Y} = [\mathbf{Y}_{N_{sn}}, \mathbf{y}] \in \mathbb{R}^{N \times (N_{sn} + 1)}$, from (2.11), the DMD operator norm error is bounded from*

above,

$$\|\mathbf{K} - \mathbf{K}_{N_{sn}}\|_2^2 \leq c^2 \|\mathbf{x}\|_2^2 \left[1 + \frac{\|\mathbf{x}\|_2^2}{\sigma_{min}^2} \right] (\sigma_{max}^2 + \|\mathbf{y}\|_2^2) + \frac{\|\mathbf{y}\|_2^2}{\sigma_{min}^2}. \quad (2.24)$$

Here, c is defined in Lemma 1; $\mathbf{y}$ is an additional column representing the next snapshot of $\mathbf{Y}_{N_{sn}}$; and $\sigma_{min}(\mathbf{X}_{N_{sn}})$ and $\sigma_{max}(\mathbf{Y}_{N_{sn}})$ are, respectively, the minimum and maximum singular values of the corresponding matrix. If $\mathbf{x}$ is orthogonal to the range of $\mathbf{X}_{N_{sn}}$, then the bound is refined to

$$\|\mathbf{K} - \mathbf{K}_{N_{sn}}\|_2^2 \leq \frac{\sigma_{max}^2(\mathbf{Y}_{N_{sn}}) + \|\mathbf{y}\|_2^2}{\|\mathbf{x}\|_2^2} + \frac{\sigma_{max}^2(\mathbf{Y}_{N_{sn}}) + 2\|\mathbf{y}\|_2^2}{\sigma_{min}^2(\mathbf{X}_{N_{sn}})}. \quad (2.25)$$

Proof. The 2-norm difference between the DMD operators $\mathbf{K}_{N_{sn}} = \mathbf{Y}_{N_{sn}} \mathbf{X}_{N_{sn}}^\dagger$ and $\mathbf{K} = \mathbf{Y} \mathbf{X}^\dagger$ satisfies the triangle inequality

$$\begin{aligned} \|\mathbf{K} - \mathbf{K}_{N_{sn}}\|_2^2 &= \|\mathbf{Y} \mathbf{X}^\dagger - \mathbf{Y} \hat{\mathbf{X}}_{N_{sn}}^\dagger + \mathbf{Y} \hat{\mathbf{X}}_{N_{sn}}^\dagger - \hat{\mathbf{Y}}_{N_{sn}} \hat{\mathbf{X}}_{N_{sn}}^\dagger\|_2^2 \\ &\leq \|\mathbf{Y}\|_2^2 \|\mathbf{X}^\dagger - \hat{\mathbf{X}}_{N_{sn}}^\dagger\|_2^2 + \|\hat{\mathbf{X}}_{N_{sn}}^\dagger\|_2^2 \|\mathbf{Y} - \hat{\mathbf{Y}}_{N_{sn}}\|_2^2, \end{aligned} \quad (2.26)$$

where

$$\hat{\mathbf{X}}_{N_{sn}}^\dagger \equiv \begin{bmatrix} \mathbf{X}_{N_{sn}}^\dagger \\ \mathbf{0}_{1 \times N} \end{bmatrix} \in \mathbb{R}^{(N_{sn}+1) \times N} \quad \text{and} \quad \hat{\mathbf{Y}}_{N_{sn}} \equiv \begin{bmatrix} \mathbf{Y}_{N_{sn}} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{N \times (N_{sn}+1)}.$$

According to Lemma 1,

$$\mathbf{X}^\dagger - \hat{\mathbf{X}}_{N_{sn}}^\dagger = -c \begin{bmatrix} \mathbf{X}_{N_{sn}}^\dagger \mathbf{x} \\ 1 \end{bmatrix} ((\mathbf{I} - \mathbf{X}_{N_{sn}} \mathbf{X}_{N_{sn}}^\dagger) \mathbf{x})^\top. \quad (2.27)$$

Since $\|\mathbf{X}_{N_{sn}}^\dagger\|_2 = 1/\sigma_{min}(\mathbf{X}_{N_{sn}})$ by singular value decomposition from (2.20), we obtain the

2-norm upper bound,

$$\left\| \begin{bmatrix} \mathbf{X}_{N_{\text{sn}}}^\dagger \mathbf{x} \\ 1 \end{bmatrix} \right\|_2^2 \leq 1 + \frac{\|\mathbf{x}\|_2^2}{\sigma_{\min}^2(\mathbf{X}_{N_{\text{sn}}})}. \quad (2.28)$$

Furthermore, being a projection matrix, its eigenvalues are upper-bounded by 1 [218], implying

$$\|\mathbf{I} - \mathbf{X}_{N_{\text{sn}}} \mathbf{X}_{N_{\text{sn}}}^\dagger\|_2 \leq 1.$$

It follows from (2.27) and (2.28) that

$$\|\mathbf{X}^\dagger - \hat{\mathbf{X}}_{N_{\text{sn}}}^\dagger\|_2^2 \leq c^2 \|\mathbf{x}\|_2^2 \left(1 + \frac{\|\mathbf{x}\|_2^2}{\sigma_{\min}^2(\mathbf{X}_{N_{\text{sn}}})} \right).$$

Combined with the fact that $\|\mathbf{Y}\|_2 = \sigma_{\max}(\mathbf{Y})$ in the proof of Proposition 1, these results transform (2.26) into (2.24). Finally, if $\mathbf{x}$ is orthogonal to $\text{range}(\mathbf{X}_{N_{\text{sn}}})$, such that $\mathbf{X}_{N_{\text{sn}}}^\top \mathbf{x} = 0$, then, according to Lemma 1, $\|\mathbf{x}\|_2^2 = c^{-1}$ and (2.24) reduces to (2.25). $\square$

The bound (2.24) is derived for deletion of a single column/snapshot from the data matrices. An analogous bound can be obtained for addition or deletion of multiple snapshots by using the Woodbury matrix identity [83]. In particular, the bound can be evaluated without explicitly forming the updated operator $\mathbf{K}$, and can be sequentially updated as new data pairs $(\mathbf{x}, \mathbf{y})$ become available.

The results in Proposition 1 and (2) rely only on general linear algebra operations, regardless of the underlying system dynamics.

While the use of short time windows enables the local Lagrangian DMD to capture nonautonomous and transient behaviors that global DMD does not, it also reduces the compressive and interpretive power traditionally associated with global Koopman eigenmodes. In the limit of vanishing window size, the extracted modes may resemble the data itself, thereby limiting generalization. This trade-off is a feature of nonautonomous systems, whose dynamics cannot be captured with a single global linear model. Proposition 2 guides the selection of

window length to balance fidelity with robustness. Thus, the loss in global spectral structure is compensated for by the gain in adaptability to time-varying inputs.

2.5.3 System-specific upper bounds

To illustrate the general formulation in Section 2.2 on a tractable yet representative dynamical system, we consider a linear time-dependent ODE arising from the spatial discretization of the advection-diffusion equation (2.1). This setting captures the essential structure of translation-invariant, nonautonomous dynamics while admitting an explicit solution representation through a flow map. This system allows us to derive sharp, system-specific bounds on the DMD reconstruction error and to precisely characterize the impact of time-dependent forcing and linear operator variation.

Consider a system of time-varying, nonautonomous ODEs

$$\frac{d\mathbf{u}}{dt} = \mathbf{C}(t)\mathbf{u}(t) + \mathbf{f}(t), \quad \mathbf{u}(0) = \mathbf{u}_0, \quad (2.29)$$

whose direct integration yields an explicit flow map (2.2) for any time $t \geq 0$,

$$\mathbf{u}(t) = \Phi_t(\mathbf{u}_0; 0) = \exp\left(\int_0^t \mathbf{C}(s)ds\right)\mathbf{u}_0 + \int_0^t \exp\left(\int_s^t \mathbf{C}(s')ds'\right)\mathbf{f}(s)ds. \quad (2.30)$$

This example also provides an explicit dependence of the time-shifted data matrix $\mathbf{Y}$ on $\mathbf{X}$ via the discretized flow map $\Phi_{\Delta t}$, which gives rise to the following proposition.

Proposition 3 (Upper bound for 2-norm of time-shifted data matrix). *Let $m \geq 1$ be the number of intervals, i.e., $[\tau_0, \tau_1], \dots, [\tau_{m-1}, \tau_m]$, with uniform spacing Δt . Furthermore, let the matrix-valued function $\mathbf{C} : [0, \infty) \rightarrow \mathbb{R}^{N_e \times N_e}$, and $\mathbf{f} : [0, \infty) \rightarrow \mathbb{R}^{N_e}$ be at least piecewise continuous on the time intervals. Finally, let $\mathbf{Y}$ denote a time-shifted (by Δt) version of $\mathbf{X}$, containing snapshots according to the flow map defined in (2.30). Then, the 2-norms of the*

data matrices $\mathbf{Y}$ and $\mathbf{X}$ are related by

$$\|\mathbf{Y}\|_2 \leq \exp\left(\frac{1}{2}\gamma^2\Delta t\right) \sqrt{\frac{mf^2}{\gamma^2} + \sum_{i=1}^m \sigma_i^2(\mathbf{X})}, \quad (2.31a)$$

with the constants

$$\gamma = \max_{1 \leq i \leq m} \gamma_i, \quad \gamma_i = \max_{t \in [\tau_{i-1}, \tau_i]} \|\mathbf{C}(t)\|_2 \quad (2.31b)$$

and

$$f = \max_{1 \leq i \leq m} f_i, \quad f_i = \max_{t \in [\tau_{i-1}, \tau_i]} \|\mathbf{f}(t)\|_2. \quad (2.31c)$$

Proof. Let us define the time-dependent matrix, in accordance with the flow map (2.30) under consideration,

$$\mathbf{M}_{\Delta t}^{(i)} = \mathbf{M}_{\Delta t}(\tau_i) \equiv \exp\left(\int_{\tau_i}^{\tau_i + \Delta t} \mathbf{C}(s) ds, \right) \quad (2.32a)$$

and the time-dependent vector,

$$\mathbf{g}_{\Delta t}^{(i)} = \mathbf{g}_{\Delta t}(\tau_i) \equiv \int_{\tau_i}^{\tau_i + \Delta t} \exp\left(\int_s^{\tau_i + \Delta t} \mathbf{C}(s') ds'\right) \mathbf{f}(s) ds. \quad (2.32b)$$

Then, (2.30) takes the form of an explicit linear transformation

$$\mathbf{u}_{i+1} = \mathbf{M}_{\Delta t}^{(i)} \mathbf{u}_i + \mathbf{g}_{\Delta t}^{(i)}, \quad 1 \leq i \leq m. \quad (2.33)$$

Let

$$\mathbf{X} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_m \\ | & | & \cdots & | \end{bmatrix}, \mathbf{Y} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{u}_2 & \mathbf{u}_3 & \cdots & \mathbf{u}_{m+1} \\ | & | & \cdots & | \end{bmatrix}.$$

Iterative application of the recurrence relation (2.33) to $\mathbf{X}$ yields an explicit dependence for each data column,

$$\mathbf{y}_i = \mathbf{M}_{\Delta t}^{(i)} \mathbf{x}_i + \mathbf{g}_{\Delta t}^{(i)}, \quad 1 \leq i \leq m,$$

so that the time-shifted data matrix $\mathbf{Y}$ is rewritten as

$$\mathbf{Y} = \begin{bmatrix} | & | & \cdots & | \\ \mathbf{M}_{\Delta t}^{(1)} \mathbf{x}_1 + \mathbf{g}_{\Delta t}^{(1)} & \mathbf{M}_{\Delta t}^{(2)} \mathbf{x}_2 + \mathbf{g}_{\Delta t}^{(2)} & \cdots & \mathbf{M}_{\Delta t}^{(m)} \mathbf{x}_m + \mathbf{g}_{\Delta t}^{(m)} \\ | & | & \cdots & | \end{bmatrix}.$$

Its 2-norm is

$$\|\mathbf{Y}\|_2^2 \leq \left\| \begin{bmatrix} \mathbf{M}_{\Delta t}^{(1)} & & & \\ & \mathbf{M}_{\Delta t}^{(2)} & & \\ & & \ddots & \\ & & & \mathbf{M}_{\Delta t}^{(m)} \end{bmatrix} \right\|_2^2 \|\mathbf{X}\|_F^2 + \sum_{i=1}^m \|\mathbf{g}_{\Delta t}^{(i)}\|_2^2. \quad (2.34)$$

Recalling the definitions in (2.32), the terms in the above expression satisfy the following inequalities

$$\|\mathbf{M}_{\Delta t}^{(i)}\|_2^2 \leq \exp \left(\int_{\tau_i}^{\tau_i + \Delta t} \|\mathbf{C}(s)\|_2^2 ds \right) \leq e^{\Delta t} \gamma_i^2, \quad (2.35a)$$

$$\|\mathbf{g}_{\Delta t}^{(i)}\|_2^2 \leq f_i^2 \int_{\tau_i}^{\tau_i + \Delta t} e^{\gamma_i^2(\tau_i + \Delta t - s)} ds = \frac{f_i^2}{\gamma_i^2} \left(e^{\gamma_i^2 \Delta t} - 1 \right), \quad (2.35b)$$

where $i = 1, \dots, m$; and the constants γ_i and f_i are defined in (2.31b) and (2.31c), respectively. Since $\mathbf{C}(t)$ and $\mathbf{f}(t)$ are assumed to be piecewise continuous on each $[t_i, t_{i+1}]$, the

attainability of γ_i and f_i is guaranteed. (The former is because the spectra of $\mathbf{C}(t)$ is a continuous map of time [13, 55].)

Combination of (2.34) and (2.35) yields the error bound (2.31). $\square$

Standard DMD is designed for linear, autonomous, and homogeneous systems, i.e., for ODEs (2.29) with $\mathbf{C}(t) \equiv \mathbf{C}$ for a constant matrix, and $\mathbf{f} \equiv \mathbf{0}$. Suppose $\mathbf{C}$ admits an eigendecomposition $\mathbf{C} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^{-1}$, whose largest eigenvalue is denoted by λ_1 . Then, the solution to these ODEs is

$$\mathbf{u}(t) = \mathbf{Q}e^{t\mathbf{\Lambda}}\mathbf{Q}^{-1}\mathbf{u}_0,$$

and the upper bound in Proposition 3 becomes

$$\|\mathbf{Y}\|_2 \leq \kappa_2(\mathbf{Q})e^{\lambda_1\Delta t}\sigma_{\max}(\mathbf{X}),$$

where $\kappa_2(\cdot)$ denotes the 2-norm condition number.

We illustrate our general (2) and bespoke (2.31) operator norm error bounds on a nonautonomous system modified from [226, Example 5.2],

$$\frac{d\mathbf{u}(t)}{dt} = \mathbf{C}(t)\mathbf{u}(t) + \mathbf{f}(t), \quad \mathbf{u}(0) = [1, 0]^\top, \quad (2.36a)$$

where

$$\mathbf{C}(t) = \begin{bmatrix} 0 & 1 + \epsilon t \\ -1 - \epsilon t & 0 \end{bmatrix}, \quad \mathbf{f}(t) = \begin{bmatrix} -\sin(\epsilon t) \\ \cos(\epsilon t) \end{bmatrix}. \quad (2.36b)$$

Temporal snapshots of $\mathbf{u}(t)$ are generated by solving (2.36) with $\epsilon = 0.1$ on the temporal domain $t \in (0, 1]$ with time step $\Delta t = 10^{-3}$, yielding $N_{\text{sn}} = 1000$ snapshots. We apply the derived bounds (2) and Proposition 3 on the constructed system (2.36), in order to provide an a priori estimate of sensitivity to data deletion of the local DMD operator (2.12).

Figure 2.1 exhibits the impact of dataset size used to train DMD operators on the op-




Figure 2.1: Left: Operator 2-norm upper bound (3) for time-shifted data matrix $\mathbf{Y}$, compared with the actual operator 2-norms, for the forced linear system (2.36). Right: Perturbation upper bounds (2.24) as function of the training-dataset size. From largest to smallest: the generic upper bound of Proposition 2, the physics-based upper bound (3), and the actual DMD perturbation.

erator 2-norm of the time-shifted data matrix $\mathbf{Y}$ and, subsequently, on the upper bound of perturbation error, Eq. (2.24), due to column removal. Incorporation of additional training data stabilizes DMD operators against perturbations. While each local DMD uses fewer data snapshots than standard DMD, a larger training sample size N_{sn} mitigates the perturbation error [136].

When one does not have access to the underlying equations (i.e., when data is available only as images and the underlying system is unknown), perturbations in the DMD operator must be estimated from linear algebra operations. Ignoring the underlying system dynamics results in an overestimate of the range of DMD perturbations (see Figure 2.1), thus yielding an overestimation of snapshots needed to accurately recover a ROM. With additional information about the underlying system, e.g., if a model such as (2.29) is available, one can obtain refined bounds by treating columns of $\mathbf{X}$ and $\mathbf{Y}$ as ordered time-shifts of the initial condition along the flow map (2.3). Both strategies provide a priori estimates of operator 2-norm perturbation, which guide the selection of training dataset size in DMD algorithms.

2.6 Numerical experiments

We use numerical experiments to test the ability of DMD Algorithm 2 to handle translation invariant and nonautonomous problems, incompressible flow past a cylinder (see Section 2.6.1) and advection-dominated transport of a passive scalar (see Section 2.6.2). The unknown temporally varying flow velocity in the latter experiment is estimated from tracking the trajectory of the mode, rendering the procedure fully data-driven and representative of the equation-discovery task.

The accuracy of a DMD prediction, $\mathbf{u}_{\text{DMD}}(t)$, of the “exact” (high-resolution numerical) solution, $\mathbf{u}(t)$, is reported in terms of the relative L^2 -error

$$\epsilon(t) = \frac{\|\mathbf{u}_{\text{DMD}}(t) - \mathbf{u}(t)\|_2}{\|\mathbf{u}(t)\|_2}. \quad (2.37)$$

2.6.1 Incompressible Navier-Stokes equations

Consider fluid flow past an impermeable cylinder. The fluid has density $\rho = 1.0$ (kg/m³) and dynamic viscosity $\nu = 1/600$ kg/(m·s); the choice of parameters yields a dimensionless Reynolds number $\text{Re} = 1200$. Furthermore, the two-dimensional flow domain $\mathcal{D}$ is a rectangle $[0, 2] \times [0, 1]$ with a circular exclusion of radius $R = 0.1$ centered at point $\mathbf{x}_{\text{circ}} = [0.3, 0.5]^\top$. Fluid pressure, $p(\mathbf{x}, t)$, and velocity, $\mathbf{v}(\mathbf{x}, t) = [v_1, v_2]^\top$, satisfy the Navier-Stokes and continuity equations

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \mathbf{v}, \quad \nabla \cdot \mathbf{v} = 0; \quad \mathbf{x} \in \mathcal{D}. \quad (2.38a)$$

These equations are subject to the initial conditions $\mathbf{v}(\mathbf{x}, 0) \equiv [0, 0]^\top$, $p(\mathbf{x}, 0) \equiv 0$. On the rectangle’s exterior, the boundary conditions for pressure p are

$$p(2, x_2, t) = 0, \quad \frac{\partial p}{\partial x_1}(0, x_2, t) = \frac{\partial p}{\partial x_2}(x_1, 0, t) = \frac{\partial p}{\partial x_2}(x_1, 1, t) = 0, \quad (2.38b)$$

and the boundary conditions for the longitudinal, $v_1(x_1, x_2, t)$, and transverse, $v_2(x_1, x_2, t)$, components of the flow velocity $\mathbf{v}(x_1, x_2, t)$ are

$$v_1(0, x_2, t) = 1, \quad \frac{\partial v_1}{\partial x_1}(2, x_2, t) = 0, \quad v_1(x_1, 0, t) = v_1(x_1, 1, t) = 0 \quad (2.38c)$$

and

$$v_2(0, x_2, t) = 0, \quad \frac{\partial v_2}{\partial x_1}(2, x_2, t) = 0, \quad v_2(x_1, 0, t) = v_2(x_1, 1, t) = 0. \quad (2.38d)$$

Finally, the conditions on the surface of the circular inclusions are consistent with its treatment as an impermeable boundary.




Figure 2.2: Spatial distributions of the speed of fluid flow past a cylinder, predicted via the numerical solution of Navier-Stokes equations (top row) and the standard (middle row) and local (bottom row) DMD algorithms, at times $t = 0.15$ (left column), 0.25 (middle column) and 0.5 (right column). The standard DMD (middle row) yielded an unstable solution. All DMD interpolations are obtained non-intrusively from an existing numerical solver.

The quantity of interest (QoI) in this example is the spatiotemporal evolution of the magnitude of the flow velocity, $v(\mathbf{x}, t) = |\mathbf{v}(\mathbf{x}, t)|$, i.e., flow speed. We apply the finite

difference method on a staggered grid [50] to solve the nonlinear initial boundary-value problem in (2.38). The spatial domain $\mathcal{D}$ is discretized with a uniform mesh consisting of $N_{\text{gp}} = 5000$ elements of size $\Delta x_1 = \Delta x_2 = 0.02$; the snapshots used for DMD training are collected every $\Delta t = 1 \times 10^{-3}$. Focusing on the period of vortex-street formation (i.e., transition from steady flow to periodic vortex shedding), we carry out the simulation over time interval $t \in (0, 0.5]$, which results in $M = 500$ training snapshots of the vectorized QoI $\mathbf{u}(t) = \{v(\mathbf{x}_1, t), \dots, v(\mathbf{x}_{N_{\text{gp}}}, t)\}$ whose dimension $N = N_{\text{gp}} = 5000$. The standard DMD strategy is implemented on the entire training set; whereas the local DMD strategy is updated in every sub-interval of size $r = 10$. For each DMD strategy, we set the SVD truncation level to $\epsilon = 1.0 \times 10^{-2}$. The results are then generated and error measured in interpolation mode, respectively, at times $t_k = 5 \times 10^{-4}k, k = 1, \dots, N_{\text{sn}}$, for $N_{\text{sn}} = 1000$ snapshots.

Figure 2.2 provides a visual comparison of the predictions of $v(\mathbf{x}, t)$, at selected times $t = 0.15, 0.25$ and 0.50 , alternatively computed with the numerical solution of (2.38), the standard DMD, and the local DMD (see Algorithm 2). As expected, the standard DMD yields an inaccurate reduced-order model under rapid time changes due to its inherent assumption of time-invariance [131]. In particular, the Navier-Stokes solutions are chaotic and turbulent during vortex-street formation, especially at high Reynolds numbers, posing a significant modeling challenge for ROMs that do not consider temporal interpolations [96], which is the central motivation for considering a time-inhomogeneous modeling approach. The distinct and exponentially-growing patterns produced by the standard DMD (middle row of Figure 2.2) demonstrate a distorted representation of the spectrum resulted from sparse sampling, which leads to unstable frequencies and eigenmode pairs. Similar misrepresented modes are observed and attributed to due to numerical noise in [39]. The numerical error propagates through the solution as time marches. Nevertheless, if the training snapshots include a full periodic solution (i.e., at least one period of vortex shedding), the standard

DMD might capture the dominant frequencies well. In contrast, by updating the operators, even at relatively sparse time steps, the local DMD retains its representational accuracy during the chaotic changes, providing more stability in solution propagation.




Figure 2.3: Left: Relative interpolation errors of the standard and local DMD ($N_T = 100$) with respect to the reference solution. The standard-DMD error grows exponentially due to the inherent model inaccuracy. Right: Relative errors (zoomed on the interval $10^{-8} < \epsilon < 10^2$) of the local DMD with different numbers of sub-intervals, N_T . The error tends to improve as N_T increases.

A more quantitative assessment of the accuracy of these two DMD variants is shown in Figure 2.3 in terms of the relative L^2 -error $\epsilon(t)$. While the error of the local DMD ($N_T = 100$) remains relatively stable, on the order of 10^{-3} , as time t increases, the error of the standard DMD increases nearly exponentially with t (see Figure 2.3a). The latter finding is consistent with the previous observation [130] of the standard DMD's failure to handle translation invariant (advection-dominated) problems even when the inputs do not change with time. On the other hand, the local DMD copes with the translation-invariance challenge without the adoption of the Lagrangian frame of reference. Figure 3b illustrates a numerical study of local DMD with varying number of sub-intervals $N_T = 2, 50, 100, 200$, corresponding to number of snapshots $r = 200, 20, 10, 5$ used to construct each DMD operator locally. A lower error is expected as the operator is updated more frequently, which is in line with the time-varying nature of the fluid flow velocity. However, one needs to strike

a balance between the overall interpolation accuracy and the added computational cost of frequent operator updates.

2.6.2 Advection-dominated transport

Consider two-dimensional advection-diffusion transport of a passive scalar in a time-varying velocity field $\mathbf{v}(t) = [v_1, v_2]^\top$. The spatiotemporal evolution of the passive scalar's concentration, $c(\mathbf{x}, t)$, in simulation domain $\mathcal{D} = [-10, 10] \times [-10, 10]$ is described by an advection-diffusion equation,

$$\frac{\partial c}{\partial t} + \mathbf{v} \cdot \nabla c = D \nabla^2 c, \quad \mathbf{x} \in \mathcal{D}, \quad (2.39a)$$

subject to initial condition $c(\mathbf{x}, 0) = e^{-\mathbf{x} \cdot \mathbf{x}}$, and boundary conditions

$$c(\pm 10, x_2, t) = c(x_1, \pm 10, t) = 0.$$

In the simulations reported below, we set $D = 0.001$ and

$$v_1(t) = 0.5 \cos t, \quad v_2 = -0.4 \sin t. \quad (2.39b)$$

This choice of parameters ensures that the transport phenomenon described by equation (2.39) is advection-dominated. Furthermore, the spatial domain $\mathcal{D}$ is discretized with a uniform mesh consisting of $N_{x_1} = N_{x_2} = 50$ nodes in each direction. Hence, the total number of nodes $\mathbf{x}_i$ ($i = 1, \dots, N$) is $N = 2500$, which is also the dimension of the state vector $\mathbf{u}(t) = \{c_1, \dots, c_N\}$, where $c_i(t) \equiv c(\mathbf{x}_i, t)$. The initial boundary-value problem (2.39) is solved numerically using a modified centered-time, centered-space (Du-Fort Frankel) method [103]. To generate the data, numerical simulation is carried out over the time horizon $t \in [0, 8]$ that is discretized with time step 0.01. If the snapshots $\mathbf{u}_m \equiv \mathbf{u}(t_m)$ ($m = 1, \dots, M$) are collected at each time step, $\Delta t = 0.01$, this yields a total of $M = 800$

training snapshots. The performance is then evaluated in the interpolation mode by estimating $\mathbf{u}_k$ at times $t_k = 5 \times 10^{-3}k, k = 1, \dots, N_{\text{sn}}$, for a total of $N_{\text{sn}} = 1600$ snapshots.




Figure 2.4: Spatial distribution of solute concentration $c(\mathbf{x}, t)$ predicted with the numerical (reference) solution of (2.39) and the three DMD variants (standard, Lagrangian, and local Lagrangian), at time $t = 4$.

In this numerical experiment, we utilize the M snapshots $\mathbf{u}_m$ ($m = 1, \dots, M$) to pursue two complementary goals. The first is to evaluate the relative performance of four DMD variants: standard DMD, Langrangian DMD, local DMD, and local Lagrangian DMD. The second goal is to estimate the advection velocity $\mathbf{v}(t)$ by tracking the trajectory of the plume's center of mass $\bar{\mathbf{x}}(t) = [\bar{x}_1, \bar{x}_2]^\top$, i.e., the mode of the solution $c(\mathbf{x}, t)$:

$$\bar{\mathbf{x}}(t) \equiv \frac{1}{\mathcal{M}} \int_{\mathcal{D}} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} c(\mathbf{x}, t) dx_1 dx_2, \quad \mathcal{M}(t) \equiv \int_{\mathcal{D}} c(\mathbf{x}, t) dx_1 dx_2, \quad (2.40)$$

and using a centered-difference scheme to approximate the first-order time-derivative (i.e. the velocity at the center-of-mass) $\mathbf{v}(t) = [v_1(t), v_2(t)]^\top = d\bar{\mathbf{x}}(t)/dt$ from the M snapshots

$\bar{\mathbf{x}}_m \equiv \bar{\mathbf{x}}(t_m)$. The estimated velocity along the characteristic curve $\bar{\mathbf{x}}(t)$ is visualized from $t = 0$ to $t = 8$ in Figure 2.5.




Figure 2.5: Left: Estimated advection-velocity along the trajectory of center-of-mass (2.40). Right: Temporal evolution of the relative prediction errors $\epsilon(t)$, of the four DMD variants (standard DMD, local DMD, Lagrangian DMD, and local Lagrangian DMD). The interpolation errors are small at early times due to small advection velocity, and remain bounded except for the standard and local DMD variants due to violation of their assumptions (i.e., time-inhomogeneity and Lagrangian dynamics).

Figure 2.4 provides a visual comparison between the reference solution and predictions from three DMD variants at time $t = 4$. Respectively, the standard DMD and Lagrangian DMD were constructed with the entire training set of $M = 800$ snapshots (the Lagrangian DMD also including the two-dimensional moving grid); the local Lagrangian DMD variant, on the other hand, was constructed and updated in sub-intervals of size $r = 30$. As displayed in the figure, the poor performance of the standard DMD is due to its inability to cope with both translation invariance (i.e., dominant advective behavior) and the input's time-dependence. The latter is also the reason for the comparatively large prediction error of the Lagrangian DMD. The local Lagrangian DMD handles well both of the challenges to the ROM construction.

As a quantitative measure, the temporal evolution of the relative prediction error, $\epsilon(t)$, of all four tested DMD strategies is presented in Figure 2.5. The error of the standard

DMD increases exponentially with time t , while the errors of the other three DMD variants remain relatively stable. Thanks to the presence of small diffusion ($D \neq 0$), the local DMD achieves the accuracy comparable to that of its Lagrangian counterparts, even though its error increases sub-exponentially with t . Although not shown here, we found the local DMD to fail in the pure advection regime, $D = 0$. The two Lagrangian DMD variants exhibit stable errors over the entire simulation time frame, with the local Lagrangian DMD (see Algorithm 2) being up to two orders of magnitude more accurate than the Lagrangian DMD.

2.7 Summary and conclusions

We developed a local Lagrangian DMD algorithm for the discovery, from temporal snapshots of either state variables or quantities of interest, of reduced-order models of nonautonomous translation-invariant (advection-dominated) phenomena. The theoretical component of our investigation comprises i) prediction-error guarantee for a time-dependent approximation of the Koopman operator, and ii) analysis of the effects of SVD truncation and number of data snapshots on DMD operator error. The experimental component verifies the performance of our local Lagrangian DMD on two problems, fluid flow past a solid cylinder (incompressible Navier-Stokes equations) and transport of a passive scalar (advection-diffusion equation). Our investigation leads to the following major conclusions:

- Standard DMD fails for advection-dominated problems (see, also, [130]) and/or for nonautonomous problems (problems with time-dependent inputs).
- Lagrangian DMD performs well for translation-invariant problems, but struggles in the presence of time-dependent inputs.
- Local DMD is capable of dealing with temporal variability of model inputs, but struggles in advection-dominated regimes.

- Local Lagrangian DMD is bespoke for nonautonomous translation-invariant problems, and provides a mechanism for real-time control and reduced-order predictions in advection-dominated systems.

In contrast to related techniques for the construction of reduced-order models, our framework is entirely data-driven. For example, like ours, transported snapshot model order reduction (TSMOR) employs Lagrangian coordinates to align the reduced basis with advective features. However, TSMOR requires access to an explicit parametric form of the governing PDE, including its transport coefficients [154].

One of the possible future directions is the identification of closures for characterizing the time-evolution of a quantity of interest that depends on the states of another dynamical system [46]. Instead of relying on an equation-free model, deriving and learning explicit forms of the reduced-order dynamics provides a principled analysis tool for uncertainty propagation and control system design, as well as extrapolation capabilities.

We briefly investigated the possibility of a full data-driven model by assuming the advection coefficients are unknown and estimated by mode-tracking. Although such a method is effective in capturing the macroscopic behavior, it is far from being sufficient for velocities that have nonlinear dependence in both spatial variables and the solution itself.

Future explorations will focus on parameterizations for the advection and diffusion coefficients, which are identified simultaneously as the optimal linear operator is constructed. Such a scenario can be potentially considered in a constrained optimization [161] or Bayesian inversion setting [109].

Further reduction of computational complexity is another possible path of future exploration due to the curse of dimensionality for advection-dominated problems associated with moderate- to high-dimensional datasets. An added layer of dimensionality reduction must be adopted in such cases where storing and operating with data snapshots and the Lagrangian moving grid are intractable. A potential solution in the DMD setting is by using low-rank

tensor-networks to approximate multidimensional linear operators [113, 53].

CHAPTER 3

DATA-DRIVEN CLOSURES AND ASSIMILATION FOR STIFF MULTISCALE RANDOM DYNAMICS

3.1 Introduction

Randomness is inherent to most, if not all, complex phenomena described by ordinary differential equations (ODEs)—it enters such models in two ways (a) stochastic forcing terms accounting for internally generated or externally imposed “sub-grid” fluctuations (i.e., noise), and (b) probabilistic representations of uncertain coefficients and initial/boundary data. Owing to simplicity of implementation and parallelizability, multilevel Monte Carlo (MC) simulations [71] and its variants (e.g., see [200] and the references therein) remain as common approaches for solving stochastic systems. uncertainty quantification (UQ) of random ODEs (RODEs) and stochastic differential equations (SDEs). However, MC simulations shed little light on a system’s probabilistic dynamics and are often burdened by slow convergence rates, requiring significant computational resources.

The search for efficient alternatives has led to the development of moment ODEs (MODEs) [28], polynomial chaos expansions (PCEs) [206], Mori-Zwanzig formalism (MZF) [42], and the method of distributions (MoD) [219], each having its strengths and weaknesses. For example, MODEs limit random inputs to be Gaussian or of small variation and are capable of providing only a few statistical moments [28], which are usually not sufficient to characterize a system’s probabilistic nature. PCEs do give rise to probability density functions (PDFs) by using a finite number of uncorrelated random variables to approximate temporally varying random inputs, e.g., Karhunen-Loève (KL) expansions. However, they are inappropriate for models whose random sources have short-range correlations [210]. MZF, on the other hand, is a step in the right direction for high-dimensional RODEs by seeking non-Markovian reduced-order Langevin equations for low-dimensional observables/quantities of interest (QoIs). Such

equations contain nonlocal/memory and noise terms, where the former requires a closure approximation and the latter must satisfy the second fluctuation-dissipation theorem (FDT) [42, 68, 138]. This is problematic for many stiff, multiscale systems lacking scale separation, which is typically necessary for efficient memory kernel approximation (see, e.g., [129]).

The aforementioned approaches fall short because they cannot simultaneously tackle high-dimensionality, stiffness, multiple scales, and colored noise. For low-dimensional systems exhibiting these traits, the MoD, comprised of PDF and cumulative distribution function (CDF) methods, has been highly successful via the derivation and learning of closed-form deterministic partial differential equations (PDEs) for joint PDFs/CDFs of system states [142, 143, 211]. The approach has also been adapted to quantify parametric uncertainty in hyperbolic [181, and the references therein] and parabolic [24] PDEs. Its major strength relies on random input fields being treated exactly, which is in contrast to implementations based on KL expansions [206], meaning that, unlike PCEs, the MoD is well-suited for systems with short-range colored noise. In what follows, we limit our study to PDF methods.

The standard MoD approach is infeasible for high-dimensional RODEs since it results in a PDE with as many spatial dimensions as there are states/equations in the system. While there have been advancements in numerical integration of high-dimensional PDEs, including ideas from multilevel MC [60], deep learning [86, 58], and tensor-networks [180], most techniques are tailored to specific applications and do not generalize well for considerably high-dimensional systems unless the dynamics are relatively straightforward or inherently lie on a low-dimensional manifold. While there have been advancements in numerical integration of high-dimensional PDEs [180], they typically are not well-suited for complex multiscale dynamics of exceptionally large dimension. Similar to MZF, we instead consider low-dimensional QoIs, albeit directly for their PDF dynamics instead of their Langevin ones, leading to the reduced-order MoD. More precisely, we derive exact, albeit unclosed, reduced-order PDF (RoPDF) equations for low-dimensional QoIs. Unlike MZF, the RODE noise

need not satisfy the second FDT, nor does a memory kernel need approximating. Moreover, we allow stiffness and multiscale properties to hold for the entire system, not just the subspace orthogonal to the QoI, and we do not require scale separation, i.e., a fast/slow variable decomposition as in [139].

Unclosed terms in our RoPDF equations take the form of state-dependent conditional expectations, henceforth referred to as regression functions, and are estimated from state data (via parametric or non-parametric methods) at discrete times. Such estimation may require a fully observable system, i.e., all system states, but is generally application dependent as seen in our experiments. When a QoI has slow dynamics with respect to the observation time intervals, the learned regression functions produce negligible discrepancies, and solutions to the resulting RoPDF equations are accurate, assuming enough data has been injected. When the RODE is at least partially separable with respect to the QoI (in the sense of [31]), parts of the regression functions are known analytically. Thus, part of the RoPDF equation is known *a priori* and is physics-informed, which is a means for variance reduction. Knowing a portion of the RoPDF dynamics exactly reduces the need for large amounts of data during learning. This was studied in [31] for ODEs with random initial conditions.

For systems of stiff, multiscale RODEs, especially highly oscillatory systems lacking scale separation, QoIs may vary rapidly over observation windows, making the available state data temporally sparse with respect to the QoI's timescale. In this setting, regression discrepancies amount to model/PDE misspecification and compound during RoPDF evolution, producing inaccurate RoPDFs. When state data is synthetically generated via MC simulations, data availability for regression may be limited due to large computational costs associated with the numerical integration of (possibly high-dimensional) stiff RODEs and necessarily coarse time steps. To overcome these challenges, we introduce an *a priori* unknown source term to the RoPDF equation for capturing model defects. It is inferred by post-processing the limited QoI data via fast, robust kernel density estimation (KDE) to form low-fidelity, tem-

porally sparse RoPDF estimates, which are then assimilated into the RoPDF equation. We give a head-to-head comparison of two assimilation procedures: nudging (a.k.a., Newtonian relaxation (NR)) and deep neural networks (DNNs). The former dynamically steers the RoPDF equation solution (a.k.a, the observer) towards the RoPDF estimates via a tuned (finite) relaxation rate. The latter, however, can be interpreted as instantaneous, albeit not dynamic, relaxation.

Chapter 3 is organized as follows. We introduce RODEs and derive deterministic PDEs for their RoPDFs in Section 3.2, and Section 3.3 discusses the data-assimilation procedures nudging and DNNs used for RoPDF inference. Details on numerics, training, and computational complexity are given in Section 3.4. Numerical experiments are presented in Section 3.5 for a stiff linear system and a power grid model of transmission failures, both driven by Ornstein-Uhlenbeck (OU) noise. Concluding remarks and future directions are summarized in Section 3.6.

3.2 Reduced-order Method of Distributions

Consider the RODE system

$$\frac{d\mathbf{x}(t)}{dt} = \mathbf{v}(\mathbf{x}(t, \omega), t; \boldsymbol{\xi}(t, \omega)), \quad \mathbf{x}(0, \omega) = \mathbf{x}^0(\omega), \quad (3.1)$$

to be solved on a time interval $(0, T_f]$ and holds for almost every $\omega \in \Omega$, where $(\Omega, \mathcal{F}, \mathbb{P})$ is an appropriate probability space. The solution $\mathbf{x}(t, \omega) : [0, T_f] \times \Omega \rightarrow \mathbb{R}^N$ is an $\mathbb{R}^N$ -valued stochastic process with the initial state $\mathbf{x}^0$, defined as an N -dimensional random vector with joint PDF $f_{\mathbf{x}^0}(\mathbf{X}) : \mathbb{R}^N \rightarrow \mathbb{R}^+$. The phase space for Equation (3.1) is taken as $\mathbb{R}^N$ for notational convenience; however, this can be altered, with little effect on the arguments below, to account for almost surely bounded processes. The given deterministic function $\mathbf{v} = [v_1, \dots, v_N]^\top : \mathbb{R}^N \times [0, T_f] \rightarrow \mathbb{R}^N$, parameterized with a set of N_p random

coefficients $\boldsymbol{\xi}(t, \omega) = [\xi_1(t, \omega), \dots, \xi_{N_p}(t, \omega)]^\top$, satisfies conditions guaranteeing the existence of a unique pathwise solution $\mathbf{x}(t, \omega)$ (see [87]). We assume without loss of generality that the random processes $\boldsymbol{\xi}(t, \omega)$ are zero-mean and characterized by a prescribed single-time joint PDF $f_{\boldsymbol{\xi}}(\boldsymbol{\Xi}; t)$. We frequently use the shorthand $\mathbf{v}(\mathbf{x}(t), t; \omega)$ for the right-hand side of Equation (3.1), use $\mathbb{E}[\cdot]$ and $\langle \cdot \rangle$ interchangeably to denote the ensemble mean, and omit ω in our notation when possible.

Let $\mathbf{X} = [X_1, \dots, X_N]^\top$ denote a phase-space variable in $\mathbb{R}^N$. For any fixed $t > 0$, the system is (partially) characterized by the single-time joint CDF $F_{\mathbf{X}}(\mathbf{X}; t) \triangleq \mathbb{P}[\mathbf{x}(t) \leq \mathbf{X}]$. If $F_{\mathbf{X}}(\mathbf{X}; t)$ is differentiable with respect to all components X_i , the system is equivalently characterized by the single-time joint PDF $f_{\mathbf{X}}(\mathbf{X}; t)$. When the state dimension N is large, deriving or learning a PDF equation for $f_{\mathbf{X}}(\mathbf{X}; t)$ is intractable since the result would be an N -dimensional PDE. We instead consider a low-dimensional QoI $\mathbf{z}(t) \equiv \mathbf{z}(\mathbf{x}(t)) \in \mathbb{R}^{N_{\text{RO}}}$, $N_{\text{RO}} < N$, where $\mathbf{z} : \mathbb{R}^N \rightarrow \mathbb{R}^{N_{\text{RO}}}$ is a continuously differentiable phase-space function guaranteeing the existence of single-time PDF $f_{\mathbf{z}}(\mathbf{Z}; t)$ of $\mathbf{z}(t)$. Here, $\mathbf{Z} \in \mathbb{R}^{N_{\text{RO}}}$ is a phase-space variable for $\mathbf{z}(t)$. We seek a deterministic PDE governing the evolution of $f_{\mathbf{z}}(\mathbf{Z}; t)$, referred to as the RoPDF equation. In practice, $N_{\text{RO}} \leq 3$ to avoid the challenges of high-dimensional PDEs and KDE. We restrict our formulation to the setting of marginal PDF equations, meaning that QoIs take the form $z(t) \triangleq x_k(t)$ for $k \in \{1, \dots, N\}$. The RoPDF equation then reduces to a one-dimensional PDE for the marginal PDF $f_{x_k}(X_k; t)$.

We begin by defining an auxiliary functional or “raw PDF” [199]

$$\Pi_{x_k}(X_k, t) \triangleq \delta(x_k(t) - X_k), \quad (3.2)$$

where $\delta(\cdot)$ is the Dirac delta function. We show in Theorem 3 that Π_{x_k} weakly satisfies the random advection equation (see Equation (3.4)). Moreover, by the Dirac delta’s sifting

property, for any given time $t > 0$, the ensemble mean of Π_k is f_{x_k} :

$$\langle \Pi_{x_k} \rangle (X_k, t) \triangleq \int_{\mathbb{R}} \delta(Y - X_k) f_{x_k}(Y; t) dY = f_{x_k}(X_k; t). \quad (3.3)$$

Given this relationship, an exact, albeit unclosed, RoPDF equation for f_{x_k} is found by stochastically homogenizing the equation for Π_{x_k} . In the setting of joint PDF equations, an analogous procedure for $f_{\mathbf{x}}$ has been the subject of several investigations [115, 143].

To proceed, we require a slight change in notation by denoting $\mathbf{v}(x_k(t), \mathbf{x}_{-k}(t), t; \omega) \equiv \mathbf{v}(\mathbf{x}(t), t; \omega)$, where $\mathbf{x}_{-k}(t) \triangleq [x_1(t), \dots, x_{k-1}(t), x_{k+1}(t), \dots, x_N(t)]^\top$, to emphasize the QoI in the velocity field. A heuristic derivation of the raw PDF equation for Π_{x_k} can be done by weakly differentiating Π_{x_k} with respect to t and employing the sifting property. However, by means of a mollifier argument, we give the formal derivation in the following theorem.

Theorem 3. $\Pi_{x_k}(X_k, t; \omega)$ almost surely obeys, in the sense of distributions, the linear conservation law

$$\frac{\partial \Pi_{x_k}}{\partial t} + \frac{\partial}{\partial X_k} \left[v_k(X_k, \mathbf{x}_{-k}(t), t; \omega) \Pi_{x_k} \right] = 0, \quad \Pi_{x_k}(X_k, 0) = \delta(\mathbf{x}_k^0(\omega) - X_k). \quad (3.4)$$

Proof. Define $\Pi_\epsilon(X_k, t)$, a regularized version of $\Pi_{x_k}(X_k, t)$ in Equation (3.2), as

$$\Pi_\epsilon(X_k, t) \triangleq (\eta_\epsilon \star \Pi_{x_k})(X_k, t) \triangleq \int_{\mathbb{R}} \eta_\epsilon(X_k - Y) \Pi_{x_k}(Y, t) dY = \eta_\epsilon(X_k - x_k(t)), \quad (3.5)$$

where the last equality holds by the definition of $\Pi_{x_k}(Y, t)$ and the sifting property of the Dirac distribution. The standard positive mollifier $\eta_\epsilon \in \mathcal{C}_c^\infty(\mathbb{R})^1$ satisfies the conditions of

1. $\mathcal{C}_c^k(\mathcal{X})$, $0 \leq k \leq \infty$, denotes the space of compactly supported, k -times continuously differentiable real-valued functions on domain $\mathcal{X}$.

symmetry, $\eta_\epsilon(X_k - x_k(t)) = \eta_\epsilon(x_k(t) - X_k)$, and scaling

$$\eta_\epsilon(Y) \triangleq \frac{\epsilon^{-1}}{\int \eta dY} \eta\left(\frac{Y}{\epsilon}\right), \quad \text{where } \eta(Y) \triangleq \begin{cases} \exp\left(\frac{1}{|Y|^2-1}\right) & \text{if } |Y| < 1 \\ 0 & \text{if } |Y| \geq 1. \end{cases} \quad (3.6)$$

Following standard arguments from [62], one can show that Π_ϵ is a smooth approximation of Π_{x_k} . Let $\phi(X_k, t) \in \mathcal{C}_c^1(\mathbb{R} \times [0, \infty))$. It follows from Equation (3.5) that

$$\mathcal{J} \triangleq \int_0^\infty \int_{\mathbb{R}} \Pi_\epsilon(X_k, t) \frac{\partial \phi}{\partial t}(X_k, t) dX_k dt = \int_0^\infty \int_{\mathbb{R}} \eta_\epsilon(X_k - x_k(t)) \frac{\partial \phi}{\partial t}(X_k, t) dX_k dt. \quad (3.7)$$

Integrating by parts in t and applying the sifting property gives

$$\begin{aligned} \mathcal{J} &= \int_0^\infty \int_{\mathbb{R}} \dot{\eta}_\epsilon(X_k - x_k(t)) v_k(x_k(t), \mathbf{x}_{-k}(t), t; \omega) \phi(X_k, t) dX_k dt \\ &\quad - \int_{\mathbb{R}} \eta_\epsilon(X_k - x_k^0) \phi(X_k, 0) dX_k, \\ &= \int_0^\infty \int_{\mathbb{R}} \int_{\mathbb{R}} \dot{\eta}_\epsilon(X_k - Y) v_k(Y, \mathbf{x}_{-k}(t), t; \omega) \Pi_{x_k}(Y, t) \phi(X_k, t) dY dX_k dt \\ &\quad - \int_{\mathbb{R}} \Pi_\epsilon(X_k, 0) \phi(X_k, 0) dX_k, \end{aligned}$$

where $\dot{\eta}_\epsilon(\cdot)$ is the derivative of $\eta_\epsilon(\cdot)$. According to the Gauss-Ostrogradsky theorem in X_k ,

$$\mathcal{J} = - \int_0^\infty \int_{\mathbb{R}} (\eta_\epsilon \star v_k \Pi_{x_k})(X_k, t) \frac{\partial \phi}{\partial X_k}(X_k, t) dX_k dt - \int_{\mathbb{R}} \Pi_\epsilon(X_k, 0) \phi(X_k, 0) dX_k. \quad (3.8)$$

It follows from Equation (3.8) and Equation (3.7) that for any $\phi \in \mathcal{C}_c^1(\mathbb{R} \times [0, \infty))$,

$$\int_0^\infty \int_{\mathbb{R}} \Pi_\epsilon \frac{\partial \phi}{\partial t} dX_k dt + \int_0^\infty \int_{\mathbb{R}} (\eta_\epsilon \star v_k \Pi_{x_k}) \frac{\partial \phi}{\partial X_k} dX_k dt + \int_{\mathbb{R}} \Pi_\epsilon(X_k, 0) \phi(X_k, 0) dX_k = 0.$$

By standard arguments, taking the limit $\epsilon \rightarrow 0$ gives

$$\int_0^\infty \int_{\mathbb{R}} \Pi_{x_k} \frac{\partial \phi}{\partial t} dX_k dt + \int_0^\infty \int_{\mathbb{R}} (v_k \Pi_{x_k}) \frac{\partial \phi}{\partial X_k} dX_k dt + \int_{\mathbb{R}} \Pi_{x_k}(X_k, 0) \phi(X_k, 0) dX_k dt = 0;$$

hence, Π_{x_k} is the distributional solution to Equation (3.4), which completes the proof. $\square$

Taking the ensemble mean of Equation (3.4) and applying the sifting property gives

$$\frac{\partial f_{x_k}}{\partial t} + \frac{\partial}{\partial X_k} \int_{\mathbb{R}^{N-1}} \int_{\mathbb{R}^{N_p}} v_k(X_k, \mathbf{X}_{-k}, t; \Xi) f_{\mathbf{x}, \xi}(\mathbf{X}, \Xi; t) d\Xi d\mathbf{X}_{-k} = 0, \quad (3.9)$$

where $f_{\mathbf{x}, \xi}(\mathbf{X}, \Xi; t)$ denotes the joint PDF of system states $\mathbf{x}(t)$ and random coefficients $\xi(t)$. RoPDF equation (see Equation (3.9)) is exact, but unclosed, since it depends on the generally unknown $f_{\mathbf{x}, \xi}$ and not on f_{x_k} alone. However, factoring $f_{\mathbf{x}, \xi}$ into the product of the marginal PDF f_{x_k} and conditional PDF $f_{\mathbf{x}_{-k}, \xi | x_k}$, Equation (3.9) can be expressed in terms of the regression function $\mathcal{R}$:

$$\frac{\partial f_{x_k}}{\partial t} + \frac{\partial}{\partial X_k} (\mathcal{R}(X_k, t) f_{x_k}) = 0, \quad f_{x_k}(X_k; 0) = \int_{\mathbb{R}^{N-1}} f_{\mathbf{x}^0}(\mathbf{X}) d\mathbf{X}_{-k}, \quad (3.10)$$

together with vanishing boundary conditions, where

$$\mathcal{R}(X_k, t) \triangleq \langle v_k(X_k, \mathbf{x}_{-k}(t), t; \omega) \mid x_k(t) = X_k \rangle. \quad (3.11)$$

is to be estimated from data.

In its current form, Equation (3.10) is fully data-driven, completely relying on accurate estimation of Equation (3.11). However, many applications produce a regression function that is partially, if not fully, separable in the QoI $x_k(t)$. By this, we mean the k -th velocity component has the form $v_k(\mathbf{x}, t; \omega) = \sum_{i \in I} g_i(x_k, t) h_i(\mathbf{x}, t; \omega)$, where I is finite and the real-valued functions $\{g_i, h_i\}_{i \in I}$ are given and smooth. Then, each $g_i(X_k, t)$ may be pulled

outside the conditional expectation Equation (3.11) and need not be estimated, giving the following physics-informed representation of Equation (3.10):

$$\frac{\partial f_{x_k}}{\partial t} + \frac{\partial}{\partial X_k} \left[\left(\sum_{i \in I} g_i(X_k, t) \mathcal{R}_i(X_k, t) \right) f_{x_k} \right] = 0, \quad (3.12)$$

with new regression functions $\mathcal{R}_i(X_k, t) \triangleq \langle h_i(X_k, \mathbf{x}_{-k}(t), t; \omega) \mid x_k(t) = X_k \rangle$.

Note, if h_i has no dependence on x_k , i.e., $h_i \equiv h_i(\mathbf{x}_{-k}, t; \omega)$, for all $i \in I$, then the regression function Equation (3.11) is considered fully separable with respect to the QoI (e.g., see [31, Eq. 23] for a more restricted definition). Well-known examples of fully separable systems include the Kraichnan-Orszag, the Lorenz-63, and the Lorenz-96 systems. For both fully and partially separable systems, part of the advection coefficient is known in closed-form, reducing the amount of data needed for accurate RoPDF solutions, as was investigated for non-stiff, noiseless RODEs in [31].

As is typical in UQ, uncertainty in Equation (3.1) has been fully prescribed. Hence, corresponding state data to be used for regression is synthetically (and independently) generated by numerically integrating Equation (3.1). However, since we are concerned with stiff, multiscale RODEs, costly implicit schemes are required for this data generation, leading to limited data availability. In other words, regression is performed in the small-sample regime, which calls for more expensive, robust algorithms. This issue is exacerbated for systems exhibiting strong nonlinearities, where nonparametric methods must be employed as in Section 3.5.2. Moreover, regression functions associated with multiscale RODEs may vary considerably on short timescales and induce a large RoPDF equation Courant number, requiring regression estimates at an unusually large number of discrete times. However, we drastically reduce our computational overhead by considering only simple, non-robust regression at sparse observation times. Naturally, this simplification amounts to misspecifying the governing RoPDF equation and introduces nonnegligible errors, which we control by sparsely

assimilating in low-fidelity RoPDF estimates. The result is a method whose computational demand is almost entirely associated with the overhead of synthetic state-data generation via (relatively few) MC realizations of Equation (3.1), while preserving the qualitative behavior of the original equation.

3.3 Data Assimilation

Data assimilation procedures for hyperbolic PDEs are categorized into variational and sequential approaches. The former, popularized by 4D-Var [124], is carried out via a global optimization problem that balances observation discrepancy and an *a priori* confidence in the governing model. The optimization is performed through gradient descent and adjoint model integration, which is computationally demanding. On the other hand, sequential approaches directly introduce a correction term to the PDE based on observation discrepancy. Its solution is referred to as an observer, which is designed to converge in time to the true solution as observations are dynamically incorporated. The most notable of these approaches is the ensemble Kalman filter (EnKF), which happens to be equivalent to the variational approach for linear systems. However, when PDE observations are noisy and the discrepancies large in magnitude or highly non-Gaussian, the EnKF performs poorly [105, 125]. Moreover, it suffers from the curse of dimensionality, usually making it infeasible for PDEs upon spatial discretization. However, one workaround, and the first assimilation procedure under consideration is nudging (NR), where the correction term is designed to converge quickly to zero in one forward simulation. Moreover, the nudging appellation motivates our second approach, where we make use of DNNs. Although it is not dynamic assimilation, DNNs can be viewed as instantaneous relaxation, which can address some of NR’s shortcomings.

To reduce error introduced by sparse observation times associated with stiff, multiscale systems, we frame the RoPDF method as a data assimilation problem. Arguably, the two most commonly employed assimilation procedures for hyperbolic PDEs are 4D-Var [124]

and the ensemble Kalman filter (EnKF). However, neither are particularly well-suited for the RoPDF method. The former relies on a computationally demanding global optimization procedure and the latter suffers from the curse of dimensionality, making it ill-suited for discretized PDEs. Moreover, the EnKF performs poorly when PDE observations are noisy with large and/or highly non-Gaussian errors [105, 125]. However, one workaround, and the first assimilation procedure under consideration is nudging (NR), where the PDE correction term is designed to converge quickly to zero in one forward simulation. Moreover, the nudging appellation motivates our second, global approach, where we make use of DNNs. Although it is not dynamic assimilation, DNNs can be viewed as instantaneous relaxation, which can address some of NR's shortcomings, such as (temporal) sparsity of available data.

To formulate the assimilation problems, suppose we have generated $N_{\text{MC}}^{\text{tr}}$ MC realizations of Equation (3.1) (i.e., training data) so that $\mathcal{R}$ may be approximated by a smooth estimator $\hat{\mathcal{R}}$. Letting $\mathcal{E}(X_k, t) \triangleq \mathcal{R}(X_k, t) - \hat{\mathcal{R}}(X_k, t)$ denote the corresponding residual arising from the regression, the RoPDF equation Equation (3.10) can be identically written

$$\frac{\partial f_{x_k}}{\partial t} + \frac{\partial}{\partial X_k}(\hat{\mathcal{R}} f_{x_k}) = \langle \mathcal{M} \rangle, \quad (3.13)$$

where $\langle \mathcal{M} \rangle \equiv -\partial_{X_k}(\mathcal{E} f_{x_k})$, referred to as the model defect/discrepancy, is unknown *a priori*. Note that we have used the fully data-driven RoPDF representation Equation (3.10) simply for notational brevity. In practice, the advection coefficient takes the form of the physics-informed version Equation (3.12), albeit with $\hat{\mathcal{R}}_i$ in place of $\mathcal{R}_i$. By means of NR and DNNs, we learn the model defect by assimilating in RoPDF observations $H(f_{x_k})$, where $H(\cdot)$ represents a given RoPDF observation map.

3.3.1 Nudging

To reduce RoPDF discrepancy, NR assumes the model defect can be described by a simple correction. The resulting PDE for the observer/estimator $\hat{f}_{x_k}^{\text{NR}}$ takes the form

$$\frac{\partial \hat{f}_{x_k}^{\text{NR}}}{\partial t} + \frac{\partial}{\partial X_k}(\hat{\mathcal{R}} \hat{f}_{x_k}^{\text{NR}}) = \lambda(H(f_{x_k}) - \hat{f}_{x_k}^{\text{NR}}), \quad \hat{f}_{x_k}^{\text{NR}}(X_k; 0) = f_{x_k}(X_k; 0), \quad (3.14)$$

with boundary conditions identical to those in Equation (3.10). Here, the observation map $H(\cdot)$ accounts for data availability and sparsity, i.e., when observations of f_{x_k} are possibly noisy and known only on a subset of spatiotemporal locations of the domain. Taking $H(\cdot)$ to be the identity map implies that complete, exact observations are available. The NR coefficient $\lambda > 0$ acts as a finite learning rate that dynamically relaxes the observer towards the observations, controlling the convergence of $\hat{f}_{x_k}^{\text{NR}}$ to f_{x_k} . The choice of λ is largely empirical and typically requires some level of manual tuning. This is in contrast to the EnKF, which takes λ to be the Kalman gain matrix, requiring Gaussian error distributions. In practice, the observations of f_{x_k} are typically noisy to some degree. If they are indeed assumed to be perfectly random, it can be shown that the correction in Equation (3.14) is equivalent to scaled white noise in a stochastic PDE, as discussed in [27] and the references therein. This equivalence was originally given by Jarvisnky [105] between nudged ODEs and SDEs.

At a high level, this explains why the Kalman gain matrix is optimal when error distributions are Gaussian, assuming the underlying dynamics are linear. However, practical NR ignores this introduced uncertainty to a certain level, allowing λ to be manually tuned to fit the data or used for forecasting when the criteria for EnKF are not met. Moreover, when observations are sparse, λ can be constructed to vary in space and/or time, classically comprised of weight functions. Another option is to interpolate observations to the full computational domain. General strategies for constructing λ are reviewed in [187].

3.3.2 Deep Neural Networks

A potential drawback of NR is that qualitative properties of the true RoPDF cannot be guaranteed for the observer, particularly when the observations are noisy and/or sparse. Although these issues are not present for the applications in Section 3.5, this generally may not be the case. For Equation (3.14), its solution $\hat{f}_{x_k}^{\text{NR}}$ is not guaranteed to have the PDF properties of nonnegativity and unit mass.

One alternative is the use of DNNs for direct RoPDF inference from observations. To reformulate the NR problem Equation (3.14) as an instantaneous one via a DNN, we introduce an optimization problem over the (sparse) spatiotemporal observation points of the domain (X_k, t) . We denote the vector of N_{obs} -many discrete observation locations by $\tilde{\mathbf{X}}_k^\nu = [\mathbf{X}_k^\top, \mathbf{T}_\nu^\top]^\top$, where $\mathbf{X}_k$ and $\mathbf{T}_\nu$ represent the spatial and temporal components, respectively. The subscript $\nu \in \mathbb{N}$ denotes the level of temporal sparsity of the data, which is formally defined in Section 3.4. The optimization problem is then defined by minimizing the discrepancy between the observer $\hat{f}_{x_k}^{\text{DNN}}$ and observations via the following loss function:

$$\mathcal{L} \triangleq \left\| \hat{f}_{x_k}^{\text{DNN}}(\mathbf{X}_k; \mathbf{T}_\nu) - H(f_{x_k}(\mathbf{X}_k; \mathbf{T}_\nu)) \right\|, \quad (3.15)$$

where $\|\cdot\|$ is an appropriate norm over $\mathbb{R}^{N_{\text{obs}}}$. Directly approximating $\hat{f}_{x_k}^{\text{DNN}}$ in Equation (3.15) as a DNN is a possibility, but given that the result would be solely data-driven, utilizing the partially known dynamics (i.e., the separable advection term) of Equation (3.12) and Equation (3.13) gives better results due to increased statistical power. To incorporate such dynamics and render the loss Equation (3.15) physics-informed, we utilize the RoPDF equation's linearity.

The solution f_{x_k} to Equation (3.13) can be decomposed into the sum of its homogeneous and particular (defect) solutions $f_{x_k}^{\text{h}}$ and $f_{x_k}^{\text{d}}$, respectively, such that $f_{x_k} = f_{x_k}^{\text{h}} + f_{x_k}^{\text{d}}$. Naturally, $f_{x_k}^{\text{h}}$ is the solution to the homogeneous equation Equation (3.13) (i.e., when

$\langle \mathcal{M} \rangle \equiv 0$), while $f_{x_k}^d$ accounts for the defect's contribution. This fact can be established by applying the method of characteristics to Equation (3.13) via the terminal value problem

$$\frac{d\chi(s)}{ds} = \hat{\mathcal{R}}(\chi(s), s), \quad \chi(t) = X_k, \quad (3.16)$$

and its associated flow $\chi(s) \equiv \Phi(s; X_k, t)$ for $0 \leq s < t$. By restricting f_{x_k} along the characteristic curves, the RoPDF equation can be solved via integrating factor, resulting in

$$\begin{aligned} f_{x_k}^h(X_k; t) &= \mathcal{J}(0; X_k, t) f_{x_k}(\Phi(0; X_k, t); 0), \\ f_{x_k}^d(X_k; t) &= \int_0^t \langle \mathcal{M} \rangle(\chi(r), r) \mathcal{J}^{-1}(r; X_k, t) dr, \end{aligned} \quad (3.17)$$

where $\mathcal{J}(s; X_k, t) \triangleq \exp\left(-\int_s^t \partial_\chi \hat{\mathcal{R}}(\chi(r), r) dr\right)$.

The homogeneous solution $f_{x_k}^h$ in Equation (3.17) is directly computed via numerical integration. This is done by separating the advection coefficient into its known and unknown terms as in Equation (3.12), approximating $\mathcal{R}_i$ with smooth estimators $\hat{\mathcal{R}}_i$ on the spatial mesh $\mathbf{X}_k$ for each (sparse) observation time in $\mathbf{T}_\nu$. When $\nu > 1$, the learned $\hat{\mathcal{R}}_i$ may be interpolated to the dense temporal grid required by the homogeneous PDE discretization to improve performance. Having $f_{x_k}^h$ at our disposal, we construct an instantaneous observer $\hat{f}_{x_k}^{\text{DNN}} \triangleq f_{x_k}^h + \hat{f}_{x_k}^d$ for f_{x_k} by estimating $f_{x_k}^d$ with a fully connected feedforward DNN $\hat{f}_{x_k}^d$ containing N_{lay} layers:

$$\hat{f}_{x_k}^d(X_k; t) \triangleq \mathbf{A}_{N_{\text{lay}}} \circ \phi \circ \mathbf{A}_{N_{\text{lay}}-1} \circ \cdots \circ \phi \circ \mathbf{A}_1 \tilde{\mathbf{X}}_k^\nu, \quad (3.18)$$

where ϕ is a nonlinear activation function applied recursively to each of the $N_{\text{lay}} - 1$ hidden layers. Note that since ϕ is typically bounded from above and/or below, it is not applied to the output layer $\mathbf{A}_{N_{\text{lay}}}$ since our intended purpose is regression. Since $f_{x_k}^d$ accounts for the defect's contributions to the RoPDF f_{x_k} , its qualitative behavior can be complex, i.e.,

nonperiodic with steep gradients.

DNNs are an expressive hypothesis class, and are known to learn complex function behavior, which is the reasoning behind this choice of observer. Moreover, partial separability of the advection coefficient ensures that $f_{x_k}^h$ and therefore $\hat{f}_{x_k}^{\text{DNN}}$ is physics-informed, resulting in increased predictive power.

Since there are no existing theoretical results for the convergence of the DNN observer, the choice of norm is not as restrictive as in NR. We take the standard mean squared error (MSE) since it gives a smooth, convex loss, significantly reducing computational costs compared to the nondifferentiable L_1 loss, but it is not without caveats. Employing the MSE loss for DNN regression may result in poor training convergence if the underlying error distributions in the observations $H(f_{x_k})$ are not close to being independent, identical, and Gaussian. The MSE can be replaced with a more general loss function to address errors that strongly violate these properties. For example, to account for non-constant variance, one can use the generalized least squares (GLS) loss as in [121, Eq. 6], where DNNs (specifically physics-informed neural networks (PINNs) [176]) were trained with the GLS loss to improve training convergence. Moreover, if the resulting observer does not have the desired properties of a PDF, regularity terms may be added to the loss to enforce nonnegativity, unit mass, and boundary conditions. Since the learning targets are PDFs, another option is to formulate the loss in Equation (3.15) with a distribution metric. Natural choices include the total variation distance, Wasserstein distances, and f -divergences such as the commonly used Kullback–Leibler (KL) divergence.

For the applications in Section 3.5, implementing regularity terms, a GLS loss, and a KL divergence loss did not provide major improvements compared to the MSE, but all significantly increased training costs. Therefore, the experiments presented use the standard MSE loss but with problem-specific transformations applied (before training) to both the predictor (observation location) and response ($H(f_{x_k})$) data to account for vastly differing

scales and a variety of complex error distributions, making the MSE loss just as effective. Although the MSE is used for DNN training/validation, (testing) error comparisons are always given over L_1 , which is pessimistic in testing PDF accuracy compared to L_2 and distribution metrics, since the existing NR theory is established in this space.

3.4 Implementation and Analysis

Via the order 1.5 implicit strong Taylor scheme from [87, Ch. 10.2], state data is synthetically generated by MC simulations of Equation (3.1) and collected on a set of uniform times $\mathbf{T}_1 = \{t_m\}_{m=0}^M \triangleq \{m\Delta t\}_{m=0}^M$, where $t_M = T_f$. For each $t_m \in \mathbf{T}_1$, a large number of N_{MC} MC samples of the QoI $x_k(t_m)$ are post-processed with robust, adaptive-like KDE [26] to form a MC marginal PDF solution $f_{\text{MC}}(X_k; t_m)$, which is treated as a yardstick for testing the RoPDF method. Naturally, N_{MC} is problem dependent and is determined by means of a convergence study for each experiment.

We introduce the notion of sparse data via a sparsity factor $\nu \in \{1, \dots, M\}$ and its associated observation times $\mathbf{T}_\nu = \{t_{m_l}\}_{l=0}^{M_\nu} \triangleq \{\nu l \Delta t\}_{l=0}^{M_\nu}$, where $M_\nu \leq M$. Hence, $\nu = 1$ implies complete observations, $\nu = 2$ implies every other observation is available, and so on. Independent of the trials for f_{MC} , we perform $N_{\text{MC}}^{\text{tr}} \ll N_{\text{MC}}$ MC simulations of Equation (3.1) to collect state (i.e. $\mathbf{x}$) and, if required, noise ($\boldsymbol{\xi}$) data for training. For a given ν , at each time $t_{m_l} \in \mathbf{T}_\nu$, these samples are used to learn $\hat{\mathcal{R}}_i(X_k, t_{m_l})$ via linear (ordinary least squares (OLS)) regression and Gaussian local linear regression (GLLR) [91] for the linear and power systems applications in Section 3.5, respectively. Additionally, the $x_k(t_{m_l})$ samples are post-processed with KDE [26] to compute a low-fidelity RoPDF observation $H(f_{x_k}(X_k; t_{m_l}))$ for each observation time. Since $N_{\text{MC}}^{\text{tr}} \ll N_{\text{MC}}$, these RoPDF observations are inherently noisy, which is exacerbated in the temporal domain for $\nu > 1$. We henceforth denote these observations by $f_{\text{MC}}^{\text{tr}, \nu}$, to identify their dependence on the training sample size, KDE, and the sparsity level. Much of our analysis will focus on how training size and sparsity

level affect observer accuracy.

In all experiments that follow, the homogeneous RoPDF equations are solved on the set of dense times $\mathbf{T}_1$ via a Lax-Wendroff finite volume discretization with a monotized central limiter. When observation times are sparse, this is done by interpolating the learned regression functions $\hat{\mathcal{R}}_i$ to the dense spatiotemporal grid $(\mathbf{X}_k, \mathbf{T}_1)$ via 2D modified Akima interpolation [3]. Although the phase space is unbounded in our formulation, the computational spatial domain is taken to be a sufficiently large (bounded) interval so that vanishing boundary conditions at $\pm\infty$ may be approximated with homogeneous Dirichlet conditions.

3.4.1 Assimilation Training

The nudged equation in Equation (3.14) is solved by successively considering the homogeneous advection and source equations via Strang operator splitting, where the source equation is integrated with a Crank-Nicolson discretization. Similar to $\hat{\mathcal{R}}_i$, for $\nu > 1$, observations $f_{\text{MC}}^{\text{tr},\nu}$ are also interpolated to the dense mesh before numerically integrating, which avoids the laborious tuning of NR weight functions.

We take the relaxation rate $\lambda \equiv \lambda_\nu(t)$ to be piecewise constant over observation intervals $[t_{m_l}, t_{m_{l+1}})$, which is tuned in an online fashion to reduce predictive error at the subsequent observation times. In particular, for a given interval with $t \in [t_{m_l}, t_{m_{l+1}})$ and $\nu > 1$, we consider two possible values: $\lambda_\nu(t) \equiv 0$ and $\lambda_\nu(t) \equiv \nu$. Supposing λ_ν and $\hat{f}_{x_k}^{\text{NR}}$ have been computed for $t < t_{m_l}$, we solve Equation (3.14) up to the following observation time $t_{m_{l+1}}$ for both possible values of λ_ν . Whichever produces the lowest (L_1) prediction error between the observer $\hat{f}_{x_k}^{\text{NR}}$ and the observation $f_{\text{MC}}^{\text{tr},\nu}$ at time $t_{m_{l+1}}$ is taken as $\lambda_\nu(t)$ on $t \in [t_{m_l}, t_{m_{l+1}})$. This approach to tuning ensures that observations are assimilated into RoPDF dynamics only when necessary.

For both applications that follow, in the DNN formulation, we represent the defect solution $\hat{f}_{x_k}^{\text{d}}$ as a fully connected DNN with a ReLU activation function. For each combination

of ν and $N_{\text{MC}}^{\text{tr}}$, we train a DNN via the standard MSE loss in Equation (3.15) using a 30% holdout set for model validation. The MSE is minimized via the L-BFGS optimizer in PyTorch v1.13.0 with a maximum of 5×10^3 iterations. To prevent overfitting, we implement an early-stopping criterion by imposing a 10^{-8} gradient tolerance, which allowed training to terminate in at most 10^3 iterations. We consider 3 to 10 equally sized hidden layers, ultimately choosing the network depth that minimizes validation MSE. For Section 3.5.1, the network width is fixed at 20 neurons, which is subsequently increased to 32 neurons for Section 3.5.2. In both applications, $N_{\text{MC}}^{\text{tr}}$ has little effect on optimal network depth so long as it is not overwhelmingly small relative to dynamic complexity, e.g., greater than 250 and 10^3 for the linear and powers systems, respectively. $\nu \geq 1$, on the other hand, is much more influential on optimal depth. This is not surprising considering that sparse $\hat{\mathcal{R}}_i$ may introduce large RoPDF errors for systems that are multiscale and/or rapidly oscillating, resulting in defects of greater complexity. After the training period, for each ν , an $\hat{f}_{x_k}^{\text{d}}$ prediction is computed on the set of complete times $\mathbf{T}_1$ required by the homogeneous equation's discretization. It is added to the homogeneous solution $f_{x_k}^{\text{h}}$ to obtain the instantaneous observer $\hat{f}_{x_k}^{\text{DNN}}$.

3.4.2 Computational Complexity

Unlike DNNs, it is straightforward to compute the homogeneous and NR equations $\mathcal{O}(\cdot)$ complexities. For simplicity, suppose the 1D spatial domain is discretized with a fixed uniform mesh $\mathbf{X}_k$ containing N_{x_k} cells. Likewise, the dense temporal grid $\mathbf{T}_1$ contains $M + 1$ nodes. Assume $N_{\text{MC}}^{\text{tr}}$ MC realizations of Equation (3.1) have been computed and stored at the $M_\nu \approx \lceil M/\nu \rceil$ -many times $\mathbf{T}_\nu$ ($\nu \geq 1$).

The homogeneous equation coincides with NR in Equation (3.14) when $\lambda \equiv 0$ for all $t \in [0, T_f]$. Given an advection coefficient, a Lax-Wendroff time step requires $\mathcal{O}(N_{x_k})$ operations, and therefore a total of $\mathcal{O}(MN_{x_k})$ operations over $[0, T_f]$. Given the Courant-Friedrichs-Lewy (CFL) condition to ensure numerical stability, this can be expressed as $\mathcal{O}(N_{x_k}^2)$. However, we

must account for the cost of learning $\hat{\mathcal{R}}$. Consider the more expensive nonparametric GLLR regression. For a given $t_{m_l} \in \mathbf{T}_\nu$ and bandwidth parameter, GLLR fitting and evaluation has $\mathcal{O}(N_{\text{MC}}^{\text{tr}} N_{x_k})$ complexity [91, Ch. 6.9]. A typical cross-validation (CV) procedure for bandwidth selection increases this cost from linear to quadratic in $N_{\text{MC}}^{\text{tr}}$. However, we avoid CV by transforming the data so that the simple plug-in estimator (see Equation (3.28)) is sufficiently accurate, keeping GLLR $\mathcal{O}(N_{\text{MC}}^{\text{tr}} N_{x_k})$, and therefore $\mathcal{O}(M_\nu N_{\text{MC}}^{\text{tr}} N_{x_k})$ over all observation times.

The cost of employing any piecewise cubic Hermite interpolating polynomial is at most $\mathcal{O}(M N_{x_k})$, giving a total complexity of

$$\mathcal{O}\left(M_\nu N_{\text{MC}}^{\text{tr}} N_{x_k} + M N_{x_k} + M N_{x_k}\right) = \mathcal{O}\left(\left(N_{\text{MC}}^{\text{tr}}/\nu + 2\right) N_{x_k}^2\right) \quad (3.19)$$

for the homogeneous RoPDF equation.

In addition to Equation (3.19), to solve the nudged equation in Equation (3.14), we must account for Strang splitting and the KDE/interpolation procedure for $f_{\text{MC}}^{\text{tr},\nu}$. However, the cost of computing $f_{\text{MC}}^{\text{tr},\nu}$ at times $\mathbf{T}_\nu$ via KDE and interpolating to the dense mesh is identical to the advection coefficient procedure. Moreover, a single time step of the source equation takes $\mathcal{O}(N_{x_k})$ operations, and taking M to be even only requires an additional M time steps for Strang splitting (compared to the standard $2M$ additional steps). Thus, the computational complexity of solving the nudged equation in Equation (3.14) is twice that of the homogeneous equation.

Circling back to the overhead of generating $N_{\text{MC}}^{\text{tr}}$ realizations of the N -dimensional RODE Equation (3.1), for a given path realization, the majority of costs corresponds to the nonlinear/implicit solve required at each time step. If we assume that the mean velocity field's Jacobian is known exactly, then each iteration of a Newton-type method can be done in $\mathcal{O}(N^2)$ to $\mathcal{O}(N^3)$ operations, depending on the system and matrix factorization. Recalling $\mathcal{O}(M) = \mathcal{O}(N_{x_k})$, this amounts to an overhead of $\mathcal{O}(N_{\text{MC}}^{\text{tr}} N_{x_k} N^3)$ operations for

an arbitrary stiff, nonlinear system. However, we have not accounted for multiple iterations during nonlinear solves nor Jacobian approximations. Hence, true overhead costs may be considerably larger, especially for very stiff and/or high-dimensional RODEs. Regardless, even for moderate N , sampling (Equation (3.1)) dominates Equation (3.19), and therefore the cost of the nudged RoPDF method.

3.5 Numerical Experiments

We test the proposed RoPDF approaches on two applications, both with random initial data and driven by OU noise. The first in Section 3.5.1 is a stiff, 2D (i.e., $N = 2$) linear system, included as a proof-of-concept. It highlights the need for data assimilation in RoPDF equations associated with sparsely observed stiff RODEs, even when the underlying dynamics are relatively simple. Our second application in Section 3.5.2 is to power systems, where the RoPDF method is used for UQ of transmission/line failures in an electrical power grid. The governing model is a 47D nonlinear system. It is highly stiff across the entire system and lacks scale separation, similar to the (simpler) model in [129]. We remark that, in addition to the QoI, only 4 of the 47 states need to be observed for our approaches.

3.5.1 Stiff Linear System

We first consider the following linear RODE system:

$$\begin{aligned} \dot{x}_1 &= -2x_1 + x_2 + 2\sin(t), & x_1(0) &\sim \mathcal{N}\left(2, 0.15^2\right), \\ \dot{x}_2 &= (\alpha - 1)x_1 - \alpha x_2 + \alpha(\cos(t) - \sin(t)) + \sigma\xi(t), & x_2(0) &\sim \mathcal{N}\left(3, 0.15^2\right), \end{aligned} \quad (3.20)$$

to be solved up to $T_f = 10$, where the Gaussian initial conditions and noise (i.e., $x_1(0)$, $x_2(0)$, $\xi(t)$) are all taken independent of one another. The driving colored noise is taken as




Figure 3.1: (Left) Learned $\hat{\mathcal{R}}$ from $N_{\text{MC}}^{\text{tr}} = 5 \times 10^2$ MC realizations of Equation (3.20) at sparse times $\mathbf{T}_\nu$ with $\nu = 2 \times 10^2$. (Middle) The learned advection coefficient of Equation (3.23). (Right) Evolution of the homogeneous solution $f_{x_1}^h$ to Equation (3.23).

an OU process defined as the solution to the Itô SDE

$$d\xi(t) = -\frac{\xi(t)}{\tau}dt + \sqrt{\frac{2}{\tau}}dW(t), \quad \xi(0) \sim \mathcal{N}(0, 1), \quad (3.21)$$

where $W(t)$ is a standard Wiener process independent of $\xi(0)$. Given this initial condition, $\xi(t)$ is an exponentially correlated stationary Gaussian process with correlation length $\tau > 0$. Its solution is conditionally given by the scaled, time-transformed Wiener process

$$\xi(t) = \xi(0)e^{-t/\tau} + W\left(1 - e^{-2t/\tau}\right). \quad (3.22)$$

Hence, paths of ξ can be directly sampled from the laws of $\xi(0)$ and W , which is more accurate and efficient than numerically integrating Equation (3.21). The intensity of ξ is denoted by $\sigma > 0$. Lastly, $\alpha > 0$ serves as a stiffness parameter, making Equation (3.20) stiff when $\alpha \gg 1$. In the experiments that follow, we set $\alpha = 999$, $\sigma = 100$, $\tau = 0.1$, and $k = 1$ such that the QoI is $x_1(t)$.

RoPDF Equation & Numerics

The RoPDF equation in Equation (3.12) takes the form

$$\frac{\partial f_{x_1}}{\partial t} + \frac{\partial}{\partial X_1} \left[\left(-2X_1 + \mathcal{R}(X_1, t) + 2\sin(t) \right) f_{x_1} \right] = 0, \quad (3.23)$$

where the initial condition $f_{x_1}(X_1; 0)$ is the univariate Gaussian PDF of $x_1(0)$ and $\mathcal{R}(X_1, t) \triangleq \langle x_2(t) | x_1(t) = X_1 \rangle$. The spatial mesh $\mathbf{X}_1$ is taken uniformly on $[-1.85, 3.15]$ with 5×10^2 cells such that $\Delta X_k = 10^{-2}$. This is a fairly dense mesh for the given dynamics; however, it demonstrates one manner in which sparse observations arise even when the dynamics are straightforward. For the given ΔX_k and estimated coefficient, the CFL condition requires $\Delta t \lesssim 1.4 \times 10^{-3}$. Since this estimated bound is dependent upon the specific regression and interpolation methods, we take a slightly smaller time step $\Delta t = 10^{-3}$ for the dense grid $\mathbf{T}_1$. Given the simple structure of Equation (3.20), $\mathcal{R}$ is linear in X_1 . Hence, the estimator $\hat{\mathcal{R}}$, which is displayed in Figure 3.1 along with its associated RoPDF evolution, is computed via OLS regression from $N_{\text{MC}}^{\text{tr}}$ MC realizations of $\mathbf{x}$ at each $t_{m_l} \in \mathbf{T}_\nu$.

To be consistent with the notation in our DNN formulation, we denote the solution to Equation (3.23) (with $\hat{\mathcal{R}}$) by $f_{x_1}^h$, which signifies that no RoPDF observations were assimilated. We incorporate varying degrees of data sparsity by considering $\nu \in \{1, 2, 5\} \times 10^2$, which corresponds to data availability at time increments of 0.1, 0.2, and 0.5. Several training sample sizes were also considered in our experiments, but apart from our scalability results (Figure 3.4), we limit our head-to-head comparisons to the $N_{\text{MC}}^{\text{tr}} = 5 \times 10^2$ case, which is considerably fewer than the $N_{\text{MC}} = 1.5 \times 10^4$ realizations needed to compute yardstick solution f_{MC} .

As mentioned in Section 3.3.2, the loss in Equation (3.15) in the DNN formulation is not actually minimized over the observation locations $\tilde{\mathbf{X}}_1^\nu$. Since the PDFs are near-Gaussian, for a given ν , we compute the mean and standard deviation of $f_{x_1}^h(\mathbf{X}_1; t_{m_l})$ for each $t_{m_l} \in$




Figure 3.2: (Left) Temporal evolution of $f_{x_1}^h$ L_1 error against the yardstick f_{MC} for $\nu \in \{1, 2, 5\} \times 10^2$ and $N_{MC}^{tr} = 5 \times 10^2$. (Middle) Snapshot of f_{MC} , $f_{x_1}^h$, $\hat{f}_{x_1}^{NR}$, and $\hat{f}_{x_1}^{DNN}$ for $\nu = 2 \times 10^2$ at time $t = 7.2$. (Right) Defects $\hat{f}_{x_1}^d$ corresponding to the middle plot for DNN and NR observers. The latter is computed *ex post facto* as $\hat{f}_{x_1}^{NR} - f_{x_1}^h$.

$\mathbf{T}_\nu$. The spatial observation locations are then shifted and scaled by the corresponding mean and standard deviation for each time. Both $f_{x_1}^h$ and $f_{MC}^{tr,\nu}$ are also scaled by these standard deviations. The resulting transformations result in PDFs that are nearly standard Gaussian for all observation times, albeit defined on varying/moving spatial grids. This method of standardization allows us to omit, for all $t_{m_l} \in \mathbf{T}_\nu$, any transformed spatial location with magnitude greater than four, i.e., where (transformed) $f_{x_1}^h$ and $f_{MC}^{tr,\nu}$ are within machine epsilon. This improves DNN costs by reducing training input size and allows training to converge with shallower networks. After training, the $\hat{f}_{x_1}^d$ prediction on dense $\mathbf{T}_1$ is transformed back to original scale on $\mathbf{X}_1$.

Recall from Section 3.4.1 that all DNNs have 20 neurons per hidden layer for this application, but the number of layers is chosen to minimize validation MSE. Across all choices of N_{MC}^{tr} , the optimal network depth ranges from 3 to 4 layers for smaller values of ν and 5 to 6 layers for larger ν . Specifically, for $N_{MC}^{tr} = 5 \times 10^2$, the networks corresponding to each $\nu \in \{1, 2, 5\} \times 10^2$ have 3, 3, and 5 hidden layers, respectively.

Error Analysis

Figure 3.2 (middle) reveals the RoPDF solutions to be overwhelmingly Gaussian. This is no surprise since Equation (3.20) is linear with Gaussian noise and initial conditions. Moreover,

the snapshot of $f_{x_1}^h$ for $\nu = 2 \times 10^2$ at time $t = 7.2$ is a close match to the yardstick f_{MC} even though the L_1 error is approximately 11%, as seen in left subfigure. Upon closer inspection, there are deviations in the mean, variance, and left tail from f_{MC} , though they appear minimal. This behavior with respect to f_{MC} is similar at other times and for other combinations of ν and N_{MC}^{tr} but is more pronounced as ν increases.

Although it is pessimistic for PDFs, we limit our error to L_1 since theoretical NR convergence is established in this metric (see Section 3.7 and [27]). However, regardless of the metric, there is always a sharp increase in $f_{x_k}^h$'s error at early times, where the error magnitude is largely determined by ν . This is expected given the dynamics' initial transience, where the RoPDF quickly transitions away from the initial condition to the dominant periodic evolution seen in Figure 3.1 (right). Naturally, if ν is too large, even with $\hat{\mathcal{R}}$'s interpolation, the advection coefficient cannot properly account for this transience, and $f_{x_k}^h$ is perturbed away from the true dynamics without any means for correction, even if the coefficient is correctly estimated at later times. This is where our proposed assimilation methods pick up the slack.

Figure 3.3 (left) is the temporal L_1 error evolution of the NR observer against the observations used during assimilation, which helps visualize the NR procedure. The tick marks along the horizontal axis denote the relatively few time periods when $\lambda > 0$ and RoPDF observations are assimilated into the dynamics via Equation (3.14). They typically correspond to small magnitudes and sharp decreases in error, showing that observations are assimilated in quickly when it serves to increase predictive power. This figure also shows the error associated with the NR (middle) and DNN (right) observers against f_{MC} , revealing that both approaches perform well compared to the homogeneous solution (Figure 3.2, left). The most striking result is that both observers, save for the initial transience, are relatively unaffected by temporal sparsity as long as ν is not unreasonably large. This fact can also be seen in our convergence rates in Figure 3.4. Overall, the DNN slightly outperforms NR, which we attribute to relatively simple error distributions and defects. The latter can be seen in




Figure 3.3: Evolution of L_1 error for $\nu \in \{1, 2, 5\} \times 10^2$ and $N_{\text{MC}}^{\text{tr}} = 5 \times 10^2$. (Left) $\hat{f}_{x_1}^{\text{NR}}$ against the assimilated $f_{\text{MC}}^{\text{tr},\nu}$. Green ticks on the t -axis represent short assimilation periods, i.e., when $\lambda_\nu(t) > 0$. (Middle) $\hat{f}_{x_1}^{\text{NR}}$ against the yardstick f_{MC} . (Right) $\hat{f}_{x_1}^{\text{DNN}}$ against f_{MC} .

Figure 3.2 (right).

Figure 3.4 provides convergence (in the normalized L_1 norm over space and time) of the assimilated observations (left), the NR observer (middle), and the DNN observer (right) as $N_{\text{MC}}^{\text{tr}}$ increases. For complete data ($\nu = 1$), we recover the standard MC convergence rate of $\mathcal{O}(1/\sqrt{N_{\text{MC}}^{\text{tr}}})$ for the observations. However, as data becomes sparse, this convergence considerably degrades due to the observations' construction via interpolation. For $\nu = 1 \times 10^2$, the error of $f_{\text{MC}}^{\text{tr},\nu}$ increases in magnitude and the convergence slows to $\mathcal{O}(1/\sqrt[3]{N_{\text{MC}}^{\text{tr}}})$. For $\nu > 2 \times 10^2$, the error magnitude continues to increase and the rate is nearly constant. The NR observer, on the other hand, surpasses standard and quasi-MC rates with $\mathcal{O}(1/N_{\text{MC}}^{\text{tr}})$ convergence. The remarkable feat is that this rate is nearly independent of ν . This also holds for the DNN observer, but with slightly smaller magnitudes and sharper rates.

Overall, both approaches to assimilation are effective and cut costs of the MC approach by a factor of 4. This speedup is significant but not drastic given that the MC approach is on the scale of minutes in CPU time, which is due to linear dynamics and low dimensionality. We note that the experiments were performed with an Apple M2 Max chip in parallel on 12 CPU cores.




Figure 3.4: Convergence rates, on a log-log scale, for the spatiotemporal L_1 error of the observations $f_{MC}^{tr,\nu}$ (left), $\hat{f}_{x_1}^{NR}$ (middle), and $\hat{f}_{x_1}^{DNN}$ (right) for $\nu = \{1, 2, 5\} \times 10^2$ as N_{MC}^{tr} increases.

3.5.2 Power System Cascade Outages

We study how our method applies to electrical power systems, particularly in characterizing cascading failure modes dependent on stochastic sources. As the grid sees an increase in renewable generation sources and electric vehicles, the risk of stochastic fluctuations triggering a cascade of failures rises, potentially leading to significant economic impacts and safety risks.

The power system consists of $N_{bus} = N_g + N_l$ buses, comprised of N_g -many generators and N_l -many loads, and a network of N_{line} -many transmission lines. Sudden perturbations of the steady state can lead to the overloading of transmission lines, resulting in their sudden trip or disconnection. This disconnection may trigger further line disconnections in a cascade fashion. Therefore, to characterize the risk of cascades, the RoPDF method is employed, which can be used to compute the probability of line outages in response to stochastic perturbations.

To address this issue in a computationally tractable fashion, Zheng and DeMarco proposed a port-Hamiltonian model to represent the potential of such cascading outages that incorporates line tripping by means of “smooth bistable” variables [52]. We show the com-

plete model:

$$\begin{aligned}
\dot{\boldsymbol{\omega}}_g &= -\mathbf{M}_g^{-1} \mathbf{D}_g \boldsymbol{\omega}_g - \mathbf{M}_g^{-1} \mathbf{U}_1^\top \mathbf{f}(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}), \\
\dot{\boldsymbol{\alpha}} &= \mathbf{U}_1 \boldsymbol{\omega}_g - \left[\mathbf{U}_2 \mathbf{D}_l^{-1} \mathbf{U}_2^\top \right] \mathbf{f}(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}), \\
\dot{\mathbf{V}}_l &= -\mathbf{D}_v^{-1} \mathbf{g}(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}), \\
\dot{\boldsymbol{\gamma}} &= -\mathbf{D}_\gamma^{-1} \mathbf{h}(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}),
\end{aligned} \tag{3.24}$$

where

$$\mathbf{U} \triangleq [-\mathbf{e} \mid \mathbf{I}_{N_{\text{bus}}}] = [\mathbf{U}_1 \mid \mathbf{U}_2], \quad \mathbf{U}_1 \in \mathbb{R}^{N_{\text{bus}} \times (N_g + 1)}, \quad \mathbf{U}_2 \in \mathbb{R}^{N_{\text{bus}} \times N_l}, \quad \mathbf{e} \triangleq [1, \dots, 1]^\top \in \mathbb{R}^{N_{\text{bus}}},$$

and $\mathbf{I}_{N_{\text{bus}}}$ is the $N_{\text{bus}} \times N_{\text{bus}}$ identity matrix. The system states $\mathbf{x}(t) \triangleq [\boldsymbol{\omega}_g^\top, \boldsymbol{\alpha}^\top, \mathbf{V}_l^\top, \boldsymbol{\gamma}^\top]^\top$ are comprised of $(N_g + 1)$ generator speeds $\boldsymbol{\omega}_g$ (including a non-physical reference/slack bus), N_{bus} non-slack angles $\boldsymbol{\alpha}$, N_l load voltage magnitudes $\mathbf{V}_l$, and N_{line} indicator-like bistable variables $\boldsymbol{\gamma}$ representing the operating status of each line. Hence, Equation (3.24) is an $(N_g + 1 + N_{\text{bus}} + N_l + N_{\text{line}})$ -dimensional system. Here, $\mathbf{M}_g \in \mathbb{R}^{(N_g + 1) \times (N_g + 1)}$ is the generator mass/inertia matrix and $\mathbf{D}_g \in \mathbb{R}^{(N_g + 1) \times (N_g + 1)}$, $\mathbf{D}_l \in \mathbb{R}^{N_l \times N_l}$, $\mathbf{D}_v \in \mathbb{R}^{N_l \times N_l}$, and $\mathbf{D}_\gamma \in \mathbb{R}^{N_{\text{line}} \times N_{\text{line}}}$ are the states' various damping matrices. $\mathbf{f}(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}) \in \mathbb{R}^{N_{\text{bus}}}$ represents the net power at each of the non-slack buses, where the sign convention takes absorbed power as positive. In other words,

$$f_i(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}) \triangleq \tilde{f}_i(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}) - P_i^0, \quad i \in \{1, \dots, N_{\text{bus}}\},$$

where $\mathbf{P}^0 = [P_1^0, \dots, P_{N_{\text{bus}}}^0]^\top \in \mathbb{R}^{N_{\text{bus}}}$ represents the prescribed active mechanical power from the generators and the negative active load power demands. $\mathbf{Q}^0 \in \mathbb{R}^{N_l}$ is the reactive

power analog of $\mathbf{P}^0$, but defined only at the loads, and $\mathbf{g}(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}) \in \mathbb{R}^{N_l}$ is defined as

$$g_i(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}) \triangleq V_{l,i}^{-1} \left(\tilde{g}_i(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}) - Q_i^0 \right), \quad i \in \{N_g + 1, N_g + 2, \dots, N_{\text{bus}}\}.$$

The system (see Equation (3.24)) approximates tripping dynamics via the bistable states $\boldsymbol{\gamma}$ and their corresponding velocity field

$$\mathbf{h}(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}) \triangleq \tilde{\mathbf{h}}(\boldsymbol{\alpha}, \mathbf{V}_l) - \mathbf{H} \odot \boldsymbol{\theta}(\boldsymbol{\gamma}), \quad (3.25)$$

where $\tilde{\mathbf{h}}(\boldsymbol{\alpha}, \mathbf{V}_l) \in \mathbb{R}^{N_{\text{line}}}$ denotes the lines' energies and $\boldsymbol{\theta}(\boldsymbol{\gamma}) \in \mathbb{R}^{N_{\text{line}}}$ comprises the corresponding thresholding functions with magnitudes $\mathbf{H} \in \mathbb{R}^{N_{\text{line}}}$. Here, $\odot$ is the element-wise vector product. The smooth thresholding functions

$$\theta_k(\boldsymbol{\gamma}) \triangleq 2 \left[-\exp(-20\gamma_k) + \exp(-200\gamma_k) + \exp(20(\gamma_k - 1)) - \exp(200(\gamma_k - 1)) \right], \quad (3.26)$$

for $k \in \{1, \dots, N_{\text{line}}\}$, are constructed so that upon integrating Equation (3.24), two potential wells are created very close to zero and one, where the latter has height $\approx H_k$. When the line energy $\tilde{h}_k$ exceeds the threshold H_k , $h_k(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma})$ becomes very large, driving γ_k in Equation (3.24) quickly to zero, effectively removing the line from the system. Moreover, due to Equation (3.26), once γ_k transitions to zero, it stays there, save for small fluctuations around zero.

To account for stochastic fluctuations at the loads, we add $N_p = 2N_l$ OU noise processes $\boldsymbol{\xi}(t)$ to $\mathbf{P}^0$ and $\mathbf{Q}^0$ such that

$$\begin{aligned} f_i(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}, \boldsymbol{\xi}) &\triangleq \tilde{f}_i(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}) - P_i^0 - \sigma_{P,i} \xi_{P,i}, \\ g_i(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}, \boldsymbol{\xi}) &\triangleq V_{l,i}^{-1} \left(\tilde{g}_i(\boldsymbol{\alpha}, \mathbf{V}_l, \boldsymbol{\gamma}) - Q_i^0 - \sigma_{Q,i} \xi_{Q,i} \right), \quad i \in \{N_g + 1, N_g + 2, \dots, N_{\text{bus}}\}, \end{aligned}$$

where the components of $\boldsymbol{\xi}(t) \triangleq [\boldsymbol{\xi}_P^\top(t), \boldsymbol{\xi}_Q^\top(t)]^\top$ are defined by Equation (3.22) and taken to be uncorrelated. We take all noise processes to have identical correlation length of $\tau = 10^{-2}$ and set all $\sigma_{P,i} \approx 2.19$ and $\sigma_{Q,i} \approx 1.55$. For our experiments below, this puts the RODE in the high-noise regime. Many methodologies that use large deviation arguments to obtain asymptotic transmission failure rates, which typically require the existence of a nice closed-form stationary measure such as a Gibbs measure, usually do not perform well in this setting [182].

All experiments that follow are over the time interval $[0, T_f]$ with $T_f = 0.5$ for the IEEE 14-Bus System, giving a 47-dimensional RODE system. An equilibrium point of the deterministic power system is found by solving the optimal power flow via MATPOWER [234]. The equilibrium point is treated as a deterministic initial condition for the RODE, which is burned in via N_{MC}^{tr} MC simulations over the entire time horizon. During this burn-in, all tripping thresholds are set to $\mathbf{H} \equiv 1$ (equivalent to a line rating of 200 megavolt amperes) so that no lines are tripped. The resulting samples of $\mathbf{x}(T_f)$ are then treated as independent samples of the random initial condition $\mathbf{x}^0$ at time $t = 0$, which are used in generating MC realizations of the RODE system over the time interval $(0, T_f]$ as well as post-processed with KDE [26] to compute the RoPDF for the QoI at $t = 0$. After the noise burn-in period, we perturb the system out of its quasi-equilibrium by manually removing line 15 at time $t = 0$. Additionally, at time $t = 0$, we reduce the thresholds for lines 12 and 17 to $H_{12} = 0.0135$ and $H_{17} = 0.0125$, respectively, to mimic so-called “weak lines,” which cannot afford normal load flow, that occur in physical power systems under various circumstances, e.g., bad weather.

We set $\text{diag}(\mathbf{M}_g) = 5.3 \times 10^{-2}$, $\text{diag}(\mathbf{D}_g) = 5 \times 10^{-2}$, $\text{diag}(\mathbf{D}_l) = 5 \times 10^{-3}$, and $\text{diag}(\mathbf{D}_v) = 10^{-2}$, which are the same parameter choices for the experiments in [182]. We determined (via convergence studies) that $\text{diag}(\mathbf{D}_\gamma) = 10^{-3}$ is the largest possible value that achieves realistic tripping dynamics.

RoPDF Equation & Numerics

Since we are interested in quantifying the uncertainty concerning line failures in the power grid, we consider the real-valued QoI to be $z(\mathbf{x}(t)) = \gamma_k(t)$, which represents the operational status of the k -th power line. Since the phase space of γ_k is technically unbounded, we let $Z_k \in \mathbb{R}$ represent a variable in its phase space. Following the derivation in Section 3.2, the exact RoPDF equation for the marginal PDF $f_{\gamma_k}(Z_k; t)$ of $\gamma_k(t)$ is given by

$$\begin{aligned} \frac{\partial f_{\gamma_k}}{\partial t} + \frac{\partial}{\partial Z_k} \left(-D_{\gamma, kk}^{-1} (\mathcal{R}(Z_k, t) - H_k \theta_k(Z_k)) \right) &= 0, \\ f_{\gamma_k}(Z_k; 0) &= f_{\gamma_k}^0(Z_k), \end{aligned} \quad (3.27)$$

with vanishing boundary conditions, where the regression function is the conditional expectation $\mathcal{R}(Z_k, t) \triangleq \langle \tilde{h}_k(\boldsymbol{\alpha}, \mathbf{V}_l) \mid \gamma_k(t) = Z_k \rangle$. Since the thresholding function θ_k depends only on the QoI γ_k , the advection coefficient in Equation (3.27) is partially separable, and thus $\theta_k(Z_k)$ has been pulled out of the conditional expectation. In the experiments that follow, $\mathcal{R}$ is always estimated by $\hat{\mathcal{R}}$ via GLLR for each $t_{m_l} \in \mathbf{T}_\nu$. In all MC simulations, three lines underwent tripping dynamics, including both weak lines. Out of these three, the RoPDF for line 12 had the most complex dynamics. Hence, we limit our presentation to the $k = 12$ case.

As seen in Figure 3.6 and Figure 3.7, the dynamics of Equation (3.27) transition the RoPDF from unimodal to bimodal, where the essential support of the modes is quite small. To accurately capture these dynamics, we take the spatial mesh $\mathbf{Z}_{12}$ to be fixed but nonuniform with ΔZ_{12} ranging from 10^{-4} near the mode locations $Z_{12} \approx 0$ and 1 to 5×10^{-2} in between the modes, resulting in approximately 850 grid cells. If uniform time stepping is used for the Lax-Wendroff discretization, the CFL condition requires $\Delta t = 10^{-6}$. Even though variable time stepping and/or different PDE discretizations may be used to reduced the number of time steps, leaving Δt uniform in our discretization serves to demonstrate

one way in which temporal sparsity can arise. Another comes from the MC simulations and the RODE discretization. For the given stiff power system, an explicit RODE discretization would require time steps as small as 10^{-9} to ensure stability. Our strong-order implicit time stepping can be taken much larger, but a small $\Delta t = 10^{-5}$ to 10^{-4} is still required to accurately capture quick transitions during tripping dynamics. However, to reduce memory requirements, we only store the MC training samples at time increments of 10^{-3} . Hence, the sparsity factor with respect to the RODE discretization is 10 to 10^2 but is 10^3 compared to the PDE discretization. Given our choice of notation, our sparsity factor ν refers to the latter, i.e., $\nu = 10^3$. For this application, we do not consider additional sparsity factors since $\nu = 10^3$ is considerably large for the given dynamics, and it has arisen naturally due to memory limitations. Similar to our presentation in Section 3.5.1, we limit our head-to-head comparisons to a single sample size of $N_{\text{MC}}^{\text{tr}} = 2 \times 10^3$, but vary this samples size to determine overall convergence rates. The total number of MC trials required to compute the yardstick solution f_{MC} for $\gamma_{12}(t)$ is $N_{\text{MC}} = 10^5$. Recall that in the DNN formulation, all hidden layers have 32 neurons, but, as in Section 3.5.1, network depth is chosen optimally. For $\nu = 10^3$ and all $N_{\text{MC}}^{\text{tr}}$ considered, 7 hidden layers is the optimal depth.




Figure 3.5: $N_{\text{MC}}^{\text{tr}} = 2 \times 10^3$ MC realizations (blue dots) of $(\gamma_{12}, \tilde{h}_{12})$ at time $t = 0.1$. (Left) GLLR estimate of $\hat{\mathcal{R}}$ using 10-fold CV for bandwidth selection. (Middle) Data transformed to standard Gaussian variates and corresponding GLLR fit with simple plug-in bandwidth. (Right) Fit from the middle plot transformed back to the original scale to obtain a more accurate estimate of $\hat{\mathcal{R}}$.

Regression

Regarding the regression estimates $\hat{\mathcal{R}}$, when the k -th line in the system is tripped and γ_k is the QoI, the underlying dynamics make regression on the MC sample data difficult. However, at any given time, the response (line energy) data $\tilde{h}_k(\boldsymbol{\alpha}, \mathbf{V}_l)$ is always nonnegative and nicely right-skewed, allowing these variates to be efficiently transformed into standard normal variates via the one-parameter Box-Cox transformation. The transformation parameter is selected via maximum likelihood estimation with the Shapiro-Wilk goodness-of-fit test as was done in [11]. The qualitative behavior of the underlying predictor data γ_k associated with $\hat{\mathcal{R}}$ drastically changes during the transition period after the line is tripped, and therefore no single parametric transformation can be expected to perform well. Since we must already compute the observations $f_{\text{MC}}^{\text{tr}, \nu}$ for the assimilation procedures, we can easily convert these PDFs into CDFs via inexpensive quadrature. Evaluating these CDFs at the $N_{\text{MC}}^{\text{tr}}$ -many variates of γ_k via 1D interpolation gives approximately uniform variates on $[0, 1]$. Applying the inverse standard normal CDF to these variates gives approximately standard Gaussian variates. Given that the predictor and response data have each been transformed to (univariate) standard Gaussian variates, an optimal plug-in bandwidth estimator for Gaussian data can be employed for GLLR [107, Ch. 3], avoiding costly CV procedures. Letting $\text{med}(\cdot)$ denote the median operator, the optimal, robust estimator is given by $\hat{s}_k \triangleq \sqrt{\hat{s}\{\gamma_k\}\hat{s}\{\tilde{h}_k\}}$, where $\hat{s}\{\cdot\}$ is defined as

$$\hat{s}\{y\} \triangleq \left(\frac{4}{3N_{\text{MC}}^{\text{tr}}} \right)^{0.2} \text{med}(|y - \text{med}(y)|) / 0.6745. \quad (3.28)$$

Figure 3.5 (left) displays $N_{\text{MC}}^{\text{tr}} = 2 \times 10^3$ MC realizations (blue dots) of $(\gamma_{12}, \tilde{h}_{12})$ at time $t = 0.1$. On this original scale, the data near $Z_{12} \approx 0$ (the RoPDF's left mode) nearly forms a vertical line. In order to remotely capture this behavior with GLLR, $\hat{\mathcal{R}}$ becomes negative and too rough between the modes, even with 10-fold CV for bandwidth selection

(black curve). However, after applying the transformations and performing GLLR with the much cheaper plug-in bandwidth (middle), the fit becomes excellent. The inverse transforms are applied to obtain a much cheaper and more accurate $\hat{\mathcal{R}}$ (right). Its temporal evolution, in addition to the full advection coefficient, is given in Figure 3.6.




Figure 3.6: (Left) Learned regression function $\hat{\mathcal{R}}(Z_{12}, t)$ estimated from $N_{\text{MC}}^{\text{tr}} = 2 \times 10^3$ MC realizations of equation (3.25) at sparse observation times $\mathbf{T}_\nu$ with $\nu = 10^3$. The behavior of $\hat{\mathcal{R}}$ is difficult to distinguish using a single scale (bottom). Hence, we partition the domain about the (red dashed) line $Z_{12} \approx 0$ and plot the left and right sides on difference scales (top). (Middle) The partially separable advection coefficient based on the learned $\hat{\mathcal{R}}$. Likewise, the domain is split about the (gray dotted) line $Z_{12} \approx 0.95$. The left side is omitted on top as it closely resembles the bottom plot. (Right) Evolution of the homogeneous RoPDF $f_{\gamma_{12}}^h$ for $\gamma_{12}(t)$. The evolution is displayed only near the phase space boundaries since the middle portion of the domain always corresponds to relatively low probability states.

Error Analysis

The temporal evolution of the solution $f_{\gamma_{12}}^h$ to the homogeneous equation in Equation (3.27) with estimated $\hat{\mathcal{R}}$ is given in Figure 3.6 (right) as well as snapshots at times $t = 5 \times 10^{-3}$, 0.1, and 0.4 in Figure 3.7. At $t = 0$, $f_{\gamma_{12}}^h$ is unimodal and nearly symmetric around its original deterministic equilibrium point, indicating that the line is fully operational. Due to the line's low power rating together with stochastic fluctuations at the loads, as time evolves, the probability of transmission failure increases and quickly skews the density left. The small value of $\mathbf{D}_\gamma$ causes the RoPDF to transition extraordinarily fast to form a new mode near $Z_{12} \approx 0$ —its mass is approximately the line's failure probability at any given

time. The scale at which this transience occurs in addition to the complexity of $\mathcal{R}$ from multiphysics is precisely why the accuracy of $f_{\gamma_{12}}^h$ degrades when data is temporally sparse. This was also true for the linear system in Section 3.5.1; however, it is much more apparent in this application. Due to the nature of $\hat{\mathcal{R}}$ near zero (see Figure 3.5 (right) and Figure 3.6 (left)), small shifts away from the true $\mathcal{R}$ can result in large changes in magnitude, which is exacerbated by interpolation. Moreover, $\mathcal{R}$ also rapidly oscillates in time (on small scales), which is not captured well by the estimate $\hat{\mathcal{R}}$ due to data sparsity. The combined difficulties introduce error to the homogeneous solution that compounds over time, as seen in the RoPDF snapshot Figure 3.7 (right) and the L_1 error evolution in Figure 3.8 (middle).

Also seen in Figure 3.7, the nudged and DNN observers perform well over the timecourse compared to the homogeneous solution. Both observers capture all qualitative aspects of the yardstick solution exceptionally well during early and middle times. At late times (right), both struggle to fully capture the right mode; however, this can be attributed to the vastly differing scales of the two modes. Note, these modes at $t = 0.4$ (right) are displayed on different scales to emphasize this discrepancy at the right mode. Given that the mass of the right mode is only a fraction of the RoPDF’s total mass, such discrepancies do not significantly influence either observer’s error against the yardstick MC solution, as seen in Figure 3.8 (middle).

While small at late times, the right mode of $f_{\gamma_{12}}$ does not vanish, even if the final time T_f is increased significantly. Hence, there is a nonzero, albeit small, probability that the line does not trip. If one desires a highly refined estimate of this probability, observer discrepancy at the RoPDF’s right mode must be reduced. When the sparsity level ν is fixed, the only surefire way to improve the nudged observer is to increase N_{MC}^{tr} , i.e., the amount of training data. This is also true for the DNN observer, but to a lesser extent by means of normalization. Unlike the linear RODE in Section 3.5.1, there is no clear-cut approach to normalizing the RoPDFs and observations locations for improved training. However,




Figure 3.7: Comparison of the yardstick f_{MC} (solid black), the homogeneous solution $f_{\gamma_{12}}^h$ (dot-dashed red), the nudged observer $\hat{f}_{\gamma_{12}}^{\text{NR}}$ (dotted blue), and the DNN observer $\hat{f}_{\gamma_{12}}^{\text{DNN}}$ (dashed purple) for $\nu = 10^3$ and $N_{\text{MC}}^{\text{tr}} = 2 \times 10^3$ at times $t = 0.005$ (left), 0.1 (middle), and 0.4 , (right). Similar to the PDFs in Figure 3.6, the phase space is restricted near the boundaries to emphasize the bimodal behavior at the latter two times. Moreover, at $t = 0.4$ (right), the modes are given on two scales, revealing that the right mode shrinks, but does not vanish.

since the CDFs corresponding to the observations $f_{\text{MC}}^{\text{tr}, \nu}$ were computed for $\hat{\mathcal{R}}$ estimation, for each observation time, we multiply the spatial grid by these CDFs. We then shift, scale, and apply the square root transform to obtain new observation locations for training. This serves to disperse the essential supports of the RoPDF modes, making them easier to learn. In addition to applying square root transforms to the RoPDFs, we also shift and scale them, along with the observation times, so that the RoPDFs and spatiotemporal locations are all on the same scale. This approach to normalization considerably improves DNN estimates of the right mode at late times. Alternative approaches to normalization may yield better results, but none are perfect—the observer still depends on the quality of $f_{\text{MC}}^{\text{tr}, \nu}$ and therefore $N_{\text{MC}}^{\text{tr}}$.

Figure 3.8 (left) demonstrates the qualitative behavior to be expected of the learned defect solution $\hat{f}_{\gamma_{12}}^{\text{d}}$ associated with the DNN observer (dashed purple). The defect corresponding to the nudged observer (dotted blue), computed as $\hat{f}_{\gamma_{12}}^{\text{NR}} - f_{\gamma_{12}}^h$, is also plotted for reference. They are nonperiodic and exhibit steep gradients. As seen in the middle plot, the errors of $\hat{f}_{\gamma_{12}}^{\text{NR}}$ and $\hat{f}_{\gamma_{12}}^{\text{DNN}}$ against the yardstick solution are quite similar, with $\hat{f}_{\gamma_{12}}^{\text{NR}}$ performing slightly better at early times, i.e., during the initial transience. As previously discussed,




Figure 3.8: (Left) The defect solutions $\hat{f}_{\gamma_{12}}^d$ for the DNN and nudged observers at time $t = 0.1$ with sparsity factor $\nu = 10^3$ and $N_{MC}^{tr} = 2 \times 10^3$ MC samples. The former is the (prediction of the) DNN in Equation (3.18) while the latter is computed *ex post facto* as $\hat{f}_{\gamma_{12}}^{NR} - \hat{f}_{\gamma_{12}}^h$. (Middle) Temporal evolution of L_1 errors for the homogeneous solution, $\hat{f}_{\gamma_{12}}^h$, the nudged observer $\hat{f}_{\gamma_{12}}^{NR}$, and the DNN observer $\hat{f}_{\gamma_{12}}^{DNN}$. (Right) Convergence rates, on a log-log scale, for the normalized spatiotemporal L_1 errors of the nudged and DNN observers as N_{MC}^{tr} increases.

DNN observer error could possibly be improved to match that of nudging via alternative normalization. However, we attribute the nudged observer's better success to its ability to dynamically overcome the highly nontrivial error distributions associated with these RoPDFs. We remark that, in this setting, the EnKF would likely perform poorly against the nudged observer due to the errors being large and non-Gaussian.

Figure 3.8 (right) displays that even though the power systems dynamics are significantly more complex, stiff, and higher dimensional than the linear system from Section 3.5.1, the RoPDF method obtains the same $\mathcal{O}(1/N_{MC}^{tr})$ convergence as the number of MC realizations increases, demonstrating the method's robustness. Moreover, given that the yardstick MC solution requires $N_{MC} = 10^5$ realizations of the stiff, 47-dimensional RODE system (even with fast, adaptive KDE) and the RoPDF method needs fewer than $N_{MC}^{tr} = 10^3$ realizations to achieve (less than) 1% L_1 error, regardless of the assimilation approach, the method requires relatively few computational resources. For this application, the computational costs of numerically integrating the RoPDF equation, including the additional cost of DNN training, is but a small fraction of total costs. Therefore, comparing N_{MC}^{tr} to N_{MC} represents the computational speedup of the method sufficiently well. In particular, we see a speedup

of at least two orders of magnitude (depending on desired error tolerance) of the RoPDF method compared to the MC approach.

3.6 Summary and Conclusions

In this work, we have developed a physics-informed framework for studying uncertainty propagation of physical quantities of interest in high-dimensional and multiscale stochastic dynamical systems. In particular, we presented a derivation of an exact RoPDF equation and a regression-based approach to closures, enabling the characterization of full probabilistic profiles at all times with low computational complexity. Furthermore, we introduced two physics-informed data assimilation procedures to address issues arising in stiff systems, namely nudging/Newtonian relaxation and deep neural networks, which assimilates in low-fidelity observations at sparse observation times with negligible cost, improving density estimates. Finally, we showed the accuracy of our method on characterizing uncertainty in both a synthetic stiff linear system and an at-scale power system cascading failure model using IEEE case data. The results of our method demonstrate promising and practical uses in the prediction of complex stochastic phenomena.

Several challenges and opportunities arise following this work. Firstly, convergence rates of the RoPDF method are dependent on three factors: (1) RoPDF observation error via KDE, (2) error from numerical PDE integration, and (3) estimation error from regression functions. A precise characterization of solution accuracy can be analyzed from a learning theory perspective. Some regularity of the RODE is required to guarantee the existence of a solution to the RoPDF equation (see Theorem 3); however, more importantly, specific regression methods typically require additional assumptions on the conditional expectations, e.g., smoothness. How these assumptions contribute to the convergence of the discrepancy $\langle \mathcal{M} \rangle$ under different choices of models, loss functions, and optimization algorithms remains to be theoretically studied. Furthermore, the notion of data stability (see, for instance, [89])

is required to formally define the error terms that depend on statistical models. The authors believe that it suffices to analyze data stability with respect to RODE regularity, whose noise process directly impacts the consistency of nonparametric estimators [204].

Secondly, we performed preliminary experimentation of using DNNs for the discovery of model defects. We believe that DNNs could have strong extrapolation power once trained with more sophisticated architectures. To name a few, Fourier-based DNNs are known to capture stiff dynamics well. Additionally, long short-term memory networks can be used to impose temporal ordering, which is more suitable for learning from (sparse) time-series observations. Particularly, when observations cease to exist, the design of reliable DNN extrapolation is necessary, which was beyond the scope of this study. Finally, the natural extension to uncertainty quantification for vector-valued QoIs is of great practical interest, such as rare-event probability estimations for multiple line failures in the power system model (Equation (3.24)). However, the issue of dimensionality returns when the reduced state space itself is high-dimensional, which may potentially be resolved via structured low-complexity methods, such as tensor-networks, flow-based generative models, and/or a combination of such strategies where an initial product measure can be formed from marginal densities solved using the 1D RoPDF method, and then optimized to approach the correct reduced-order joint density.

3.7 Appendix

3.7.1 *Nudging Convergence*

For general random hyperbolic conservation laws, the method of distributions is formulated similarly to that of Theorem 3 and [143, Eq. 3] for RoPDF and joint PDF equations (respectively) corresponding to the RODE of Equation (3.1). The kinetic description of the hyperbolic system is precisely the deterministic equation for the “raw PDF” II. How-

ever, when the governing random PDE exhibits shocks, the method of distributions for Π breaks down at singularities. This can be overcome by partitioning the domain and tracking shocks analytically, which was done in [4, 209] for the water-hammer and Buckley-Leverett equations, respectively. However, analytically tracking shocks is rarely possible for general nonlinear hyperbolic PDEs. Instead, the kinetic defect term/collision operator $\mathcal{M}$ may be introduced as a source function in the raw PDF equation, incorporating all information regarding discontinuities. When the hyperbolic system exhibits smooth solutions, $\mathcal{M}$ is unique and identically zero. Otherwise, it can be written as the partial derivative of what is known as the kinetic entropy defect measure—it is exact, albeit generally unknown *a priori*. Learning this defect in the CDF equations of nonlinear scalar conservation laws with random initial data was the focus of [25], which largely motivated our extension to the setting of reduced-order equations.

It was shown in [27] that nudging hyperbolic conservation laws at the kinetic level does not perturb the stability of the macroscopic system, which is beneficial for establishing strong convergence. In our setting of RoPDF equations, the conservation law in Equation (3.4) for Π_{x_k} is a kinetic description and is exact. Therefore, the defect $\mathcal{M}$ vanishes and (scalar) nudging ensures that the corresponding observer $\hat{\Pi}_{x_k}$ converges globally and exponentially in L_1 to Π_{x_k} with rate $\lambda > 0$ when $H(\Pi_{x_k}) \equiv \Pi_{x_k}$, i.e., when observations are complete and exact. By virtue of the triangle inequality,

$$\|\hat{f}_{x_k} - f_{x_k}\|_1 \leq \langle \|\hat{\Pi}_{x_k} - \Pi_{x_k}\|_1 \rangle, \quad (3.29)$$

the nudged observer $\hat{f}_{x_k}$ corresponding to the exact RoPDF equation enjoys the same global convergence to f_{x_k} , i.e., the solution to Equation (3.12). In the case of temporally discrete observations, i.e., when the observations are sparse in time and complete in space, global convergence cannot be obtained. However, if the observations are interpolated over finite

time intervals of length $T_w > 0$ via the correction term

$$\lambda \sum_{l \in I} \phi_{T_w}(t - t_{m_l}) \left(\Pi_{x_k}(X_k, t_{m_l}) - \hat{\Pi}_{x_k}(X_k, t_{m_l}) \right),$$

then, for any given time $T > 0$, the nudged kinetic observer has bounded L_1 convergence:

$$\|\hat{\Pi}_{x_k}(X_k, T) - \Pi_{x_k}(X_k, T)\|_1 \leq C_0 e^{\lambda L} + T_w \mathcal{I}(T), \quad (3.30)$$

where L is the number of time steps in $[0, T - T_w]$, C_0 is a constant depending on the initial condition, and $\mathcal{I}$ is a convergence-rate dependent quantity. Taking the ensemble mean provides a λ -dependent convergence bound for nudging the exact RoPDF Equation (3.12). When the observations of Π_{x_k} are noisy, under sufficient regularity conditions, one obtains a strong upper bound on the observation error in a homogeneous Sobolev norm and an optimal (scalar) nudging coefficient $\lambda > 0$ (see [27]).

In our nudging formulation, we have replaced the conditional expectations $\mathcal{R}_i$ with smooth estimators $\hat{\mathcal{R}}_i$, meaning that the kinetic defect term does not vanish. For exact but temporally sparse observations, this amounts to adding the term $\sup_{0 < t \leq T} \|\mathcal{M}(X_k, t)\|_1 / \lambda$ to the bound in Equation (3.30). However, the more concerning issue is that we have introduced observation noise at the macroscopic level (via the low-fidelity estimates $f_{\text{MC}}^{\text{tr}, \nu}$) rather than the kinetic level. Moreover, we are not aware of any existing literature that has addressed theoretical convergence in this setting. Since measurement noise often occurs on the macroscopic level in many applications, such results would be of great interest, and are indeed the focus of an ongoing body of work. In the meantime, we rely on the mounting empirical evidence for the practical nudging of ODEs/PDEs, including the new results for RoPDF equations in Section 3.5.

CHAPTER 4

DATA-DRIVEN ESTIMATION OF FAILURE PROBABILITIES IN CORRELATED STRUCTURE-PRESERVING STOCHASTIC POWER SYSTEM MODELS

4.1 Introduction

Recent blackout events and fires caused by power lines have continued to raise awareness of the need for designing robust circuit systems and devising mitigation strategies when catastrophic events occur [126]. Uncertainties in peak power demands, operative states of components, load, and seasonal factors contribute to considerable difficulties in modeling line failures [56, 156].

Failure events are typically studied through traversal/clustering of time-dependent graphs [222] or by solving differential systems that incorporate transient (e.g. line-removal) dynamics [65]. Particularly for large cases, the first class of approaches is restricted to predicting local behavior that may not yield accurate results corresponding to real-world observations [97]. On the other hand, extensive simulations are needed in order to obtain statistical information of failure distributions as tail events, making the dynamical systems approach computationally intensive. In this work, we build upon the systematic approach of [150] to incorporate uncertainty in a structure-preserving dynamics model by including the Ornstein–Uhlenbeck (OU) process as load stochastic fluctuations. In the absence of approximations, resolving the full probability profile of all random states suffers from the curse of dimensionality. Although reduction methods such as the method of moments (MoM) [141] and Karhunen–Loève (KL) expansions [174] exist, they are known to underperform standard Monte Carlo simulations and kernel density estimation (KDE) for realistic power system models. In particular, the MoM requires state data to be nearly Gaussian, while KL expansions degrade when the noise input exhibits a short correlation time scale [142].

To address this need, instead of approximating the full joint probability density, we exploit the fact that in the model of [150] the probability of failure of a transmission line depends on a scalar function that can be efficiently evaluated. Subsequently, we directly propagate its probability density function (PDF) by expressing it as a solution of a low-dimensional partial differential equation (PDE) that we solve numerically. This reduced-order formulation allows us to simultaneously consider the joint density of multiple transmission lines, whose conditional probabilistic structure is typically difficult to simulate using kinetic Monte Carlo methods [182]. Furthermore, the method is data-driven, where measurements from system states can be incorporated into estimating unknown terms in the PDE using regression methods. The resulting PDE is a linear equation that is readily solved with standard finite volume schemes, with comparable accuracy to plug-in Gaussian KDE, while using significantly fewer samples from the law of the underlying stochastic process. We also note that the previous approach in [182], while being more sample-efficient, computed only the large deviation (low-temperature limit) approximation of the failure probability whereas this approach with enough samples will converge to the true probability of failure.

The paper is organized as follows. Section 4.2 introduces the general power system model along with the correlated noise process, which is jointly considered as a diffusion process admitting a Fokker–Planck equation (FPE). In Section 4.3 we discuss the derivation of the reduced-order equation by integrating over extraneous variables in the FPE, allowing one to formulate the density of arbitrary quantities of interest. In addition, we discuss estimation of unknown terms in the PDE as a regression problem. Effectiveness of the method is tested in Section 4.4 on the classical multimachine model with an emphasis on use cases of computing tail event probabilities and correlation. In Section 4.6 we summarize our work and discuss possible directions for extension.

4.2 Stochastic Power System Model

A stochastic power system model is generally described by a set of index-1 governing differential algebraic equations:

$$\begin{aligned}\dot{\mathbf{x}}(t) &= \mathbf{f}(\mathbf{x}(t), \mathbf{y}(t); \boldsymbol{\eta}(t)) \\ \mathbf{0} &= \mathbf{g}(\mathbf{x}(t), \mathbf{y}(t); \boldsymbol{\eta}(t)),\end{aligned}\tag{4.1}$$

where $\mathbf{x}(t)$ are the system states at time t and $\mathbf{y}(t)$ accounts for any auxiliary or algebraic variables. Furthermore, to account for random fluctuations in the states, a noise process $\boldsymbol{\eta}(t)$ is introduced with the following general form [150]:

$$\begin{aligned}\dot{\boldsymbol{\eta}}(t) &= \mathbf{a}(\boldsymbol{\eta}(t)) + \mathbf{b}(\boldsymbol{\eta}(t)) \odot \dot{\boldsymbol{\zeta}}(t) \\ \dot{\boldsymbol{\zeta}}(t) &= \mathbf{C} \cdot \frac{d\mathbf{w}(t)}{dt},\end{aligned}\tag{4.2}$$

where $\mathbf{w}(t)$ denotes the multidimensional standard Wiener process in $\mathbb{R}^n$, such that $d\mathbf{w}/dt$ is formally interpreted as white noise. In essence, the second equation of (4.2) defines a set of correlated Wiener processes with correlation matrix $\mathbf{R} := \mathbf{C}\mathbf{C}^T$. We assume Lipschitz continuity of $\mathbf{a}, \mathbf{b}$ in (4.2), such that (4.1) admits the following form of diffusion processes, where we define $\mathbf{z}(t) := [\mathbf{x}^T(t), \mathbf{y}^T(t), \boldsymbol{\eta}^T(t)]^T \in \mathbb{R}^N$, with $N \gg n$:

$$\begin{aligned}d\mathbf{z}(t) &= \boldsymbol{\mu}(\mathbf{z}(t), t)dt + \boldsymbol{\sigma}(\mathbf{z}(t), t)d\mathbf{W}(t) \\ \mathbf{z}(0) &\sim f_{\mathbf{z}_0},\end{aligned}\tag{4.3}$$

where $\mathbf{W}(t)$ is the $\mathbb{R}^N$ -valued standard Wiener process. The initial condition $\mathbf{z}_0$ may be either random with probability density $f_{\mathbf{z}_0}$ or deterministic, in which case $f_{\mathbf{z}_0} \equiv \boldsymbol{\delta}_{\mathbf{Z}^*}$ is the (multidimensional) Dirac delta distribution at a fixed point $\mathbf{Z}^* \in \mathbb{R}^N$. Furthermore, let $\boldsymbol{\mu}_t \equiv \boldsymbol{\mu}(\mathbf{z}(t), t) : \mathbb{R}^N \times \mathbb{R}^+ \rightarrow \mathbb{R}^N$ define the vector of drifts, and let $\boldsymbol{\sigma}_t \equiv \boldsymbol{\sigma}(\mathbf{z}(t), t) : \mathbb{R}^N \times \mathbb{R}^+ \rightarrow \mathbb{R}^{N \times N}$ define the diffusion matrix. The velocity and noise terms of (4.3) are

assumed to satisfy sufficient regularity, which guarantee the existence of solutions whose paths are almost surely continuous.

The Itô process (4.3) admits a full-order Fokker–Planck equation [179] taking the following form:

$$\begin{aligned} \frac{\partial f_{\mathbf{z}}}{\partial t} + \sum_{i=1}^N \frac{\partial}{\partial Z_i} (\mu_i(\mathbf{Z}, t) f_{\mathbf{z}}) &= \sum_{i,j=1}^N \frac{\partial^2}{\partial Z_i \partial Z_j} (\mathcal{D}_{ij}(\mathbf{Z}, t) f_{\mathbf{z}}) \\ f_{\mathbf{z}}(\mathbf{Z}, 0) &= f_{\mathbf{z}_0}(\mathbf{Z}), \end{aligned} \tag{4.4}$$

The advection and diffusion coefficients of Equation (4.4) are assumed to decay sufficiently fast at $\pm\infty$ so that vanishing boundary conditions can be employed. Furthermore, we denote $\mathcal{D} := \frac{1}{2} \boldsymbol{\sigma}_t \boldsymbol{\sigma}_t^T \in \mathbb{R}^{N \times N}$. When drift and diffusion are time-dependent, it is generally not the case that (4.4) possesses a stationary solution. We refer the reader to [162] for specific investigated cases where an analytic derivation is possible.

4.3 Proposed Method

To estimate probabilities of certain events defined using a forward model, one may either simulate Monte Carlo trajectories of (4.3) and post-process with KDE, or solve (4.4) and perform numerical integration on a computational mesh. Both approaches suffer from the curse of dimensionality when the effective system size N becomes large [186, 66]. Instead of seeking to resolve the full-state dynamics, it often suffices for practical purposes to observe the evolution of a low-dimensional quantity of interest (QoI) over the (finite) domain of numerical simulations. Therefore, we consider deriving the PDF dynamics for the QoI directly, and solve only a low-dimensional FPE-like equation (referred to as the “reduced-order PDF equation” of the QoI). In this section, we present a general extension for formally deriving the PDF equation for QoIs of arbitrary dimensions. We begin by reviewing the derivation for scalar QoIs in the next section.

4.3.1 Reduced-Order PDF Equations

We let $u : \mathbb{R}^N \rightarrow \mathbb{R}$ be a deterministic, second-order continuously differentiable function of the system states, representing a QoI evaluated from the states of system (4.3). Furthermore, let $u(t) \equiv u(\mathbf{z}(t))$ denote the state of the QoI at time t , and let $U \in \mathbb{R}$ denote the corresponding phase space variable. Applying Itô's lemma, we see that the QoI follows an associated process:

$$du(t) = ((\nabla_{\mathbf{z}} u)^T \boldsymbol{\mu}_t + \frac{1}{2} \text{tr}(\boldsymbol{\sigma}_t^T (H_{\mathbf{z}} u) \boldsymbol{\sigma}_t)) dt + (\nabla_{\mathbf{z}} u)^T \boldsymbol{\sigma}_t d\mathbf{W}_t, \quad (4.5)$$

where we let $\nabla_{\mathbf{z}} := [\frac{\partial}{\partial z_1}, \dots, \frac{\partial}{\partial z_N}]^T$. $H_{\mathbf{z}}(\cdot)$ denotes the Hessian operator with respect to states $\mathbf{z}(t)$. For simplicity of notations, we let μ^u denote the drift term of the process (4.5). To derive a deterministic governing equation for the probability density f_u , we augment the states by considering the joint vector $\tilde{\mathbf{z}}(t) := [\mathbf{z}^T(t), u(t)]^T \in \mathbb{R}^{N+1}$. Similarly, let $\widetilde{\mathbf{W}}(t) := [\mathbf{W}^T(t), w(t)]^T$ denote an $(N+1)$ -dimensional standard Wiener process, with $w(t)$ independent of $\mathbf{W}(t)$. We arrive at the following system:

$$d\tilde{\mathbf{z}}(t) = \tilde{\boldsymbol{\mu}}_t dt + \tilde{\boldsymbol{\sigma}}_t d\widetilde{\mathbf{W}}(t) = \begin{bmatrix} \boldsymbol{\mu}_t \\ \mu^u \end{bmatrix} dt + \begin{bmatrix} \boldsymbol{\sigma}_t & \mathbf{0} \\ (\nabla_{\mathbf{z}} u)^T \boldsymbol{\sigma}_t & 0 \end{bmatrix} d\widetilde{\mathbf{W}}(t). \quad (4.6)$$

Since the function u is deterministic, it does not introduce additional diffusion. Similar to (4.4), the density of the augmented process (4.6) follows the $(N+1)$ -dimensional FPE for the joint states with the same form as (4.4). After factoring the joint density in terms of marginal and conditional densities, that is, the decomposition $f_{\tilde{\mathbf{z}}} = f_{\mathbf{z}|u} \cdot f_u$, we marginalize over the state space of $\mathbf{z}$ and arrive at the following reduced-order PDF equation in u only:

$$\frac{\partial f_u}{\partial t} + \frac{\partial}{\partial U} (\mathbb{E}[\mu^u | u(t) = U] f_u) = \frac{\partial^2}{\partial U^2} (\mathbb{E}[\mathcal{D}^u | u(t) = U] f_u) \quad (4.7)$$

$$f_u(U, 0) = f_{u_0}(U), \quad (4.8)$$

where f_{u_0} denotes the initial density of the QoI, which may be either analytically known or estimated from a KDE using evaluated samples, and

$$\mu^u := (\nabla_{\mathbf{z}} u)^T \boldsymbol{\mu}_t + \frac{1}{2} \text{tr}(\boldsymbol{\sigma}_t^T (H_{\mathbf{z}} u) \boldsymbol{\sigma}_t) \quad (4.9)$$

$$\mathcal{D}^u := (\nabla_{\mathbf{z}} u)^T \mathcal{D} (\nabla_{\mathbf{z}} u) = \frac{1}{2} [\tilde{\boldsymbol{\sigma}}_t \tilde{\boldsymbol{\sigma}}_t^T]_{(N+1)(N+1)}. \quad (4.10)$$

As an illustration, we consider the special case where the QoI denotes a projection onto a specific coordinate of $\mathbf{z}$. Let $u(\mathbf{z}(t)) = z_k(t)$ for some fixed $1 \leq k \leq N$. We have

$$\nabla_{\mathbf{z}} u \equiv \mathbf{e}_k, \quad H_{\mathbf{z}} u \equiv \mathbf{0}, \quad (4.11)$$

where $\mathbf{e}_k := [0, \dots, 0, 1, 0, \dots, 0]^T$ is the k th standard basis vector in $\mathbb{R}^N$. Substituting (4.11) in the derivations of (4.9) and (4.10), we obtain

$$\mu^u \equiv \mu_k, \quad \mathcal{D}^u \equiv \mathcal{D}_{kk}, \quad (4.12)$$

where μ_k is the k th coordinate of the drift vector $\boldsymbol{\mu}_t$, such that the PDF equation becomes

$$\begin{aligned} \frac{\partial f_{z_k}}{\partial t} + \frac{\partial}{\partial Z_k} (\mathbb{E}[\mu_k(\mathbf{z}(t), t) | z_k(t) = Z_k] f_{z_k}) \\ = \frac{\partial^2}{\partial Z_k^2} (\mathbb{E}[\mathcal{D}_{kk}(\mathbf{z}(t), t) | z_k(t) = Z_k] f_{z_k}), \end{aligned}$$

which is consistent with the results presented in [31] for single components of the state vector.

4.3.2 Joint Reduced-Order PDF Equations

In applications where conditional events are involved, it is helpful to consider the joint probabilistic structure of vector-valued QoIs. The derivations presented in Section 4.3.1 can be readily extended to multiple dimensions. To do so, we define a multivariate mapping with second-order continuously differentiable components, $\mathbf{u}(\mathbf{z}(t)) = [u_1(\mathbf{z}(t)), \dots, u_{N_R}(\mathbf{z}(t))]^T$, with $1 < N_R \ll N$. Each component $u_j, 1 \leq j \leq N_R$ of the vector-valued mapping $\mathbf{u}$ is understood as a mapping from $\mathbb{R}^N$ to $\mathbb{R}$. We apply the multivariate extension of Itô's lemma and observe that the vector of QoIs, $\mathbf{u}(t) \equiv \mathbf{u}(\mathbf{z}(t))$, follows a diffusion process using the same derivation in (4.5) for each component:

$$d\mathbf{u}(t) = (\mathbf{J}_\mathbf{u}\boldsymbol{\mu}_t + \mathbf{h}_\mathbf{u})dt + \mathbf{J}_\mathbf{u}\boldsymbol{\sigma}_t d\mathbf{W}_t,$$

where $\mathbf{J}_\mathbf{u} \in \mathbb{R}^{N_R \times N}$ is the Jacobian matrix of the vector-valued function $\mathbf{u}(\cdot)$ with respect to the states $\mathbf{z}(t)$, and

$$\mathbf{h}_\mathbf{u} := \frac{1}{2} \begin{bmatrix} \text{tr}(\boldsymbol{\sigma}_t^T (H\mathbf{z}u_1)\boldsymbol{\sigma}_t) \\ \vdots \\ \text{tr}(\boldsymbol{\sigma}_t^T (H\mathbf{z}u_{N_R})\boldsymbol{\sigma}_t) \end{bmatrix}.$$

With a similar argument to (4.7), we arrive at the reduced-order PDF equation joint in multiple QoIs, which is a deterministic PDE of N_R -dimensions:

$$\frac{\partial f_\mathbf{u}}{\partial t} + \sum_{i=1}^{N_R} \frac{\partial}{\partial U_i} (\mathbb{E}[\mu_i^\mathbf{u} | \mathbf{u}(t) = \mathbf{U}] f_\mathbf{u}) = \sum_{i,j=1}^{N_R} \frac{\partial^2}{\partial U_i \partial U_j} (\mathbb{E}[\mathcal{D}_{ij}^\mathbf{u} | \mathbf{u}(t) = \mathbf{U}] f_\mathbf{u}) \quad (4.13)$$

$$f_\mathbf{u}(\mathbf{U}, 0) = f_{\mathbf{u}_0}(\mathbf{U}), \quad (4.14)$$

where we defined the vector of drifts for the chosen subset of QoIs as the following:

$$\boldsymbol{\mu}^\mathbf{u} = [\mu_1^\mathbf{u}, \dots, \mu_{N_R}^\mathbf{u}]^T := \mathbf{J}_\mathbf{u}\boldsymbol{\mu}_t + \mathbf{h}_\mathbf{u}$$

and the augmented matrix similar to that of (4.10) such that

$$\mathcal{D}^{\mathbf{u}} := \mathbf{J}_{\mathbf{u}} \mathcal{D} \mathbf{J}_{\mathbf{u}}^T \in \mathbb{R}^{N_{\mathbf{R}} \times N_{\mathbf{R}}},$$

where $\mathcal{D}$ is defined in (4.4). In particular, the state-dependent conditional expectations in the reduced-order PDF equation (4.13) are $N_{\mathbf{R}}$ -variate functions of the QoIs $\mathbf{u}(t)$.

4.3.3 Approximation of Closure Terms

The analytic expressions of the advection and diffusion terms in equation (4.13) involve high-dimensional integrals and generally do not admit elementary expressions. Numerical integration or expansion-based approximations (e.g., generalized cumulant expansions [118]) are strongly dependent on the original system dimension N and become intractable when dimensions become large.

However, a direct data-driven approximation of the unknown terms is possible by recognizing the estimation of such conditional expectations as a minimization problem under the L^2 -norm induced by the probability distribution of $\mathbf{u}(t)$ (see [91]),

$$\mathcal{R}(\mathbf{U}; t) = \mathbb{E}[h(t) | \mathbf{u}(t) = \mathbf{U}] \in \underset{g \in L^2(\mathbf{u}; t)}{\operatorname{argmin}} \mathbb{E}[\|h(t) - g(\mathbf{u}(t))\|^2], \quad (4.15)$$

where $h(t) \equiv h(\mathbf{z}(t); \omega)$ is a random field depending on the states $\mathbf{z}(t)$ and is ensembled under the conditional law of $\mathbf{u}(t)$ as a result of solving (4.15). The minimization problem (4.15) can be approximately solved by using regression-based methods from simulated trajectories of $h(t)$ and $\mathbf{u}(t)$. The estimated coefficients are henceforth referred to as the “regression functions” of the reduced-order equations. With the main steps outlined, we summarize the procedure for the reduced-order PDF equation in the following section.

4.3.4 Implementation of the Reduced-Order PDF Method

The implementation of (4.13) requires a choice of numerical PDE scheme (e.g., finite difference, finite volume, finite element) coupled with a routine for identifying the coefficient terms (4.15) (e.g., parametric or nonparametric regression methods). For completeness, we present a practical implementation in Algorithm 3 based on the finite volume method, which involves an explicit scheme resolving (4.13) and estimation of advection and diffusion coefficients through the general minimization problem (4.15).

Algorithm 3 Data-driven finite volume scheme with adaptive time-stepping

Inputs: Monte Carlo trajectories of $\mathbf{z}(t) \in \mathbb{R}^N$ from (4.1) at times $t_1, \dots, t_{N_t}$; derived data $\mathbf{u}(\mathbf{z}(t))$, response variables $\boldsymbol{\mu}^{\mathbf{u}}(t)$ and $\mathcal{D}^{\mathbf{u}}(t)$; initial density $f_{\mathbf{u}_0}$.

1. If $f_{\mathbf{u}_0}$ is not known analytically, estimate it via kernel density estimation.
2. For $i = 1, 2, \dots, N_t$, estimate regression functions at time t_{i-1} by approximately solving (4.15), yielding

$$\hat{\mathcal{R}}_{\mu_i^{\mathbf{u}}}(\mathbf{U}; t_{i-1}) \approx \mathbb{E}[\mu_i^{\mathbf{u}} | \mathbf{u}(t_{i-1}) = \mathbf{U}], \quad \hat{\mathcal{R}}_{\mathcal{D}_{ij}^{\mathbf{u}}}(\mathbf{U}; t_{i-1}) \approx \mathbb{E}[\mathcal{D}_{ij}^{\mathbf{u}} | \mathbf{u}(t_{i-1}) = \mathbf{U}].$$

3. Choose the time step Δt to satisfy the Courant–Friedrichs–Lewy (CFL) condition for explicit schemes.
4. Interpolate the coefficients $\hat{\mathcal{R}}_{\mu_i^{\mathbf{u}}}$ and $\hat{\mathcal{R}}_{\mathcal{D}_{ij}^{\mathbf{u}}}$ in time as needed on $[t_{i-1}, t_i]$.
5. Advance $f_{\mathbf{u}}$ from t_{i-1} to $t_i = t_{i-1} + \Delta t$ with an explicit finite-volume method (e.g., upwind) for the approximate PDE

$$\frac{\partial f_{\mathbf{u}}}{\partial t} + \sum_{i=1}^{N_R} \frac{\partial}{\partial U_i} (\hat{\mathcal{R}}_{\mu_i^{\mathbf{u}}} f_{\mathbf{u}}) = \sum_{i,j=1}^{N_R} \frac{\partial^2}{\partial U_i \partial U_j} (\hat{\mathcal{R}}_{\mathcal{D}_{ij}^{\mathbf{u}}} f_{\mathbf{u}}).$$

6. Repeat steps 2–5 for $i = 1, 2, \dots, N_t$.
-

The main source of error in this procedure lies in the estimation and interpolation of advection and diffusion coefficients to the refined temporal grids (as constrained by the CFL condition), where Monte Carlo data is not necessarily available. Apart from interpolating at discrete times, it is also possible to fit a time-varying regression function beforehand and query the advection and diffusion coefficients at arbitrary time points as needed [227], though at a significantly higher computational overhead. For simplicity, we estimate the regression

functions at discrete times and linearly interpolate them to required refinement levels in the experiments of Section 4.4, which we present below.

4.4 IEEE Cases

We consider the classical multimachine model [8] with stochastic fluctuations in the power injections. In particular, we select the noise model to be driven by an OU process, which has been reviewed in [2] to capture short-term power system dynamics. Furthermore, we impose a constant correlation structure among the connected machines.

Let $(\mathcal{V}, \mathcal{E})$ denote the set of nodes and edges describing the topology of the power network comprising n buses. Furthermore, we denote the voltage magnitude, voltage phase angle, and angular velocity of the generators, respectively, as $v_i(t)$, $\delta_i(t)$, $\omega_i(t)$ for $1 \leq i \leq n$. The dynamics is governed by the second-order swing equations with noise fluctuations:

$$\begin{aligned}\dot{\delta}_i(t) &= \omega_i(t) - \omega_R \\ \frac{2h_i}{\omega_R} \dot{\omega}_i(t) + d_i \omega_i(t) &= p_i^m - p_i^e(t) + \eta_i(t),\end{aligned}$$

where p_i^m represents the (constant) mechanical power input into each generator and $\eta_i(t)$ represents the noise injected at machine i . Furthermore, we define the following net active power injection:

$$p_i^e(t) := \sum_{j=1}^n b_{ij} v_i(t) v_j(t) \sin(\delta_i(t) - \delta_j(t)) + \sum_{j=1}^n g_{ij} v_i(t) v_j(t) \cos(\delta_i(t) - \delta_j(t)). \quad (4.16)$$

In vector notation, the complete stochastic model is specified as the following:

$$\dot{\mathbf{v}}(t) = \mathbf{0} \quad (4.17)$$

$$\dot{\boldsymbol{\omega}}(t) = \frac{\omega_R}{2} \mathbf{h}^{-1} \odot \left[-(\boldsymbol{\omega}(t) - \omega_R) \odot \mathbf{d} + \mathbf{p}^m \right] \quad (4.18)$$

$$- \left\{ [\mathbf{G} \odot \cos(\Delta) + \mathbf{B} \odot \sin(\Delta)] \mathbf{v}(t) \right\} \odot \mathbf{v}(t) + \boldsymbol{\eta}(t) \right] \quad (4.19)$$

$$\dot{\boldsymbol{\delta}}(t) = (\boldsymbol{\omega}(t) - \omega_R) \quad (4.20)$$

$$\dot{\boldsymbol{\eta}}(t) = -\theta \boldsymbol{\eta}(t) + \alpha \sqrt{2\theta} \mathbf{C} \cdot \frac{d\mathbf{W}(t)}{dt}. \quad (4.21)$$

In the system (4.17), drift and diffusion parameters are taken as $\theta = 1$ and $\alpha = 0.05$, respectively. $\odot$ denotes Hadamard (elementwise) product. Let $\mathbf{v}(t) = [v_1(t), \dots, v_n(t)]^T$ denote the voltage magnitude of the machines, which are taken as randomly perturbed around a mean operating state that does not have additional temporal dynamics, derived from the constant impedance loads assumption in [8]. $\boldsymbol{\omega}(t) = [\omega_1(t), \dots, \omega_n(t)]^T$ and $\boldsymbol{\delta}(t) = [\delta_1(t), \dots, \delta_n(t)]^T$ respectively denote the vector of speeds and angles of the machines. ω_R denotes the reference bus speed. $\boldsymbol{\eta}(t)$ denotes the OU noise processes that are correlated across the machines according to the correlation matrix $\mathbf{R} := \mathbf{C}\mathbf{C}^T$. Furthermore,

$$\mathbf{h}^{-1} = \begin{bmatrix} 1/h_1 \\ \vdots \\ 1/h_n \end{bmatrix}, \quad \mathbf{d} = \begin{bmatrix} d_1 \\ \vdots \\ d_n \end{bmatrix}, \quad \mathbf{p}^m = \begin{bmatrix} p_1^m \\ \vdots \\ p_n^m \end{bmatrix},$$

denote, respectively, the vector of inverse inertia coefficients, constant damping factors, and constant equilibrium total power injections. Finally

$$\mathbf{G} := [g_{ij}], \quad \mathbf{B} := [b_{ij}],$$

are the constant real and imaginary parts of the admittance matrix, respectively. We define

$$\cos(\Delta) := \begin{bmatrix} \cos(\Delta_{11}) & \cos(\Delta_{12}) & \cdots & \cos(\Delta_{1n}) \\ \cos(\Delta_{21}) & \cos(\Delta_{22}) & \cdots & \cos(\Delta_{2n}) \\ \vdots & \cdots & \ddots & \vdots \\ \cos(\Delta_{n1}) & \cos(\Delta_{n2}) & \cdots & \cos(\Delta_{nn}), \end{bmatrix}$$

to be the result of the cosine function applied elementwise on Δ , and $\sin(\Delta)$ is similarly defined, where

$$\Delta_{ij}(t) := \delta_i(t) - \delta_j(t)$$

denotes the pairwise angle difference. Letting $\mathbf{z}(t) := [\mathbf{v}(t)^T, \boldsymbol{\omega}(t)^T, \boldsymbol{\delta}(t)^T, \boldsymbol{\eta}(t)^T]^T$, the effective dimension of the system (4.17) is $N = 4n$.

4.4.1 Density of Line Energy

In the study of reliable power systems, stability of a transmission network can be quantified and studied through a closed-form energy function that is dependent on the system states of the dynamic model [231]. Subsequently, line failure events are typically modeled as random and rare events that occur in a network during operation. Upon stochastic perturbation, the energy function associated with a particular line may exceed a certain user-specified safety threshold, upon which a relay action is triggered, causing the line to be removed from graph $(\mathcal{V}, \mathcal{E})$. A modification of the admittance matrices reflects such removals, and in turn effectively changes the system (4.17). The post-relay dynamics continues to be monitored until the system states transition and settle in another equilibrium [231].

As a step toward efficient prediction of failure rates using a stochastic power system model, it is of interest to characterize the probability density of line energy functions. Specifically, assuming the baseline model (4.17), let $l \in \mathcal{E}$ denote an edge corresponding

to generator bus indices (i, j) . We define the quantity of interest as the magnitude of current flowing through line l as follows:

$$u(\mathbf{z}(t)) := \frac{1}{2}b_{ij}^2[(v_i(t))^2 - 2v_i(t)v_j(t)\cos(\delta_i(t) - \delta_j(t)) + (v_j(t))^2]. \quad (4.22)$$

Following the framework derived in Section 4.3.1, we can compute the following response variables explicitly, that is,

$$\mu^u := b_{ij}^2 v_i(t)v_j(t)\sin(\delta_i(t) - \delta_j(t))(\omega_i(t) - \omega_j(t)), \quad (4.23)$$

with $\mathcal{D}^u \equiv 0$ due to cancellation of the second-order terms in (4.7). The precise derivations of the response variables in (4.23) are provided in Section 4.7.1 of the Appendix. Upon substituting (4.23) into the reduced-order PDF equation (4.7), we obtain a 1D advection equation

$$\frac{\partial f_u}{\partial t} + b_{ij}^2 \frac{\partial}{\partial U}(\mathcal{R}(U, t)f_u) = 0, \quad (4.24)$$

where the advection coefficient has the following formally exact expression, according to equation (4.7), we have

$$\mathcal{R}(U, t) = \mathbb{E}[v_i v_j \sin(\delta_i - \delta_j)(\omega_i - \omega_j) | u(t) = U], \quad (4.25)$$

which is to be estimated from trajectory data depending on realizations of (4.22) via solving the minimization problem (4.15), as outlined in Algorithm 3.

4.4.2 Joint Dynamics of Line Pairs

When a line fails, its voltage load is redistributed to adjacent generators, causing an unexpected spiking of voltages in the local topology that may potentially cause expansive disruptions. For this reason, conditional failure events of adjacent machines are of great

practical interest [182]. We further consider the joint probability of two lines simultaneously exceeding their respective safety thresholds, which incorporates the conditional probabilistic structure.

Let two non-islanding lines $l_1, l_2 \in \mathcal{E}$ correspond to machines (i, j) and (j, k) with node j being an intermediate node from node i to node k . We follow the generalized formulation introduced in (4.13) and define the vector-valued quantity of interest, which is composed of the energy functions of both lines l_1, l_2 , yielding

$$\boldsymbol{\mu}_{\mathbf{u}} = \begin{bmatrix} b_{ij}^2 v_i v_j \sin(\delta_i - \delta_j)(\omega_i - \omega_j) \\ b_{jk}^2 v_j v_k \sin(\delta_j - \delta_k)(\omega_j - \omega_k) \end{bmatrix},$$

with no additional diffusion terms, similar to that of the 1D case; that is, $\mathcal{D}^{\mathbf{u}} \equiv \mathbf{0}$. In conjunction with the results of (4.13), the above derivation yields an advection equation in 2D as follows:

$$\frac{\partial f_{\mathbf{u}}}{\partial t} + b_{ij}^2 \frac{\partial}{\partial U_1} (\mathcal{R}_1(\mathbf{U}, t) f_{\mathbf{u}}) + b_{jk}^2 \frac{\partial}{\partial U_2} (\mathcal{R}_2(\mathbf{U}, t) f_{\mathbf{u}}) = 0, \quad (4.26)$$

where the phase vector $\mathbf{U} = [U_1, U_2]$, and the advection coefficients are now bivariate functions, defined as the following:

$$\mathcal{R}_1(\mathbf{U}, t) = \mathbb{E}[v_i v_j \sin(\delta_i - \delta_j)(\omega_i - \omega_j) | u_1(t) = U_1, u_2(t) = U_2], \quad (4.27)$$

$$\mathcal{R}_2(\mathbf{U}, t) = \mathbb{E}[v_j v_k \sin(\delta_j - \delta_k)(\omega_j - \omega_k) | u_1(t) = U_1, u_2(t) = U_2]. \quad (4.28)$$

The conditional expectations (4.25), (4.27), and (4.28) will be estimated jointly by regression, and the resulting advection-diffusion equations (4.24) in 1D and (4.26) in 2D will be computed numerically, thus enabling propagation of uncertainty.

4.5 Numerical Experiments

4.5.1 Experimental Setup

The experiments were conducted using three test cases: the WSCC 9-bus system, IEEE 30-bus system, and IEEE 57-bus system. To mimic the instability caused by sudden failing of a transmission line during normal operation, we first initialize the power system states with the connections intact. Namely, we disregard the noise model in (4.17) and solve the deterministic optimal power flow (OPF) problem for an equilibrium voltage magnitude $\mathbf{v}^* = [v_1^*, \dots, v_n^*]^T$ and machine angles $\boldsymbol{\delta} = [\delta_1^*, \dots, \delta_n^*]^T$ using Matpower 7.1 [234]. After obtaining equilibrium voltages and angles, we add a small random perturbation in the voltage magnitudes and treat the angles along with prespecified initial speeds at the reference level $\omega_R = 1$, as initial conditions for the random dynamical system (4.16). Specifically, the initial conditions for (4.17) are taken to be

$$\mathbf{v}(0) = |\mathcal{N}(\mathbf{v}^*, 0.01 \cdot \sigma_{\mathbf{v}^*}^2 \mathbf{I}_n)|$$

$$\boldsymbol{\omega}(0) = \omega_R \cdot \mathbf{1}$$

$$\boldsymbol{\delta}(0) = \boldsymbol{\delta}^*$$

$$\boldsymbol{\eta}(0) \sim \mathcal{N}(\mathbf{0}, \alpha^2 \mathbf{R}),$$

where the initial random voltage magnitudes are generated from a folded Gaussian distribution with mean at the equilibrium levels $\mathbf{v}^*$, and perturbed around the optimal power flow equilibrium with standard deviation $0.1 \cdot \sigma_{\mathbf{v}^*}$. In particular, $\sigma_{\mathbf{v}^*}$ is the standard deviation of equilibrium voltages across machines. Following [144], we take $\alpha = 0.05$, and the correlation matrix $\mathbf{R}$ is chosen for simplicity to be constant with moderate correlation levels, varying by application. We use $R_{ij} = R_{ji} = 0.44$ for the WSCC 9-bus case, for all $i, j \in \mathcal{V}, i \neq j$, and $R_{ij} = R_{ji} = 0.36$ in the IEEE 30-bus and 57-bus test cases.

With the noise model set, we first perform a burn-in period lasting time $T = 50$ and observe that the system (4.17) settles in an approximate (stochastic) equilibrium. We then simulate a single line failure by manually setting entries of the admittance matrices $\mathbf{B}, \mathbf{G}$ to zero. The samples at final time of the burn-in period are considered to be at equilibrium and again taken as initial conditions to be re-solved after line removal, for an additional time period $T = 10$. To investigate the energy changes caused to adjacent lines in the local topology, we compute and record the trajectories of energy function $\mathbf{u}$, defined in (4.22).

In our experiments we specifically choose non-islanding edges in the network and observe the adjacent lines connected to a shared generator. Concretely, lines 8–9 were removed for the WSCC 9-bus case, and energy of lines 4–9, 7–8 were subsequently recorded. Lines 6–8 were removed from IEEE case 30, and energy data was collected for lines 6–7, 6–9. For IEEE case 57, lines 36–37 were removed, and data was collected at lines 35–36 and 36–40. The system (4.17) was integrated by using a Milstein scheme [112] with a uniform time step size of $\Delta t = 10^{-2}$. In particular, we investigate the method’s ability to predict probabilities of random events defined with respect to the PDF. The reported 1D and 2D experiments were compared with a benchmark computed from a Gaussian kernel density estimator from $m_{\text{KDE}} = 10^4$ Monte Carlo trials, using Silverman’s rule for bandwidth selection. When reporting probabilities, we use instead the empirical cumulative distribution function, which represents a probability estimate from Monte Carlo sampling, a common approach in conventional studies. Across all test cases, the Monte Carlo sample size used for computing the initial condition, as well as the coefficient estimation (4.15), is taken to be $m_{\text{R}} = 5 \times 10^3$, for reported experiments of Section 4.5.2 and Section 4.5.3. As a quantitative measure, we define and henceforth refer to error as the space-time L^1 -norm distance from the benchmark

$$e(\hat{f}, f) \equiv \int_0^T \int_{\Omega} |\hat{f}(\mathbf{U}, t) - f(\mathbf{U}, t)| d\mathbf{U} dt, \quad (4.29)$$

where, for notation consistency, $\hat{f}$ is understood to be the result of the reduced-order PDF

method (4.13) and f is the KDE benchmark. The computational domain $\Omega \subset \mathbb{R}^{N_R}$ is determined from ranges of simulated values of $\mathbf{u}(t)$, with a padding on the boundaries of ± 0.5 –1 standard deviations. With the numerical domain set, we approximate the vanishing boundary conditions at infinities, as required by (4.4), with a Dirichlet boundary condition. Furthermore, to compare and evaluate the accuracy of probability estimations of a predefined event $E \subset \Omega$, we also report the probabilities $\hat{\mathbb{P}}(E), \mathbb{P}(E)$, where by construction $\hat{\mathbb{P}}(E) = \int_E \hat{f}$ is a prediction output from the reduced-order PDF and $\mathbb{P}$ is the benchmark result computed from empirical CDF. In Section 4.5.4 and Section 4.5.5 we examine empirically the sample scaling and correlation structure, respectively.

4.5.2 Marginal Probability Densities

To validate our proposed method, we first solve the 1D reduced-order PDF equations (4.26). The reduced-order equations marginal in selected line energies, as outlined in Algorithm 3, were solved by using a standard finite volume Lax–Wendroff method with a Monte Carlo flux limiter [127]. To resolve the unknown terms in (4.25), we used Gaussian local linear regression (GLLR) with 10-fold cross validations. The choice of regression method is dependent on the underlying system dynamics at hand and, naturally, requires a degree of manual tuning. The analysis of effects of regression methods is an open problem and can be studied via systematic model validation [227].

We show in Figure 4.1 the predicted PDF computed from the reduced-order equation by numerical integration, plotted at “peak time” where the energy level (averaged over all $m_R = 5 \times 10^3$ realizations) for each line reaches its maximum during simulation time. The predicted PDF is compared in the same plot with KDE using $m_{\text{KDE}} = 10^4$ samples. As a potential use case, we focus on the tail region defined by the exceedance of long-term line rating (specified in `Matpower` as field `rateA`). Specifically, we consider the following




Figure 4.1: (Top to bottom) Predicted PDF of line energies for WSCC case 9, IEEE case 30, and IEEE case 57, plotted at respective peak times, with exceedance thresholds zoomed in and marked in red. The 1D reduced-order marginal PDF equations (4.11) were solved with $m_R = 5 \times 10^3$ samples, compared with a $m_{KDE} = 10^4$ KDE benchmark.




Figure 4.2: Predicted reduced-order joint PDFs (4.26) with $m_R = 5 \times 10^3$ samples, at selected PDE solution times. Particularly, IEEE Case 30 and 57 show non-trivial evolution and correlation structure over time.

time-parameterized event:

$$E_l(t) \equiv \{u_l(t) > u_{l,\max}\}, \quad (4.30)$$

where $u_{l,\max}$ is understood as the individual line rating for line l . The probability of event E_l can be readily computed through numerical integration of the predicted PDF $\hat{f}$, which we report in Table 4.1. The predicted tail event probabilities were compared with those provided by an empirical CDF evaluation, which corresponds to the conventional Monte Carlo integral computed directly using samples of $u_l(t)$. In the investigated lines, lines 4–9, 7–8 of WSCC case 9 have a prespecified rating of 1 p.u. Lines 6–7, 6–9 in the IEEE case 30 have, respectively, ratings of 1.3 and 0.65 p.u.. For the IEEE case 57, the line ratings were not provided by **Matpower** but instead were estimated from solving the OPF with the Julia `PowerModels.jl` library using the `_calc_thermal_limits` function. The obtained ratings for lines 35–36, 36–40 were 16.33 and 20.27 p.u., respectively. From both Figure 4.1 and Table 4.1 we observe good agreement between our method and KDE estimation (the latter using double the number of samples as the former).

4.5.3 Joint Failure Probabilities

Following the discussions in Section 4.4.2, we solved the 2D reduced-order PDF equations joint in lines l_1, l_2 using a corner transport upwind method with second-order corrections and a van Leer flux limiter [127]. With the same computational domain setup as the 1D cases in Section 4.5.2 to approximate vanishing boundary, we applied both ordinary linear regression and locally weighted linear regression (Lowess) at each step of PDE integration to obtain the regression functions (4.27) from observations of $\mathbf{u}(t)$, with the latter being much more computationally demanding. The reduced-order equations were solved by using $m_R = 5 \times 10^3$ Monte Carlo samples and were compared with a benchmark computed from bivariate Gaussian KDE with $m_{\text{KDE}} = 10^4$ Monte Carlo samples. Figure 4.2 visualizes selected temporal snapshots of the predicted PDFs. In particular, we observe that the joint

PDF of WSCC case 9 is relatively stationary over time, whereas that of IEEE case 30 and case 57 show a high degree of variation.

To evaluate the joint probabilistic profile through a tail event, we let $U_{l_1,\max}, U_{l_2,\max}$ denote the respective constant line ratings, and we consider the region where either of the two lines fail, i.e., a pure series failure criterion:

$$E_{l_1,l_2}(t) \equiv \{u_{l_1}(t) > U_{l_1,\max} \cup u_{l_2}(t) > U_{l_2,\max}, \} \quad (4.31)$$

which can be evaluated from the predicted joint PDF $\hat{f}_{l_1 l_2}$ from Equation (4.26), by integrating the following numerically:

$$\begin{aligned} \hat{\mathbb{P}}(E_{l_1,l_2}) &= 1 - \hat{\mathbb{P}}(\{u_{l_1} \leq U_{l_1,\max} \cap u_{l_2} \leq U_{l_2,\max}\}) \\ &= 1 - \int_{-\infty}^{U_{l_1,\max}} \int_{-\infty}^{U_{l_2,\max}} \hat{f}_{l_1 l_2}(U_1, U_2, t) dU_1 dU_2. \end{aligned}$$

The predicted tail probabilities of events (4.31) for all test cases were reported in Table 4.1, in comparison with values computed from numerically integrating the benchmark KDE. Similarly to the marginal probabilities, we record the probabilities at the peak time at the larger of the two line energies.

We make a case for the consideration of joint events, rather than assuming the independence of failures. We highlight the deviating results (compared with the benchmark) in Table 4.1 computed from a product of independent marginal probabilities. Namely, the assumption

$$\hat{\mathbb{P}}(\{u_{l_1} \leq U_{l_1,\max} \cap u_{l_2} \leq U_{l_2,\max}\}) = \hat{\mathbb{P}}(\{u_{l_1} \leq U_{l_1,\max}\}) \cdot \hat{\mathbb{P}}(\{u_{l_2} \leq U_{l_2,\max}\}), \quad (4.32)$$

does not hold in general and yields inaccurate representations of failure probabilities, particularly for the cases 30 and 57. This is relevant since the approach used in [182] would

be unable to treat the occurrence of two failure events as dependent, and thus the approach here, while more expensive, presents a significant modeling advantage.

Case 9 ($t = 2.45$) 1×10^{-3}			Case 30 ($t = 0.56$) 1×10^{-2}			Case 57 ($t = 2.30$) 1×10^{-2}		
Lines	Pred.	Emp.	Lines	Pred.	Emp.	Lines	Pred.	Emp.
(4,9)	1.671	1.400	(6,7)	3.625	3.270	(35,36)	2.968	2.760
(7,8)	1.496	1.400	(6,9)	0.058	0.060	(36,40)	2.424	1.970
Joint	0.658/0.659	0.458	Joint	5.580/2.110	2.112	Joint	2.056/0.918	1.029

Table 4.1: Tail probabilities at energy-peak times for individual exceedances (4.30) and joint exceedances (4.31). Values from reduced-order PDF equations use $m_R = 5 \times 10^3$ and are compared against a KDE benchmark with $m_{\text{KDE}} = 10^4$. Grey numbers correspond to the independence assumption (4.32).




Figure 4.3: Ratio $\hat{\mathbb{P}}/\mathbb{P}$ between predicted and benchmark probabilities for the three cases (left to right). Grey markers correspond to the independence assumption; see Section 4.5.1.

4.5.4 Empirical Sample Complexity

For each test case, we investigate the reduced-order PDF method's practical scalability through parallelizing the computation over all line energies and counting the total number of samples needed for sufficient accuracy. In precise terms, we define the sample complexity measure for each line l , that is,

$$m_l^*(\gamma) := \min\{m : e(\hat{f}_l(\cdot; m), f_{l,\text{bench}}) < \gamma\},$$

where we let $\hat{f}_l(\cdot; m)$ denote the reduced-order PDF solution computed with m Monte Carlo samples, to emphasize consideration of sample size. $f_{l,\text{bench}}$ is a benchmark solution com-

puted from a high-resolution KDE using $2^{15} = 327,68$ samples. m_l^* thus quantifies the minimum number of samples required for the predicted PDF of line l to converge within $0 < \gamma < 1$ in L^1 -norm (4.29) of the benchmark KDE. Proceeding, we consider the aggregate sample complexity over all possible lines,

$$m^*(\gamma) := \sum_{l=1}^{|\mathcal{E}|} m_l^*(\gamma), \quad (4.33)$$

which is a case-dependent quantity due to the increasing number of connections in the power system. Without introducing line tripping, we simulate the system (4.17) in stochastic equilibrium for $T = 2$ with a refined time step size of $\Delta t = 5 \times 10^{-3}$ and with the same noise parameters as in Section 4.5.1. Then, sample sizes are increased from 2^{12} to 2^{14} in the computations of reduced-order PDF equations (i.e., using linear regression for closure estimation (4.25)), described in equation (4.24). As a comparison to the reduced-order PDF approach, a low-fidelity (in contrast to the benchmark) Gaussian KDE computed using the same number of samples was recorded for each example. The sequence of increasing sample sizes yield reduced-order PDF solutions and low-fidelity KDE of improved cumulative L^1 -accuracy in the sense of (4.29), when compared to the benchmark KDE. Finally, specifying the error threshold γ , we count and aggregate the number of samples required for each line, in order to obtain a final estimate of total sample complexity (4.33). This procedure for investigating sample scalability is summarized in Figure 4.4, where we observe a reduction of sample counts for an accuracy requirement of 1%. The dependence of sample complexity (4.33) on the number of lines $|\mathcal{E}|$ for the reduced-order PDF method has an estimated scaling of $\mathcal{O}(|\mathcal{E}|^{0.89})$, compared with $\mathcal{O}(|\mathcal{E}|^{1.07})$ for Gaussian KDE. While further sample complexity reduction has been noted in, for instance [144], where regression functions (4.15) admit partial separation (i.e., explicit dependence on QoI $\mathbf{u}(t)$ can be factored out due to conditioning and need not be estimated by regression) and thus variance reduction, the re-




Figure 4.4: (Top) Convergence of L^1 -norm error from benchmark (4.29), averaged by the number of lines for each case. (Bottom) Aggregate number of samples required for KDE or reduced-order PDF method to reach within $\gamma = 1\%$ L^1 -norm error from the benchmark.

duction in our case is particularly noteworthy because the response variable (4.24) is not an elementary expression of the line energy (4.22). In any case, the performance improvement of our method relative to direct simulation is significant, as we see from Figure 4.4 that the accuracy of our method with about 2^{10} samples can be achieved only at about 2^{14} samples with direct Monte Carlo. Since KDE does not naturally impose any time-dependent structure, the empirical result suggests the importance and plausibility of investigating scalable reduced-order models of time-dependent systems.

4.5.5 Mutual Information

Information transmission is studied in Bayesian reliability analysis to identify fragile components of a power network under system perturbations [232, 110]. As an alternative quantification of the correlation structure for the measured lines under load uncertainty, we consider a discrepancy measure between the predicted 2D joint PDF, and the product of marginal PDFs, which represents the density obtained from the independence assumption in Equation (4.32). The Kullback–Leibler (KL) divergence, denoted as D_{KL} , is used to assess the

strength of interdependency between pairs of random variables, also known as their mutual information. Total correlation among a larger subset of random states can also be conveniently generalized in graphical models that incorporate the network topology directly; see, for instance [197]. For our application, we define in 2D the approximate KL divergence for the energy distributions of lines $l_1, l_2 \in \mathcal{E}$,

$$D_{\text{KL}}(\hat{f}_{l_1 l_2} \| \hat{f}_{l_1} \otimes \hat{f}_{l_2}; t) := \int \int \hat{f}_{l_1 l_2}(U_1, U_2, t) \log \left(\frac{\hat{f}_{l_1 l_2}(U_1, U_2, t)}{\hat{f}_{l_1}(U_1, t) \hat{f}_{l_2}(U_2, t)} \right) dU_1 dU_2, \quad (4.34)$$

where the density $\hat{f}_{l_1} \otimes \hat{f}_{l_2}$ denotes the product probability measure from 1D reduced-order PDFs (4.23). Figure 4.5 shows the plotted curves for all cases predicted using $m_R = 5,000$ observations. In particular, Figure 4.5 suggests that the assessed lines in WSCC case 9 and IEEE case 30 are on average only weakly correlated; thus, estimating marginal densities separately may provide a relatively faithful approximation (though at the peak time such discrepancy did manifest more prominently for IEEE case 30, as shown in Table 4.1). The correlation is non-negligible for IEEE case 57, and a product measure approximation does not yield accurate results, as also seen in Table 4.1.

4.6 Conclusions and Future Work

In this work, we investigated the reduced-order PDF method to reformulate the propagation of uncertainty in a power system model as a low-dimensional numerical PDE problem with closure terms that can be modeled empirically by regression. We proposed a general framework that can flexibly incorporate correlated noise models. Furthermore, we provided an implementation and demonstrated the method's ability for describing density of arbitrary QoIs with a considerable reduction of high-dimensional MC simulations. The method is shown to accurately propagate failure probabilities. Most importantly, we demonstrate nontrivial correlation structure particularly for IEEE test cases of increasing sizes, measured




Figure 4.5: Estimated mutual information (4.34) of WSCC case 9 and IEEE cases 30 and 57 over PDE simulation time up to $T = 10$.

through the mutual information (4.34), validating the need for and advantage in considering joint densities. Although we have discussed the probabilistic profile of only two lines, the equation for several lines ($N_R \geq 3$) can be analogously derived. For future work, such a framework holds promise in the prediction of cascading failure events in more complex energy-based power system models involving line tripping dynamics, such as [231].

4.7 Appendix

4.7.1 Derivations of Reduced-Order PDF Coefficients

To make the form of the reduced-order PDF equations considered in Section 4.4.1 and Section 4.4.2 more explicit, we provide a concise computation of the advection and diffusion coefficients that appear in (4.23), following the formulae introduced in (4.9) and (4.10). Because of the elementary extension to two transmission lines, derivations of the associated coefficients in the joint equation (4.26) follow exactly.

We consider line $l \in \mathcal{E}$ and generator indices (i, j) fixed. Furthermore, we recall the

energy function defined in (4.22),

$$u(\mathbf{z}(t)) := \frac{1}{2}b_{ij}^2[(v_i(t))^2 - 2v_i(t)v_j(t)\cos(\delta_i(t) - \delta_j(t)) + (v_j(t))^2].$$

which is a function of the states evolving according to the noise-perturbed system defined in (4.17). Effectively, only a small subset of variables in $\mathbf{z}$ contributes to line l 's energy; that is, $u(\mathbf{z}) = u(v_i, v_j, \delta_i, \delta_j)$. As a result, the entries in the gradient, $\nabla_{\mathbf{z}}u$, and Hessian matrix, $H_{\mathbf{z}}u$, of the line energy vanish except the following contributions of voltage:

$$\begin{aligned}\frac{\partial u}{\partial v_i} &= -\frac{\partial u}{\partial v_j} = b_{ij}^2(v_i - v_j \cos(\delta_i - \delta_j)) \\ \frac{\partial^2 u}{\partial v_i^2} &= \frac{\partial^2 u}{\partial v_j^2} = b_{ij}^2 \\ \frac{\partial^2 u}{\partial v_i \partial v_j} &= -b_{ij}^2 \cos(\delta_i - \delta_j).\end{aligned}$$

And similarly for generator angles,

$$\begin{aligned}\frac{\partial u}{\partial \delta_i} &= -\frac{\partial u}{\partial \delta_j} = b_{ij}^2 v_i v_j \sin(\delta_i - \delta_j) \\ \frac{\partial^2 u}{\partial \delta_i^2} &= \frac{\partial^2 u}{\partial \delta_j^2} = -\frac{\partial^2 u}{\partial \delta_i \partial \delta_j} = b_{ij}^2 v_i v_j \cos(\delta_i - \delta_j).\end{aligned}$$

In vector form, we obtain more concisely the following expression:

$$\nabla_{\mathbf{z}}u = \begin{bmatrix} \nabla_{\mathbf{v}}u \\ \mathbf{0} \\ \nabla_{\boldsymbol{\delta}}u \\ \mathbf{0} \end{bmatrix}, H_{\mathbf{z}}u = \begin{bmatrix} H_{\mathbf{v}}u & \mathbf{0} & H_{\mathbf{v}\boldsymbol{\delta}}u & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ H_{\boldsymbol{\delta}\mathbf{v}}u & \mathbf{0} & H_{\boldsymbol{\delta}}u & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix},$$

where $\mathbf{0}$ denotes the vector/matrix of all 0's with appropriate dimensions. The notation $H_{\mathbf{z}_1\mathbf{z}_2}$ is understood as the matrix of mixed (second-order) derivatives involving compo-

nents of arbitrary vectors $\mathbf{z}_1, \mathbf{z}_2$. From the vector form of the stochastic differential equations (4.17), we have

$$\boldsymbol{\mu}_t = \begin{bmatrix} \mathbf{0} \\ \dot{\boldsymbol{\omega}}(t) \\ \dot{\boldsymbol{\delta}}(t) \\ \dot{\boldsymbol{\eta}}(t) - \alpha\sqrt{2\theta}\mathbf{C} \cdot \frac{d\mathbf{W}(t)}{dt} \end{bmatrix}, \quad \boldsymbol{\sigma}_t = \begin{bmatrix} \mathbf{0} & & & \\ & \mathbf{0} & & \\ & & \mathbf{0} & \\ & & & \alpha\sqrt{2\theta} \cdot \mathbf{C} \end{bmatrix}.$$

in correspondence with the notations introduced in (4.7). Consequently, we can readily compute the quantities defined formally in (4.9) and (4.10) as the following:

$$\begin{aligned} (\nabla_{\mathbf{z}u})^T \boldsymbol{\mu}_t &= (\nabla_{\boldsymbol{\delta}u})^T d\boldsymbol{\delta}(t) \\ &= b_{ij}^2 v_i v_j \sin(\delta_i - \delta_j)(w_i - w_R) - b_{ij}^2 v_i v_j \sin(\delta_i - \delta_j)(w_j - w_R) \\ &= b_{ij}^2 v_i v_j \sin(\delta_i - \delta_j)(\omega_i - \omega_j), \end{aligned}$$

and since $\boldsymbol{\sigma}_t^T (H\mathbf{z}u)\boldsymbol{\sigma}_t = \mathbf{0}$, $(\nabla_{\mathbf{z}u})^T \boldsymbol{\sigma}_t = \mathbf{0}$, we obtain

$$\begin{aligned} \mu^u &= (\nabla_{\mathbf{z}u})^T \boldsymbol{\mu}_t + \frac{1}{2} \text{tr}(\boldsymbol{\sigma}_t^T (H\mathbf{z}u)\boldsymbol{\sigma}_t) = b_{ij}^2 v_i v_j \sin(\delta_i - \delta_j)(\omega_i - \omega_j) \\ \mathcal{D}^u &= (\nabla_{\mathbf{z}u})^T \boldsymbol{\sigma}_t ((\nabla_{\mathbf{z}u})^T \boldsymbol{\sigma}_t)^T = 0. \end{aligned}$$

Here, the (stochastic) differentials $d\boldsymbol{\omega}(t), d\boldsymbol{\delta}(t), d\boldsymbol{\eta}(t)$ are defined through rewriting the system (4.17) in the Itô interpretation, consistent with the general definition in (4.3). The above derivations extend trivially to the multiple-lines case (4.13).

CHAPTER 5

OVERLAPPING SCHWARZ SCHEME FOR LINEAR-QUADRATIC PROGRAMS IN CONTINUOUS-TIME

5.1 Introduction

Optimal control problems (OCPs) constrained by differential equations are fundamental in a wide range of scientific and engineering applications, including fluid dynamics, biomedical engineering, and aerospace design [205, 81, 203]. These problems require determining an optimal control function that minimizes a given cost functional while satisfying the governing differential constraints. The numerical formulation of OCPs traditionally follows one of the two major paradigms: the direct “discretize-then-optimize” approach, and the indirect “optimize-then-discretize” approach [95, 9].

The direct approach first applies numerical discretization techniques, such as finite element [32], finite difference [128], or spectral methods [190], to approximate the governing equations. This transforms the original infinite-dimensional problem in the space of solution functions into a finite-dimensional constrained optimization problem, where the control and state variables are indexed by a discrete set of design variables. Then, to solve the resulting nonlinear program, one can employ constrained optimization techniques, including interior-point methods [213] and sequential quadratic programming (SQP) methods [72], which enforce constraints on the state and control variables jointly. Recent advancements in direct methods have focused on improving numerical efficiency and stability. For example, reduced-space Newton-Krylov solvers [93, 94] offer iterative techniques to compute the reduced Karush-Kuhn-Tucker (KKT) matrix-vector products efficiently, while penalty-barrier methods [158] introduce regularization terms on the primal-dual problem to improve stability. Despite these improvements, direct methods remain inherently dependent on a fully discretized formulation of the OCP, making them increasingly computationally intensive as

the discretization level grows.

In contrast, the indirect approach derives first-order optimality conditions in the continuous setting before applying numerical discretization. This formulation yields necessary optimality conditions in the form of a Hamiltonian system [45]. A key component of the indirect method is the adjoint equation, which enables efficient gradient computation with respect to the control variables by solving an auxiliary system of equations [61], thereby eliminating the need for explicit numerical differentiation with respect to the discretized design variables. To solve the resulting Hamiltonian system, state-of-the-art methods employ numerical techniques that preserve the structure of the continuous problem. For example, shooting methods [19] solve boundary-value problems by iteratively adjusting the control variables to satisfy the necessary conditions for optimality; and matrix-free adjoint techniques [85] alleviate memory cost by avoiding the explicit storage of large Jacobian matrices, making them effective for numerically solving large-scale partial differential equations (PDEs). By preserving the problem’s continuous formulation, optimize-then-discretize approaches offer greater flexibility in handling varying spatial and temporal resolutions, making them particularly advantageous in high-fidelity simulations.

A fundamental challenge in differential equation-constrained OCPs is scalability, as the significant number of design variables and constraints make classic *centralized* approaches computationally prohibitive for high-dimensional problems [51]. The optimization often struggles with memory limitations and long solver runtime when they are applied to large-scale problems. Since centralized methods rely on a single processor, their performance degrades rapidly in resource-constrained computing environments. To alleviate these computational bottlenecks, domain decomposition methods have been developed to divide the computational domain into smaller (and potentially overlapping) subdomains, enabling parallel workload distribution across multiple processors [57, 170].

Schwarz alternating methods have long been regarded as effective preconditioners for

solving elliptic and parabolic PDEs, leveraging iterative subdomain solutions to accelerate convergence [137]. More recently, these methods have been adapted to PDE-constrained control problems, where allowing design variables to overlap across subdomains has been shown to improve convergence rates [192, 152]. A key theoretical insight is that *perturbations at the primal-dual subdomain boundaries exhibit an exponential decay phenomenon as they propagate into the subdomain*. This property serves as the foundation for iterative Schwarz-based optimization schemes, ensuring that the concatenation of overlapping subdomain solutions ultimately converge to the full, global optimum under appropriate boundary coordinations [152]. Building on this insight, recent work has refined the local design of the Schwarz method to a globally convergent fast overlapping temporal decomposition (FOTD) method [153]. While these developments have demonstrated the effectiveness of Schwarz-based decomposition in structured settings, existing approaches have purely focused on discrete OCP formulations and relied on explicit domain partitioning.

In this work, we take the first step toward developing an overlapping Schwarz decomposition framework for general nonlinear PDE-constrained OCPs by studying the linear-quadratic setting. While our method can readily accommodate arbitrary spatial and temporal variables that arise in a PDE, we focus solely on temporal variables and concentrate our efforts on decomposing the temporal domain for clarity. This simplification allows us to clearly distinguish the continuous-time Schwarz method from aforementioned discrete-time counterparts. In particular, our method follows the indirect “optimize-then-discretize” paradigm and leverages the Pontryagin Minimum Principle (PMP) to partition the time domain into overlapping subdomains, each associated with a local OCP formulated as a Hamiltonian two-point boundary value problem in continuous-time. At each Schwarz iteration, we specify the boundary values based on the previous iteration; solve local OCPs in parallel; then concatenate the resulting local solutions to form a global trajectory. The method iterates until a stopping criterion is satisfied, ensuring convergence to the global optimum.

Performing decomposition directly on continuous-time problems offers several advantages: (i) it preserves the clean analytical structure of the optimality conditions in function space, ensuring consistency with the underlying control problem; (ii) it enables efficient computation of gradient information without requiring full discretization; and (iii) it facilitates the use of solvers with adaptive time-stepping, making it well-suited for stiff or multi-scale systems. In this work, we establish a convergence analysis for the continuous-time overlapping Schwarz method. We show that under standard controllability and coercivity conditions, *the method converges linearly with a rate that improves exponentially with the overlap size*. A key component of our analysis is to show that the exponential decay of sensitivity (EDS), previously observed in discrete-time OCPs, carries over to continuous-time settings. Our analysis draws close conceptual connections to the *turnpike property* in optimal control, which states that in long-horizon OCPs (without terminal cost), for any initial state, the optimal trajectories remain close to a steady-state trajectory for most of the horizon, deviating only near the boundaries. The fraction of time spent close to the steady-state is independent of horizon length [63, Theorem 2]. And, under suitable assumptions, the optimal trajectories approach exponentially fast. The pith of the exponential convergence behavior relies on the so-called *strict dissipativity / coercivity* of the cost functional [49, Definition 4.1, Theorem 5.6] as well as system controllability [77, Theorem 5.3], at least for the discrete setting. In our case, the construction of Schwarz subproblems requires an added terminal cost that embeds perturbations to the (unperturbed) global optimal control that contracts through updated solution parameters. In this setting, we derive explicit bounds that establish the exponential property—not with respect to steady-states, but rather global optimal trajectories without perturbation—under system controllability and coercivity of the cost functional. Numerical experiments validate our theory and demonstrate promising performance of our infinite-dimensional Schwarz framework, offering flexibility across numerical solvers and highlighting its efficiency and structure-preserving properties for general OCPs.

5.1.1 Organization

In Section 5.2, we introduce the linear-quadratic OCPs in infinite dimensions, along with the assumptions and preliminary results. In Section 5.3, we present the continuous-time overlapping Schwarz algorithm, which is derived in Section 5.8.1. Section 5.4 analyzes a generic parameterized linear-quadratic OCP and establishes the exponential decay of sensitivity result that serves as the basis of the convergence analysis of the Schwarz scheme. In Section 5.5, we prove the pointwise linear convergence of the Schwarz scheme and show that the linear rate improves exponentially with the overlap size. Numerical experiments are presented in Section 5.6, followed by conclusions and future directions in Section 5.7. We present the proofs of our key results in Section 5.8.2 and Section 5.8.3.

5.1.2 Notation

Throughout this chapter, we consider state and control trajectories $x : [0, T] \rightarrow \mathbb{R}^{n_x}$ and $u : [0, T] \rightarrow \mathbb{R}^{n_u}$ defined over the finite time interval $[0, T]$. We use $\|\cdot\|$ to denote the ℓ_2 norm for vectors and operator norm for matrices. For a symmetric matrix $A \in \mathbb{R}^{n \times n}$, $A \succ (\succeq) 0$ means that A is positive (semi-)definite. We use $\langle \cdot, \cdot \rangle$ to denote the norm-induced inner product in appropriate Banach spaces (including the Euclidean space $\mathbb{R}^n$). Let X be a Banach space and $\mathcal{J} : f \in X \rightarrow \mathbb{R}$ be a functional. We denote the total (Gateaux) derivative of $\mathcal{J}$ with respect to f by $\delta\mathcal{J}/\delta f$. When $\mathcal{J}$ depends on multiple functions, we use $\partial\mathcal{J}/\partial f$ to denote the partial derivative with respect to f . For further details on functional derivatives in infinite-dimensional spaces, we refer the reader to [224]. The notation for functional derivatives is used throughout to describe Algorithm 4, drawing an analogy with its finite-dimensional discrete counterpart.

5.2 Problem formulation

Consider the following linear-quadratic OCP of the Bolza-type objective [189]:

$$\min_{u(\cdot), x(\cdot)} \quad \frac{1}{2} \int_0^T \begin{bmatrix} x(t) \\ u(t) \end{bmatrix}^\top \begin{bmatrix} Q(t) & H^\top(t) \\ H(t) & R(t) \end{bmatrix} \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} dt + \frac{1}{2} x(T)^\top Q_T x(T), \quad (5.1a)$$

$$\text{s.t.} \quad \dot{x}(t) = A(t)x(t) + B(t)u(t), \quad t \in (0, T], \quad (5.1b)$$

$$x(0) = x_0, \quad (5.1c)$$

where $A : [0, T] \rightarrow \mathbb{R}^{n_x \times n_x}$, $B : [0, T] \rightarrow \mathbb{R}^{n_x \times n_u}$ are the time-varying system matrices; $Q : [0, T] \rightarrow \mathbb{R}^{n_x \times n_x}$, $H : [0, T] \rightarrow \mathbb{R}^{n_u \times n_x}$, $R : [0, T] \rightarrow \mathbb{R}^{n_u \times n_u}$ are the running cost matrices; $x_0 \in \mathbb{R}^{n_x}$ is the initial state; and $Q_T \in \mathbb{R}^{n_x \times n_x}$ defines the terminal cost.

For Problem (5.1), we impose the following regularity conditions. For any mapping $f : [0, T] \rightarrow \mathbb{R}^{p \times q}$, we let $\|f\|_\infty := \sup_{t \in [0, T]} \|f(t)\|$. When $\mathbb{R}^{p \times q}$ reduces to $\mathbb{R}^p$ (e.g., for $u(\cdot), x(\cdot)$), $\|\cdot\|$ is understood as the ℓ_2 norm.

Assumption 1. *We assume the time-varying matrices in the OCP (5.1) satisfy the following conditions:*

1. *The mappings $A : [0, T] \rightarrow \mathbb{R}^{n_x \times n_x}$ and $B : [0, T] \rightarrow \mathbb{R}^{n_x \times n_u}$ are continuous and uniformly bounded: $\|A\|_\infty \leq \lambda_A$, $\|B\|_\infty \leq \lambda_B$ for some constants $\lambda_A, \lambda_B > 0$ independent of T .*
2. *The mappings $Q : [0, T] \rightarrow \mathbb{R}^{n_x \times n_x}$, $H : [0, T] \rightarrow \mathbb{R}^{n_u \times n_x}$, and $R : [0, T] \rightarrow \mathbb{R}^{n_u \times n_u}$ are continuous and uniformly bounded: $\|Q\|_\infty \leq \lambda_Q$, $\|H\|_\infty \leq \lambda_H$, $\|R\|_\infty \leq \lambda_R$ for some constants $\lambda_Q, \lambda_H, \lambda_R > 0$ independent of T .*

We should mention that the uniformity of the boundedness, i.e., the independence of the constants from the terminal time T , is not required if we consider T to be fixed. Our analysis of the sensitivity of OCPs and the Schwarz method still holds without the uniformity;

however, this stronger boundedness condition will indicate that the convergence rate of the method, determined by the boundedness constants, is independent of T . This uniform result is particularly desirable for solving long-horizon OCPs (i.e. as T becomes large) [151, 152].

The regularity conditions in Assumption 1 ensure the existence and uniqueness of the *mild* state solution for any bounded control input [165, Chapter 5.1]. In particular, under Assumption 1, for any bounded control mapping $u^* : [0, T] \rightarrow \mathbb{R}^{n_u}$, we define the corresponding state trajectory $x^* : [0, T] \rightarrow \mathbb{R}^{n_x}$ as the continuous function

$$x^*(t) = \Phi_A(t, 0)x_0 + \int_0^t \Phi_A(t, s)B(s)u^*(s) ds, \quad t \in [0, T], \quad (5.2)$$

where $\Phi_A(t, s)$ is called the *linear evolution operator* associated with the homogeneous system $\dot{x}(t) = A(t)x(t)$, satisfying $x(t) = \Phi_A(t, s)x(s)$. The representation of (5.2) is commonly referred to as the *mild solution* of the state equation. The existence, uniqueness, and continuity of $\Phi_A(\cdot, \cdot)$ follow from classical results in semigroup theory under the continuity of $A(t)$; see [165, Theorems 5.1 and 5.2].

The mild solution (5.2) will serve as the foundation for analyzing how control inputs influence the state trajectory over time, which leads to the following notion of controllability.

Definition 1 (Complete Controllability) *A state $x_{t_0} \in \mathbb{R}^{n_x}$ is said to be controllable at time $t_0 \geq 0$ if there exist $t_1 > t_0$ and a control function $\bar{u} : [t_0, t_1] \rightarrow \mathbb{R}^{n_u}$, depending on t_0 and x_{t_0} , such that the mild solution of the state equation (5.1b) with the initialization $x(t_0) = x_{t_0}$ and the control input $\bar{u}$ satisfies $x(t_1) = 0$. That is,*

$$\Phi_A(t_1, t_0)x_{t_0} + \int_{t_0}^{t_1} \Phi_A(t_1, s)B(s)\bar{u}(s) ds = 0.$$

The state system (5.1b) is said to be completely controllable if every state x_{t_0} is controllable at all $t_0 \geq 0$.

We recall a fundamental result on controllability, which will be used frequently in our

later analysis.

Lemma 2 (Controllability Gramian). *The state system defined by the pair (A, B) in (5.1b) is completely controllable on an interval $[t_0, t_1] \subseteq [0, T]$ if and only if the symmetric matrix*

$$W_{A,B}(t_0, t_1) := \int_{t_0}^{t_1} \Phi_A(t_0, s) B(s) B^\top(s) \Phi_A^\top(t_0, s) ds \in \mathbb{R}^{n_x \times n_x} \quad (5.3)$$

is positive definite. The matrix $W_{A,B}$ is referred to as the controllability Gramian for the pair (A, B) . Furthermore, when $W_{A,B}(t_0, t_1)$ is nonsingular, an explicit control that steers any state x_{t_0} at time t_0 to the origin at time t_1 is given by

$$u(t; t_0, x_{t_0}) = \begin{cases} -B^\top(t) \Phi_A^\top(t_0, t) W_{A,B}^{-1}(t_0, t_1) x_{t_0}, & t \in [t_0, t_1], \\ 0, & t > t_1. \end{cases}$$

Proof. The results directly follow from [23, Proposition 5.2, (5.9), and (5.10)]. $\square$

While complete controllability ensures that the system state can be driven to the origin over some finite time horizon, our analysis requires a quantitative refinement on the time required to achieve controllability. Specifically, we impose explicit lower and upper bounds on both the controllability Gramian and the evolution operator associated with the linear dynamics.

Assumption 2 (Uniform Complete Controllability (UCC)). *We assume there exist a fixed constant $\sigma \in (0, T)$ and some positive constants $\alpha_0(\sigma), \alpha_1(\sigma), \beta_0(\sigma), \beta_1(\sigma) > 0$, possibly depending on σ , such that for every $t_0 \in [0, T - \sigma]$, there exists a time $t_1 = t_1(t_0) \in [t_0, t_0 + \sigma] \subseteq [0, T]$ for which the following bounds hold:*

$$\alpha_0(\sigma)I \preceq W_{A,B}(t_0, t_1) \preceq \alpha_1(\sigma)I \quad (5.4)$$

$$\beta_0(\sigma)I \preceq \Phi_A(t_1, t_0) W_{A,B}(t_0, t_1) \Phi_A^\top(t_1, t_0) \preceq \beta_1(\sigma)I. \quad (5.5)$$

Assumption 2 is introduced in [23, Definition 5.13] and is standard in the numerical control literature; see, for example, [152, Assumption 2] and [102, Definition 2]. It essentially states that the system, at any time t_0 and state x_{t_0} , can be controlled within an interval of length σ . Notably, we do not require controllability on the terminal segment $[T - \sigma, T]$ since we consider a finite-horizon problem.

To ensure the well-posedness of Problem (5.1), we impose the following coercivity conditions on the cost functional. We abuse the constant λ_Q from Assumption 1 for notational consistency.

Assumption 3. *There exist uniform constants $\gamma_R, \gamma_Q > 0$ independent of T such that*

$$R(t) \succeq \gamma_R I, \quad Q(t) - H^\top(t)R^{-1}(t)H(t) \succeq \gamma_Q I, \quad \forall t \in [0, T].$$

In addition, the terminal cost matrix satisfies $\gamma_Q I \preceq Q_T \preceq \lambda_Q I$ for some constants $\lambda_Q \geq \gamma_Q > 0$.

Assumption 3 guarantees the strong convexity of the cost functional, thereby ensuring the existence and uniqueness of solutions for OCP (5.1).

Similar to Assumption 1, the uniformity of the constants is not required when we treat T as a fixed constant, while the uniformity indicates that our results, including the convergence rate, are independent of T . We highlight that, under Assumption 1(b), Assumption 3 is equivalent to what is commonly assumed in the literature (e.g., [151, Assumption 3.1], [217, Assumption 2.1]), where a constant $\gamma' > 0$ independent of T is assumed such that

$$\begin{bmatrix} Q(t) & H^\top(t) \\ H(t) & R(t) \end{bmatrix} \succeq \gamma' I, \quad \forall t \in [0, T]. \quad (5.6)$$

In particular, the condition (5.6) implies that

$$R(t) \succeq \gamma' I, \quad \begin{bmatrix} Q(t) - \gamma' I & H^\top(t) \\ H(t) & R(t) \end{bmatrix} \succeq 0, \quad \forall t \in [0, T],$$

suggesting that Assumption 3 holds with $\gamma_R = \gamma_Q = \gamma'$ [225, Theorem 1.12]. On the other hand, Assumption 3 implies for any $\gamma' \leq \epsilon \gamma_R$ with $\epsilon \in (0, 1)$ that $R(t) - \gamma' I \succeq (1 - \epsilon)R(t)$. Thus, we have

$$\begin{aligned} Q(t) - H^\top(t)(R(t) - \gamma' I)^{-1}H(t) &\succeq Q(t) - H^\top(t)R(t)^{-1}H(t)/(1 - \epsilon) \\ &\succeq Q(t) - H^\top(t)R(t)^{-1}H(t) - \frac{\epsilon}{1 - \epsilon} \frac{\lambda_H^2}{\gamma_R} I \succeq \left(\gamma_Q - \frac{\epsilon}{1 - \epsilon} \frac{\lambda_H^2}{\gamma_R} \right) I. \end{aligned}$$

Thus, let ϵ be small enough such that $\gamma_Q - \frac{\epsilon}{1 - \epsilon} \frac{\lambda_H^2}{\gamma_R} \geq \epsilon \gamma_Q$, then (5.6) holds with $\gamma' = \epsilon \min\{\gamma_Q, \gamma_R\}$. Without Assumption 1, Assumption 3 is strictly weaker than (5.6) since one can see that $Q(t) = t^2 + 1$, $H(t) = t$, $R(t) = 1$ satisfies Assumption 3 but not (5.6) (the smallest eigenvalue of the quadratic matrix goes to zero as $t \rightarrow \infty$).

Given the above assumptions, we state the first-order optimality conditions for the linear-quadratic time-varying OCP (5.1). These conditions will be later used to truncate time intervals and formulate subproblems by specifying appropriate boundary conditions.

Theorem 4 (Pontryagin's Minimum Principle (PMP)) *Suppose Assumption 1 and Assumption 3 hold for the OCP in (5.1). Then, there is a unique, continuous optimal control solution $u^* : [0, T] \rightarrow \mathbb{R}^{n_u}$ along with a unique, continuously differentiable optimal state solution $x^* : [0, T] \rightarrow \mathbb{R}^{n_x}$ that solves (5.1). Furthermore, we define the Hamiltonian as*

$$\mathcal{H}(t, x(t), u(t), \lambda(t)) := \frac{1}{2} \begin{bmatrix} x(t) \\ u(t) \end{bmatrix}^\top \begin{bmatrix} Q(t) & H^\top(t) \\ H(t) & R(t) \end{bmatrix} \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} + \lambda(t)^\top (A(t)x(t) + B(t)u(t)). \quad (5.7)$$

Then, there exists a unique, continuously differentiable adjoint solution $\lambda^* : [0, T] \rightarrow \mathbb{R}^{n_x}$ such that

$$\dot{x}^*(t) = \nabla_x \mathcal{H}(t, x^*(t), u^*(t), \lambda^*(t)), \quad x^*(0) = x_0, \quad (5.8a)$$

$$\dot{\lambda}^*(t) = -\nabla_\lambda \mathcal{H}(t, x^*(t), u^*(t), \lambda^*(t)), \quad \lambda^*(T) = Q_T x^*(T), \quad (5.8b)$$

$$\mathcal{H}(t, x^*(t), u^*(t), \lambda^*(t)) \leq \mathcal{H}(t, x^*(t), u, \lambda^*(t)), \quad \forall u \in \mathbb{R}^{n_u}, \forall t \in [0, T]. \quad (5.8c)$$

Conversely, if the triple (x^*, u^*, λ^*) with x^*, λ^* being continuously differentiable and u^* being continuous satisfies the conditions (5.8), then (x^*, u^*) is the unique solution to OCP (5.1).

Proof. We refer the reader to the discussions of [6, Sections 2.3 and 3.4] for the existence and uniqueness of the optimal solution to the linear-quadratic OCP (5.1); also refer to [37, Theorems 3.9 and 3.11, and Remark 3.12] for the necessary and (Mangasarian) sufficient conditions of OCPs. $\square$

Theorem 4 generalizes the Karush-Kuhn-Tucker (KKT) conditions of nonlinear optimization problems to the setting of infinite-dimensional problems. In particular, Theorem 4 provides an *open-loop* characterization of optimality: for every t , the optimal control $u^*(t)$ is not expressed as a function of the state $x^*(t)$, but rather satisfies the minimization conditions derived from the Hamiltonian. This form is particularly useful in the design and analysis of approximation algorithms. In contrast, the *closed-loop* characterization, where the control is expressed explicitly as a function of the current state, is useful for studying stability. Such a formulation typically arises from the dynamic programming principle, which we now briefly review.

Theorem 5 (Hamilton-Jacobi-Bellman (HJB) Equation). *For each $t \in [0, T]$, we define the*

optimal cost-to-go function as

$$J^*(t, z) := \min_{u(\cdot), x(\cdot)} \frac{1}{2} \int_t^T \begin{bmatrix} x(s) \\ u(s) \end{bmatrix}^\top \begin{bmatrix} Q(s) & H^\top(s) \\ H(s) & R(s) \end{bmatrix} \begin{bmatrix} x(s) \\ u(s) \end{bmatrix} ds + \frac{1}{2} x(T)^\top Q_T x(T), \quad (5.9)$$

where the state trajectory $x(\cdot)$ satisfies the linear dynamics $\dot{x}(s) = A(s)x(s) + B(s)u(s)$, $s \in (t, T]$ and $x(t) = z$. Let the Hamiltonian be defined as in (5.7). Under Assumptions 1 and 3, $J^*(t, z)$ is the solution of the Hamilton-Jacobi-Bellman (HJB) partial differential equation:

$$\frac{\partial J^*}{\partial t}(t, z) + \min_{u \in \mathbb{R}^{n_u}} \mathcal{H}(t, z, u, \nabla_z J^*(t, z)) = 0, \quad (5.10a)$$

$$J^*(T, z) = \frac{1}{2} z^\top Q_T z. \quad (5.10b)$$

Proof. See [6, Sections 2.2 and 2.3] for the argument. We also refer the readers to the literature in optimal control theory, including [12, 18, 100, 203, 22]. $\square$

The upcoming section introduces the subproblem formulation used in our Schwarz decomposition, along with its theoretical justification, leveraging the optimality theorems presented above.

5.3 Continuous-Time Overlapping Schwarz Decomposition

In this section, we introduce a continuous-time overlapping Schwarz decomposition scheme for solving linear-quadratic OCPs. We decompose the time domain into overlapping subdomains and formulate a subproblem for each subdomain. The subproblem formulation depends on properly chosen boundary conditions. The Schwarz method enables parallel computation and achieves convergence to the global solution by iteratively updating these boundary conditions. Our subproblem formulation leverages the PMP in Theorem 4, and we

offer the flexibility to incorporate different numerical solvers for solving the continuous-time subproblems.

We begin by partitioning the time domain $[0, T]$ into m subdomains using a sequence of time points:

$$0 = t_0 < t_1 < \cdots < t_m = T, \quad \text{so that} \quad [0, T] = \bigcup_{j=1}^m [t_{j-1}, t_j].$$

We then define the overlap parameters (τ_j^0, τ_j^1) for each subdomain j , where $\tau_j^0 > 0$ denotes the backward overlap size and $\tau_j^1 > 0$ the forward overlap size. Using these parameters, we define the overlapping subdomains $[t_j^0, t_j^1]$ as

$$t_j^0 = \max\{t_{j-1} - \tau_j^0, 0\} \quad \text{and} \quad t_j^1 = \min\{t_j + \tau_j^1, T\} \quad \text{for } j = 1, \dots, m.$$

Each subproblem will be formulated and solved over its corresponding overlapping subdomain, allowing the scheme to iteratively propagate information across the entire time domain while exploiting parallelism at each iteration. To formalize the proposed scheme, we define the following parameterized subproblems. From this point onward, we denote the full OCP (5.1) as $\mathcal{P}([0, T]; x_0)$, with its optimal solution denoted by (x^*, u^*, λ^*) .

Definition 2. For each $1 \leq j \leq m$, we define the j -th subproblem on the interval $[t_j^0, t_j^1]$ with boundary parameters $(p_j, q_j) \in \mathbb{R}^{n_x} \times \mathbb{R}^{n_x}$, denoted by $\mathcal{P}_j([t_j^0, t_j^1]; p_j, q_j)$, as

$$\min_{u_j(\cdot), x_j(\cdot)} \mathcal{J}_j[u_j, x_j; q_j] := \frac{1}{2} \int_{t_j^0}^{t_j^1} \begin{bmatrix} x_j(t) \\ u_j(t) \end{bmatrix}^\top \begin{bmatrix} Q(t) & H^\top(t) \\ H(t) & R(t) \end{bmatrix} \begin{bmatrix} x_j(t) \\ u_j(t) \end{bmatrix} dt + L_j(x_j(t_j^1); q_j), \quad (5.11a)$$

$$\text{s.t.} \quad \dot{x}_j(t) = A(t)x_j(t) + B(t)u_j(t), \quad t \in (t_j^0, t_j^1], \quad (5.11b)$$

$$x_j(t_j^0) = p_j, \quad (5.11c)$$

where the terminal cost $L_j(x_j(t_j^1); q_j)$ is given by

$$L_j(x_j(t_j^1); q_j) := \begin{cases} \frac{1}{2}x_j^\top(t_j^1)Q(t_j^1)x_j(t_j^1) - x_j^\top(t_j^1)Q(t_j^1)q_j, & \text{if } 1 \leq j < m, \\ \frac{1}{2}x_j(T)^\top Q_T x_j(T), & \text{if } j = m. \end{cases} \quad (5.12)$$

In Definition 2, the parameter $p_j \in \mathbb{R}^{n_x}$ defines the initial state, and $q_j \in \mathbb{R}^{n_x}$ defines the target state, which is also used to define the terminal cost. In (5.12), the first case promotes continuity across overlapping intervals through a quadratic penalty, and the second case recovers the original terminal cost from the full problem (5.1). The inclusion of an appropriate penalty on the boundary centered at a reference point is essential to ensure boundary consistency. The following proposition establishes that, with properly chosen boundary parameters, the subproblems defined in (5.11) recover the truncated optimal solution of the full problem.

We note that it is beneficial to formulate the subproblems so that their optimality conditions take the same form as the PMP equations (5.8) for the full problem. This structural alignment ensures that a single numerical solver can be reused across parallel subproblems with minimal modification.

Proposition 4. *Suppose Assumptions 1 and 3 hold for $\mathcal{P}([0, T]; x_0)$. For each $1 \leq j \leq m$, we consider the subproblem $\mathcal{P}_j([t_j^0, t_j^1]; p_j^*, q_j^*)$ with the boundary parameters (p_j^*, q_j^*) defined as*

$$p_j^* = \begin{cases} x_0, & \text{if } j = 1, \\ x^*(t_j^0), & \text{if } j > 1, \end{cases} \quad q_j^* = \begin{cases} x^*(t_j^1) - Q^{-1}(t_j^1)\lambda^*(t_j^1), & \text{if } j < m, \\ 0, & \text{if } j = m. \end{cases} \quad (5.13)$$

Then, $\mathcal{P}_j([t_j^0, t_j^1]; p_j^*, q_j^*)$ admits a unique solution $(x_j^*, u_j^*, \lambda_j^*) : [t_j^0, t_j^1] \rightarrow \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_x}$, given by

$$x_j^*(t) = x^*(t), \quad u_j^*(t) = u^*(t), \quad \lambda_j^*(t) = \lambda^*(t), \quad t \in [t_j^0, t_j^1].$$

That is, the restriction of the full problem's solution to the interval $[t_j^0, t_j^1]$ solves the subproblem exactly.

Proof. By the necessary and sufficient conditions stated in Theorem 4, it suffices to verify that $\{x^*(t), u^*(t),$

$\lambda^*(t)\}_{t \in [t_j^0, t_j^1]}$ satisfies the PMP system for the subproblem $\mathcal{P}_j([t_j^0, t_j^1]; p_j^*, q_j^*)$ under the parameter choices given in (5.13). For any $1 \leq j \leq m$, let us define the Hamiltonian of the j -th subproblem as

$$\mathcal{H}_j(t, x_j(t), u_j(t), \lambda_j(t)) := \frac{1}{2} \begin{bmatrix} x_j(t) \\ u_j(t) \end{bmatrix}^\top \begin{bmatrix} Q(t) & H^\top(t) \\ H(t) & R(t) \end{bmatrix} \begin{bmatrix} x_j(t) \\ u_j(t) \end{bmatrix} + \lambda_j(t)^\top (A(t)x_j(t) + B(t)u_j(t)).$$

Then, the PMP system for the subproblem $\mathcal{P}_j([t_j^0, t_j^1]; p_j, q_j)$ is

$$\begin{aligned} \dot{x}_j^*(t) &= \nabla_{\lambda_j} \mathcal{H}_j(t, x_j^*(t), u_j^*(t), \lambda_j^*(t)) = A(t)x_j^*(t) + B(t)u_j^*(t), \quad x_j^*(t_j^0) = p_j, \\ \dot{\lambda}_j^*(t) &= -\nabla_{x_j} \mathcal{H}_j(t, x_j^*(t), u_j^*(t), \lambda_j^*(t)) \\ &= - \begin{bmatrix} Q(t) & H^\top(t) \end{bmatrix} \begin{bmatrix} x_j^*(t) \\ u_j^*(t) \end{bmatrix} - A^\top(t)\lambda_j^*(t), \quad \lambda_j^*(t_j^1) = \nabla_{x_j} L_j(x_j^*(t_j^1); q_j), \\ \mathcal{H}_j(t, x_j^*(t), u_j^*(t), \lambda_j^*(t)) &\leq \mathcal{H}_j(t, x_j^*(t), u, \lambda_j^*(t)), \quad \forall u \in \mathbb{R}^{n_u}, \quad \forall t \in [t_j^0, t_j^1]. \end{aligned}$$

Comparing the above system with the full PMP system in (5.8), we see that it suffices to verify $x^*(t_j^0) = p_j^*$ and $\lambda^*(t_j^1) = \nabla_{x_j} L_j(x^*(t_j^1); q_j^*)$, since all other conditions are subconditions and are directly implied by (5.8). Note that $x^*(t_j^0) = p_j^*$ is implied by (5.13). When $j = m$,

$$\lambda^*(t_m^1) = \lambda^*(T) \stackrel{(5.8b)}{=} Q_T x^*(T) = \nabla_{x_m} L_m(x^*(t_m^1); q_m^*).$$

When $1 \leq j < m$, we have

$$\nabla_{x_j} L_j(x^*(t_j^1); q_j^*) \stackrel{(5.12)}{=} Q(t_j^1) \left(x^*(t_j^1) - q_j^* \right) \stackrel{(5.13)}{=} \lambda^*(t_j^1).$$

This completes the proof. $\square$

Algorithm 4 Continuous-Time Overlapping Schwarz Decomposition

Input: Initial state, control, and adjoint trajectory $(x^{(0)}, u^{(0)}, \lambda^{(0)}) : [0, T] \rightarrow \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_\lambda}$ with $x^{(0)}(0) = x_0$; domain decomposition and its overlapping subdomains $[t_{j-1}, t_j] \subseteq [t_j^0, t_j^1]$, $1 \leq j \leq m$.

1. For $k = 0, 1, 2, \dots$:

(a) For $j = 1$ **to** m **in parallel**:

- i. Specify boundary parameters: $p_j^{(k+1)} = x^{(k)}(t_j^0)$ and $q_j^{(k+1)} = x^{(k)}(t_j^1) - Q^{-1}(t_j^1) \lambda^{(k)}(t_j^1)$.
- ii. Solve the subproblem $\mathcal{P}_j([t_j^0, t_j^1]; p_j^{(k+1)}, q_j^{(k+1)})$ to obtain $(x_j^{(k+1)}, u_j^{(k+1)}, \lambda_j^{(k+1)})$ (e.g., with warm initialization $(x_j^{(k)}, u_j^{(k)}, \lambda_j^{(k)})$ and Algorithm 5).

(b) Aggregate $(x^{(k+1)}, u^{(k+1)}, \lambda^{(k+1)})(t) := (x_j^{(k+1)}, u_j^{(k+1)}, \lambda_j^{(k+1)})(t)$ for $t \in [t_{j-1}, t_j]$ when $1 \leq j \leq m-1$ and for $t \in [t_{m-1}, t_m]$ when $j = m$.

Since each subproblem admits a unique optimal solution, Proposition 4 guarantees that, with correctly specified boundary parameters, solving $\mathcal{P}_j([t_j^0, t_j^1]; p_j^*, q_j^*)$ yields the restriction of the global solution (x^*, u^*, λ^*) to the sub-interval $[t_j^0, t_j^1]$. This observation motivates our decomposition scheme. Since in practice the optimal boundary parameters are unknown without solving full problem $\mathcal{P}([0, T]; x_0)$, we present the continuous-time overlapping Schwarz algorithm in Algorithm 4 to address this, which iteratively refines the parameters (p_j, q_j) . Starting with an imperfect set of parameters, the algorithm improves them over iterations, and ultimately drives them to converge to the full problem solution. We also present a gradient descent procedure for solving the subproblems in Appendix 5.8.1, where

the required gradient is derived via functional calculus.

In the following sections, we turn to the analysis of the properties of the linear-quadratic OCPs to investigate the theoretical foundations of the convergence of the Schwarz method.

5.4 Exponential Decay of Sensitivity in Linear-Quadratic Control

While Proposition 4 provides a theoretical foundation for recovering the full problem solution using the overlapping Schwarz scheme under ideal boundary conditions, it relies on knowledge of the exact optimal trajectory that is not available in practice. To rigorously justify the convergence behavior of Algorithm 4, we conduct a stability analysis of the OCP (5.1). In particular, we examine how errors in the boundary data influence the resulting solution, and how such perturbations propagate through the system dynamics over time.

Our sensitivity study can be made on a more general problem. We begin by introducing a parameterized linear-quadratic OCP that generalizes (5.1) by including additional dependence on an auxiliary input mapping $d : [0, T] \rightarrow \mathbb{R}^{n_d}$ and two boundary vectors $d_0 \in \mathbb{R}^{n_x}, d_T \in \mathbb{R}^{n_d}$:

$$\min_{u(\cdot), x(\cdot)} \quad \frac{1}{2} \int_0^T \begin{bmatrix} x(t) \\ u(t) \\ d(t) \end{bmatrix}^\top \begin{bmatrix} Q(t) & H^\top(t) & G^\top(t) \\ H(t) & R(t) & W^\top(t) \\ G(t) & W(t) & 0 \end{bmatrix} \begin{bmatrix} x(t) \\ u(t) \\ d(t) \end{bmatrix} dt + \frac{1}{2} \begin{bmatrix} x(T) \\ d_T \end{bmatrix}^\top \begin{bmatrix} Q_T & G_T^\top \\ G_T & 0 \end{bmatrix} \begin{bmatrix} x(T) \\ d_T \end{bmatrix}, \quad (5.15a)$$

$$\text{s.t.} \quad \dot{x}(t) = A(t)x(t) + B(t)u(t) + C(t)d(t), \quad t \in (0, T], \quad (5.15b)$$

$$x(0) = d_0. \quad (5.15c)$$

Here, $G : [0, T] \rightarrow \mathbb{R}^{n_d \times n_x}$, $W : [0, T] \rightarrow \mathbb{R}^{n_d \times n_u}$, and $G_T \in \mathbb{R}^{n_x \times n_d}$ represent additional linear contributions of the parameter trajectory to the cost; and $C : [0, T] \rightarrow \mathbb{R}^{n_x \times n_d}$ introduces a linear dependence of the dynamics on d . To ensure well-posedness, we extend

Assumption 1 to hold for parameter-dependent matrices.

Assumption 4. *We assume the mappings G, W, C are continuous and uniformly bounded: $\|G\|_\infty \leq \lambda_G$, $\|W\|_\infty \leq \lambda_W$, $\|C\|_\infty \leq \lambda_C$ for some constants $\lambda_G, \lambda_W, \lambda_C > 0$ independent of T . In addition, $\|G_T\| \leq \lambda_G$.*

In essence, Problem (5.15) allows bounded linear dependence on the mapping d and vectors d_0, d_T , which include the boundary parameterization as given by Proposition 4, where only d_0, d_T are possibly nonzero at domain boundaries of subproblems. We now state the equations that are central to our analysis.

Theorem 6. *Under Assumptions 1–4, the optimal state trajectory $x^*(t)$ of the parameterized OCP (5.15) satisfies the closed-loop system*

$$\begin{aligned} \dot{x}^*(t) = & \left\{ A(t) - B(t)R^{-1}(t)H(t) - B(t)R^{-1}(t)B^\top(t)S(t) \right\} x^*(t) - B(t)R^{-1}(t)B^\top(t)v(t) \\ & + \left\{ C(t) - B(t)R^{-1}(t)W^\top(t) \right\} d(t), \quad t \in (0, T], \end{aligned} \quad (5.16a)$$

$$x^*(0) = d_0, \quad (5.16b)$$

where the mapping $S : [0, T] \rightarrow \mathbb{R}^{n_x \times n_x}$ solves the matrix Riccati equation

$$\begin{aligned} \dot{S}(t) = & S(t)B(t)R^{-1}(t)B^\top(t)S(t) - S(t) \left\{ A(t) - B(t)R^{-1}(t)H(t) \right\} \\ & - \left\{ A(t) - B(t)R^{-1}(t)H(t) \right\}^\top S(t) - \left\{ Q(t) - H^\top(t)R^{-1}(t)H(t) \right\}, \quad t \in [0, T], \end{aligned} \quad (5.17a)$$

$$S(T) = Q_T, \quad (5.17b)$$

and the mapping $v : [0, T] \rightarrow \mathbb{R}^{n_x}$ solves the following backward linear dynamics

$$\begin{aligned} \dot{v}(t) = & - \left\{ A(t) - B(t)R^{-1}(t)H(t) - B(t)R^{-1}(t)B^\top(t)S(t) \right\}^\top v(t) \\ & + \left\{ W(t)R^{-1}(t) \left[B^\top(t)S(t) + H(t) \right] - (G(t) + C^\top(t)S(t)) \right\}^\top d(t), \quad t \in [0, T], \end{aligned} \quad (5.18a)$$

$$v(T) = G_T^\top d_T. \quad (5.18b)$$

Furthermore, the optimal control trajectory is given by

$$u^*(t) = -R^{-1}(t) \left\{ H(t)x^*(t) + B^\top(t)\lambda^*(t) + W^\top(t)d(t) \right\}, \quad (5.19)$$

where the adjoint is given by

$$\lambda^*(t) = S(t)x^*(t) + v(t).$$

Proof. The proof is provided in Appendix 5.8.2. □

The stability of the closed-loop system (5.16a) is governed by the matrix-valued function S , which evolves according to (5.17) and characterizes the sensitivity of the cost-to-go function (5.9). We will demonstrate that, if S remains positive definite and uniformly bounded (especially from below), the resulting states (5.16a) exhibit exponential stability. Thus, quantifying the spectral properties of S is essential for analyzing how perturbations in the parameters propagate through the system.

To facilitate the analysis, we introduce an auxiliary linear-quadratic control problem that possesses the same Riccati equation (5.16a) but without the linear dependence on the parameter d , which simplifies the derivation of spectral bounds.

Lemma 3 (Shifted OCP). *Suppose Q, R, H, A, B be time-varying matrices satisfying Assumptions 1 and 3. For any initial state $x_0 \in \mathbb{R}^{n_x}$, consider the following OCP on the*

interval $[0, T]$:

$$\min_{u(\cdot), x(\cdot)} \quad \frac{1}{2} \int_0^T \begin{bmatrix} x(t) \\ u(t) \end{bmatrix}^\top \begin{bmatrix} Q(t) - H^\top(t)R^{-1}(t)H(t) & 0 \\ 0 & R(t) \end{bmatrix} \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} dt + \frac{1}{2} x^\top(T) Q_T x(T), \quad (5.20a)$$

$$s.t. \quad \dot{x}(t) = (A(t) - B(t)R^{-1}(t)H(t))x(t) + B(t)u(t), \quad (5.20b)$$

$$x(0) = x_0.$$

Then, the associated Riccati matrix S of (5.20) is given by the same Riccati equation (5.17) as in Theorem 6.

Proof. See [111, Section 3.12] for the matrix Riccati equation corresponding to a linear-quadratic OCP without cross-product terms in the cost functional. By directly applying [111, (3.12-14)] to (5.20), we obtain the desired result. $\square$

We now study the properties of the Riccati matrix through the lens of the shifted problem (5.20). Note that the auxiliary problem shares the same Riccati equation as the original parameterized control problem (5.15); thus, any boundedness or decay properties of S can be analyzed through problem (5.20). We begin the discussion with an intermediate result.

Lemma 4. *Under Assumptions 1 and 3, a unique solution $S : [0, T] \rightarrow \mathbb{R}^{n_x \times n_x}$ of the matrix Riccati equation (5.17) exists and satisfies $S(t) \succ 0$ for any $t \in [0, T]$.*

Proof. For Problem (5.20), the existence and uniqueness of a positive semidefinite solution S to the Riccati equation (5.17) under Assumptions 1 and 3 is classical; see [23, Equations 6.1–6.4]. We provide an explicit verification of strict positive definiteness. For any $t_0 \in [0, T]$, we let (x^*, u^*) denote the truncated optimal solution of (5.20) on $[t_0, T]$. By the PMP necessary conditions in Theorem 4, there exists a continuously differentiable adjoint

trajectory $\lambda^* : [t_0, T] \rightarrow \mathbb{R}^{n_x}$ satisfying $\lambda^*(T) = Q_T x^*(T)$ and

$$\dot{\lambda}^*(t) = -(A(t) - B(t)R^{-1}(t)H(t))^\top \lambda^*(t) - (Q(t) - H^\top(t)R^{-1}(t)H(t))x^*(t), \quad t \in [t_0, T]. \quad (5.21)$$

In fact, applying Theorem 6 with $d = 0$ (so $v = 0$) and $H = 0$, we know the optimal adjoint is explicitly given by $\lambda^*(t) = S(t)x^*(t)$ and the optimal control is explicitly given by $u^*(t) = -R^{-1}(t)B^\top(t)\lambda^*(t)$. With these formulations and for any initial state $0 \neq x_{t_0} \in \mathbb{R}^{n_x}$, we have

$$\begin{aligned} x_{t_0}^\top S(t_0)x_{t_0} &= x^*(T)^\top Q_T x^*(T) - \int_{t_0}^T \frac{d}{dt} (x^*(t)^\top S(t)x^*(t)) dt \\ &= x^*(T)^\top Q_T x^*(T) - \int_{t_0}^T \frac{d}{dt} (x^*(t)^\top \lambda^*(t)) dt. \end{aligned} \quad (5.22)$$

On the other hand, we apply the chain rule and substitute the dynamics of x^* , u^* , and λ^* , and obtain for any $t \in (t_0, T)$,

$$\begin{aligned} \frac{d}{dt} (x^*(t)^\top \lambda^*(t)) &= \dot{x}^*(t)^\top \lambda^*(t) + x^*(t)^\top \dot{\lambda}^*(t) \\ &\stackrel{(5.20b), (5.21)}{=} \left\{ (A(t) - B(t)R^{-1}(t)H(t))x^*(t) + B(t)u^*(t) \right\}^\top \lambda^*(t) \\ &\quad - x^*(t)^\top \left\{ (A(t) - B(t)R^{-1}(t)H(t))^\top \lambda^*(t) + (Q(t) - H^\top(t)R^{-1}(t)H(t))x^*(t) \right\} \\ &= -x^*(t)^\top (Q(t) - H^\top(t)R^{-1}(t)H(t))x^*(t) - \lambda^*(t)^\top B(t)R^{-1}(t)B^\top(t)\lambda^*(t) \\ &\leq -x^*(t)^\top (Q(t) - H^\top(t)R^{-1}(t)H(t))x^*(t) \leq 0, \end{aligned}$$

where the third equality is because $u^*(t) = -R^{-1}(t)B^\top(t)\lambda^*(t)$ and the last two inequalities are due to Assumption 3. Since $x_{t_0} \neq 0$, by the continuity of x^* , we know there exists $t'_0 > t_0$ such that $x^*(t) \neq 0$ for any $t \in [t_0, t'_0]$. Within the interval $[t_0, t'_0]$, the above

derivation implies $(d/dt)(x^*(t)^\top \lambda^*(t)) < 0$ strictly (by Assumption 3). Thus, we have

$$\int_{t_0}^T \frac{d}{dt}(x^*(t)^\top \lambda^*(t))dt \leq \int_{t_0}^{t'_0} \frac{d}{dt}(x^*(t)^\top \lambda^*(t))dt < 0.$$

Combining the above display with (5.22) and noting that $x^*(T)^\top Q_T x^*(T) \geq 0$, we have $x_{t_0}^\top S(t_0)x_{t_0} > 0$. Since x_{t_0} is any nonzero vector, we have $S(t_0) \succ 0$. This completes the proof since it holds for all $t_0 \in [0, T]$. $\square$

While Lemma 4 ensures that the solution to the matrix Riccati equation exists and is positive definite on $[0, T]$, this property alone is not sufficient to determine how perturbations propagate. In our context, it is important that S is not only positive definite but also uniformly bounded from both above and below. To proceed, we require an intermediate result confirming that if the uniform controllability condition (Assumption 2) holds for the system matrix pair (A, B) , then it also holds for the pair $(A - BR^{-1}H, B)$. The latter corresponds to the shifted system in (5.20).

Lemma 5. *Consider the linear system $\dot{x}(t) = A(t)x(t) + B(t)u(t)$. Suppose $A : [0, T] \rightarrow \mathbb{R}^{n_x \times n_x}$ and $B : [0, T] \rightarrow \mathbb{R}^{n_x \times n_u}$ define a matrix pair (A, B) satisfying Assumptions 1 and 2. Then, for any continuous mapping $F : [0, T] \rightarrow \mathbb{R}^{n_u \times n_x}$ satisfying $\|F\|_\infty \leq \lambda_F$ for some constant $\lambda_F > 0$ independent of T , the system*

$$\dot{x}(t) = (A(t) + B(t)F(t))x(t) + B(t)u(t)$$

also satisfies UCC in the sense of (5.4).

In particular, there exist constants $\alpha'_0(\sigma), \alpha'_1(\sigma), \beta'_0(\sigma), \beta'_1(\sigma) > 0$, depending on σ from Assumption 2 but independent of T , such that for every $t_0 \in [0, T - \sigma]$, there exists a time

$t_1 = t_1(t_0) \in [t_0, t_0 + \sigma] \subseteq [0, T]$ for which the following bounds hold:

$$\begin{aligned}\alpha'_0(\sigma)I &\preceq W_{A+BF,B}(t_0, t_1) \preceq \alpha'_1(\sigma)I \\ \beta'_0(\sigma)I &\preceq \Phi_{A+BF}(t_1, t_0)W_{A+BF,B}(t_0, t_1)\Phi_{A+BF}^\top(t_1, t_0) \preceq \beta'_1(\sigma)I.\end{aligned}$$

(The expressions of $\alpha'_0(\sigma)$, $\alpha'_1(\sigma)$, $\beta'_0(\sigma)$, $\beta'_1(\sigma) > 0$ are provided in (5.30), (5.25), and (5.31).)

Proof. By Assumption 1, we note that for any $0 \leq s \leq t \leq T$ (see [165, Theorem 5.2(i)]),

$$\|\Phi_{A+BF}(t, s)\| \leq \exp\left(\int_s^t \|A(\tau) + B(\tau)F(\tau)\| d\tau\right) \leq e^{(\lambda_A + \lambda_B \lambda_F)(t-s)}. \quad (5.23)$$

Furthermore, we note that

$$\begin{aligned}1 &= \|I\| = \|\Phi_{A+BF}(t, s)\Phi_{A+BF}^{-1}(t, s)\| = \|\Phi_{A+BF}(t, s)\Phi_{A+BF}(s, t)\| \\ &= \|\Phi_{A+BF}(t, s)\Phi_{-(A+BF)^\top}^\top(t, s)\| \leq \|\Phi_{A+BF}(t, s)\| \cdot \|\Phi_{-(A+BF)^\top}(t, s)\| \\ &\leq \|\Phi_{A+BF}(t, s)\| e^{(\lambda_A + \lambda_B \lambda_F)(t-s)},\end{aligned}$$

where the last inequality is due to (5.23). The above inequality implies

$$\|\Phi_{A+BF}(t, s)\| \geq e^{-(\lambda_A + \lambda_B \lambda_F)(t-s)}. \quad (5.24)$$

Now, let us consider bounding the controllability Gramian for the pair $(A + BF, B)$. For any $t_0 \in [0, T - \sigma]$ and $0 \neq v \in \mathbb{R}^{n_x}$, we let $t_1 = t_1(t_0)$ from Assumption 2. Then, by the

definition (5.3) we have

$$\begin{aligned}
v^\top W_{A+BF,B}(t_0, t_1)v &= \int_{t_0}^{t_1} v^\top \Phi_{A+BF}(t_0, s) B(s) B^\top(s) \Phi_{A+BF}^\top(t_0, s) v ds \\
&= \int_{t_0}^{t_1} \|B^\top(s) \Phi_{A+BF}^\top(t_0, s) v\|^2 ds \leq \lambda_B^2 \|v\|^2 \int_{t_0}^{t_1} \|\Phi_{-(A+BF)^\top}(s, t_0)\|^2 ds \\
&\leq \lambda_B^2 \|v\|^2 \int_{t_0}^{t_1} e^{2(\lambda_A + \lambda_B \lambda_F)(s-t_0)} ds \leq \lambda_B^2 \|v\|^2 \int_{t_0}^{t_0+\sigma} e^{2(\lambda_A + \lambda_B \lambda_F)(s-t_0)} ds \\
&= \frac{\lambda_B^2 \|v\|^2}{2(\lambda_A + \lambda_B \lambda_F)} (\exp \{2\sigma(\lambda_A + \lambda_B \lambda_F)\} - 1),
\end{aligned}$$

which implies

$$\|W_{A+BF,B}(t_0, t_1)\| \leq \frac{\lambda_B^2 (\exp \{2\sigma(\lambda_A + \lambda_B \lambda_F)\} - 1)}{2(\lambda_A + \lambda_B \lambda_F)} =: \alpha'_1(\sigma). \quad (5.25)$$

For the lower bound, given any $0 \neq v \in \mathbb{R}^{n_x}$, we consider the zero-steering control defined on $[t_0, t_1]$ to be $u(t; v) := -B^\top(t) \Phi_A^\top(t_0, t) W_{A,B}^{-1}(t_0, t_1) v$ (see Lemma 2). With this choice of control, the solution of the system $\dot{x}(t) = A(t)x(t) + B(t)u(t; v)$, $t \in (t_0, t_1]$, $x(t_0) = v$, denoted by $x(t; v)$, satisfies $x(t_1, v) = 0$. Then, let us consider the system $(A+BF, B)$. Define the control trajectory

$$\tilde{u}(t; v) := u(t; v) - F(t)x(t; v), \quad \text{for } t \in [t_0, t_1], \quad (5.26)$$

and the corresponding state trajectory of the system with the above control: $\dot{x}(t) = (A(t) + B(t)F(t))x(t) + B(t)\tilde{u}(t; v)$, $t \in (t_0, t_1]$, $x(t_0) = v$, denoted by $\tilde{x}(t; v)$, satisfies $\tilde{x}(t; v) = x(t; v)$, $\forall t \in [t_0, t_1]$. To see this, we note that

$$\begin{aligned}
\dot{\tilde{x}}(t; v) - \dot{x}(t; v) &= (A(t) + B(t)F(t))\tilde{x}(t; v) + B(t)\tilde{u}(t; v) - A(t)x(t; v) - B(t)u(t; v) \\
&= (A(t) + B(t)F(t))(\tilde{x}(t; v) - x(t; v)).
\end{aligned}$$

Combining the above display with the fact that $\tilde{x}(t_0; v) - x(t_0; v) = 0$, we know $\tilde{x}(t; v) = x(t; v)$, $\forall t \in [t_0, t_1]$. In particular, $\tilde{x}(t_1; v) = x(t_1; v) = 0$. On the other hand, by [223] and Lemma 2, we know $W_{A+BF, B}(t_0, t_1)$ is positive definite, and the control trajectory

$$\bar{u}(t; v) := -B^\top(t)\Phi_{A+BF}^\top(t_0, t)W_{A+BF, B}^{-1}(t_0, t_1)v, \quad t \in [t_0, t_1] \quad (5.27)$$

steers the state trajectory $x(t)$ of $(A + BF, B)$ from v at t_0 to 0 at t_1 . In fact, by [196, Section 3.5, Theorem 5], $\bar{u}(t; v)$ is the unique control trajectory that attains the minimum square norm. That is,

$$\int_{t_0}^{t_1} \|\bar{u}(t; v)\|^2 dt \leq \int_{t_0}^{t_1} \|\tilde{u}(t; v)\|^2 dt.$$

For the left-hand side, we have

$$\begin{aligned} \int_{t_0}^{t_1} \|\bar{u}(t; v)\|^2 dt &= \int_{t_0}^{t_1} \bar{u}^\top(t; v)\bar{u}(t; v) dt \\ &\stackrel{(5.27)}{=} \int_{t_0}^{t_1} \|B^\top(t)\Phi_{A+BF}^\top(t_0, t)W_{A+BF, B}^{-1}(t_0, t_1)v\|^2 dt \\ &\stackrel{(5.3)}{=} v^\top W_{A+BF, B}^{-1}(t_0, t_1)v. \end{aligned}$$

Combining the above two displays, we apply Assumption 2 and the boundedness of $F(t)$, and obtain

$$\begin{aligned} v^\top W_{A+BF, B}^{-1}(t_0, t_1)v &\leq \int_{t_0}^{t_1} \|\tilde{u}(t; v)\|^2 dt \\ &\stackrel{(5.26)}{=} \int_{t_0}^{t_1} \|u(t; v) - F(t)x(t; v)\|^2 dt \\ &\leq 2 \int_{t_0}^{t_1} \|u(t; v)\|^2 dt + 2 \int_{t_0}^{t_1} \|F(t)x(t; v)\|^2 dt \\ &= 2v^\top W_{A, B}^{-1}(t_0, t_1)v + 2 \int_{t_0}^{t_1} \|F(t)x(t; v)\|^2 dt \\ &\leq 2\alpha_0^{-1}(\sigma)\|v\|^2 + 2\lambda_F^2 \int_{t_0}^{t_1} \|x(t; v)\|^2 dt. \end{aligned} \quad (5.28)$$

To upper bound the second term on the right-hand side, we use the following explicit expression of $x(t; v)$ (by applying (5.2)):

$$\begin{aligned} x(t; v) &= \Phi_A(t, t_0)v + \int_{t_0}^t \Phi_A(t, s)B(s)u(s; v)ds \\ &= \Phi_A(t, t_0) \left(I - W_{A,B}(t_0, t)W_{A,B}^{-1}(t_0, t_1) \right) v, \quad t \in [t_0, t_1]. \end{aligned}$$

From the definition (5.3), we have for any $t \in [t_0, t_1]$ that

$$W_{A,B}(t_0, t_1) = W_{A,B}(t_0, t) + W_{A,B}(t, t_1) \succeq W_{A,B}(t_0, t). \quad (5.29)$$

Therefore, we have

$$\begin{aligned} \int_{t_0}^{t_1} \|x(t; v)\|^2 dt &= \int_{t_0}^{t_1} \left\| \Phi_A(t, t_0) \left(I - W_{A,B}(t_0, t)W_{A,B}^{-1}(t_0, t_1) \right) v \right\|^2 dt \\ &\stackrel{(5.29)}{\leq} \int_{t_0}^{t_1} e^{2\lambda_A(t-t_0)} \cdot \left(1 + \frac{\alpha_1(\sigma)}{\alpha_0(\sigma)} \right)^2 \|v\|^2 ds \\ &\leq \int_{t_0}^{t_0+\sigma} e^{2\lambda_A(t-t_0)} \cdot \left(1 + \frac{\alpha_1(\sigma)}{\alpha_0(\sigma)} \right)^2 \|v\|^2 dt \\ &= \frac{1}{2\lambda_A} \left(1 + \frac{\alpha_1(\sigma)}{\alpha_0(\sigma)} \right)^2 \left(e^{2\lambda_A\sigma} - 1 \right) \|v\|^2. \end{aligned}$$

Plugging the above display into (5.28), we obtain

$$v^\top W_{A+BF,B}^{-1}(t_0, t_1)v \leq \left(\frac{2}{\alpha_0(\sigma)} + \frac{\lambda_F^2}{\lambda_A} \left(1 + \frac{\alpha_1(\sigma)}{\alpha_0(\sigma)} \right)^2 \{\exp(2\lambda_A\sigma) - 1\} \right) \|v\|^2.$$

Since $v \neq 0 \in \mathbb{R}^{n_x}$ is arbitrary and $W_{A+BF,B}(t_0, t_1)$ is positive definite, this proves that

$$W_{A+BF,B}(t_0, t_1) \succeq \alpha'_0(\sigma)I, \quad (5.30)$$

with

$$\alpha'_0(\sigma) = \left(\frac{2}{\alpha_0(\sigma)} + \frac{\lambda_F^2}{\lambda_A} \left(1 + \frac{\alpha_1(\sigma)}{\alpha_0(\sigma)} \right)^2 \{ \exp(2\lambda_A\sigma) - 1 \} \right)^{-1}.$$

Furthermore, we define

$$\beta'_0(\sigma) := \alpha'_0(\sigma) \cdot e^{-2(\lambda_A + \lambda_B \lambda_F)\sigma} \quad \text{and} \quad \beta'_1(\sigma) := \alpha'_1(\sigma) \cdot e^{2(\lambda_A + \lambda_B \lambda_F)\sigma}, \quad (5.31)$$

and obtain for any $0 \neq v \in \mathbb{R}^{n_x}$ that

$$\begin{aligned} v^\top \Phi_{A+BF}(t_1, t_0) W_{A+BF, B}(t_0, t_1) \Phi_{A+BF}^\top(t_1, t_0) v \\ \stackrel{(5.23)}{\leq} \alpha'_1(\sigma) \cdot e^{2(\lambda_A + \lambda_B \lambda_F)\sigma} \|v\|^2 = \beta'_1(\sigma) \|v\|^2, \end{aligned}$$

$$\begin{aligned} v^\top \Phi_{A+BF}(t_1, t_0) W_{A+BF, B}(t_0, t_1) \Phi_{A+BF}^\top(t_1, t_0) v \\ \stackrel{(5.24)}{\geq} \alpha'_0(\sigma) \cdot e^{-2(\lambda_A + \lambda_B \lambda_F)\sigma} \|v\|^2 = \beta'_0(\sigma) \|v\|^2. \end{aligned}$$

Thus, we conclude the proof. $\square$

The above lemma is helpful because we effectively require UCC of the matrix pair $(A - BR^{-1}H, B)$ to establish the upper and lower bounds of Riccati equation solution S (cf. Lemma 3) and to further analyze perturbations in d for the parameterized problem (5.15), while Assumption 2 imposes UCC only on (A, B) . The above lemma thus provides the necessary connection. We emphasize that the UCC property is not required on the interval $[T - \sigma, T]$ since this segment has length at most σ . The upper and lower bounds on relevant quantities are effectively guaranteed by the explicit relationship between the Riccati matrix S and the optimal cost functional (details in Appendix 5.8.3). We now state the boundedness result.

Lemma 6. *Under Assumptions 1–4, there exist constants $c_0(\sigma), c_1(\sigma) > 0$ independent of*

T such that the unique solution of the matrix Riccati equation (5.17) satisfies

$$c_0(\sigma)I \preceq S(t_0) \preceq c_1(\sigma)I, \quad \text{for all } t_0 \in [0, T].$$

(The expressions of $c_0(\sigma), c_1(\sigma) > 0$ are provided in (5.81).)

Proof. The proof is provided in Appendix 5.8.3. □

We have established that the Riccati matrix is uniformly bounded from above and below on the entire solution interval $[0, T]$ with constants independent of T . This result provides the foundation for analyzing sensitivity properties of the optimal state system (5.16a). The result is summarized in the next lemma.

Lemma 7 (Exponential stability of optimal state system). *Define the following time-dependent matrix:*

$$Z(t) := A(t) - B(t)R^{-1}(t)H(t) - B(t)R^{-1}(t)B^\top(t)S(t), \quad t \in [0, T]. \quad (5.32)$$

Consider the homogeneous system $\dot{x}(t) = Z(t)x(t)$. Under Assumptions 1–4, the corresponding solution evolution operator satisfies the exponential decay bound: for any $0 \leq t_0 \leq t_1 \leq T$,

$$\|\Phi_Z(t_1, t_0)\| \leq c_Z(\sigma)e^{-\rho_Z(\sigma)(t_1-t_0)}, \quad (5.33)$$

where

$$c_Z(\sigma) = \sqrt{\frac{c_1(\sigma)}{c_0(\sigma)}} \quad \text{and} \quad \rho_Z(\sigma) = \frac{\gamma_Q}{2c_1(\sigma)}. \quad (5.34)$$

Proof. Let us define the candidate Lyapunov function $V(t, x(t)) := x(t)^\top S(t)x(t)$, $t \in [t_0, t_1]$.

In particular, when $x(t)$ follows the dynamics $\dot{x}(t) = Z(t)x(t)$, we have

$$\begin{aligned}
\frac{d}{dt}V(t, x(t)) &= \dot{x}^\top(t)S(t)x(t) + x^\top(t)\dot{S}(t)x(t) + x^\top(t)S(t)\dot{x}(t) \\
&= x^\top(t)Z^\top(t)S(t)x(t) + x^\top(t)\dot{S}(t)x(t) + x^\top(t)S(t)Z(t)x(t) \\
&= -x^\top(t)S(t)B(t)R^{-1}(t)B^\top(t)S(t)x(t) - x^\top(t)(Q(t) - H^\top(t)R^{-1}(t)H(t))x(t) \\
&\leq -\gamma_Q\|x(t)\|^2 \leq -\frac{\gamma_Q}{c_1(\sigma)}V(t, x(t)),
\end{aligned}$$

where the third equality is due to (5.32) and (5.17), and the last inequality is due to Lemma 6.

By Grönwall's inequality, this implies that for any $t \in [t_0, t_1]$,

$$\begin{aligned}
c_0(\sigma)\|\Phi_Z(t, t_0)x(t_0)\|^2 &= c_0(\sigma)\|x(t)\|^2 \leq x^\top(t)S(t)x(t) = V(t, x(t)) \\
&\leq V(t_0, x(t_0)) \exp\left(-\frac{\gamma_Q}{c_1(\sigma)}(t - t_0)\right) \\
&= x^\top(t_0)S(t_0)x(t_0) \exp\left(-\frac{\gamma_Q}{c_1(\sigma)}(t - t_0)\right) \leq c_1(\sigma) \exp\left(-\frac{\gamma_Q}{c_1(\sigma)}(t - t_0)\right) \|x(t_0)\|^2.
\end{aligned}$$

This leads to

$$\|\Phi_Z(t, t_0)x(t_0)\| \leq \sqrt{\frac{c_1(\sigma)}{c_0(\sigma)}} \exp\left(-\frac{\gamma_Q}{2c_1(\sigma)}(t - t_0)\right) \|x(t_0)\|.$$

Since $x(t_0)$ can be any vector in $\mathbb{R}^{n_x}$, we apply the fact that

$$\|\Phi_Z(t, t_0)\| := \max_{0 \neq x \in \mathbb{R}^{n_x}} \|\Phi_Z(t, t_0)x\|/\|x\|$$

and complete the proof. $\square$

Lemma 7 provides a direct consequence which we call exponential decay of sensitivity (EDS). That is, the effect of a single-point perturbation on the optimal solutions of our parameterized problem (5.15) is damped exponentially fast as one moves away from the

perturbation point. We present this main result below. We need some additional notation.

Let us denote $x^*(d, d_0, d_T) : [0, T] \rightarrow \mathbb{R}^{n_x}$, $u^*(d, d_0, d_T) : [0, T] \rightarrow \mathbb{R}^{n_u}$, $\lambda^*(d, d_0, d_T) : [0, T] \rightarrow \mathbb{R}^{n_x}$ to be the unique optimal state, control, and adjoint solutions to Problem (5.15) with the parameters $d : [0, T] \rightarrow \mathbb{R}^{n_d}$ and $d_0 \in \mathbb{R}^{n_x}$, $d_T \in \mathbb{R}^{n_d}$. For any $t' \in [0, T]$, we consider a single-point perturbation of d at t' , given by

$$d_p(t) = d(t) + l_{t'}\delta(t - t'), \quad t \in [0, T], \quad (5.35)$$

where $0 \neq l_{t'} \in \mathbb{R}^{n_d}$ is the perturbation direction and $\delta(\cdot)$ is the Dirac delta functional. With the perturbed parameters (d_p, d_0, d_T) , we denote the perturbed solutions of (5.15) as $(x_p^*, u_p^*, \lambda_p^*)(d_p, d_0, d_T) : [0, T] \rightarrow \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_x}$. The next result investigates the perturbation of the solutions, characterized by $\delta x = x_p^* - x^*$, $\delta u = u_p^* - u^*$, and $\delta \lambda = \lambda_p^* - \lambda^*$.

Theorem 7. *Consider the perturbation of d at t' in (5.35) and suppose Assumptions 1–4 hold. Then, there exists a constant $\Lambda = \Lambda(\sigma) > 0$, independent of T , such that*

$$\|\delta x(t)\| + \|\delta u(t)\| + \|\delta \lambda(t)\| \leq \|l_{t'}\| \cdot \Lambda(\sigma) e^{-\rho_Z(\sigma)|t-t'|}, \quad \forall t' \neq t \in [0, T],$$

where $\rho_Z(\sigma)$ is defined in (5.34). (The expression of $\Lambda(\sigma)$ is provided in (5.45).)

Proof. We model the perturbation as a Dirac delta functional $\delta_{t'}(t) := \delta(t - t')$, which, once taken an inner product with a function, results in a point evaluation at $t' \in [0, T]$ in the following sense: for any $[t_a, t_b] \subseteq [0, T]$,

$$\int_{t_a}^{t_b} x(t) \delta_{t'}(t) dt = \begin{cases} x(t'), & \text{if } t' \in [t_a, t_b], \\ 0, & \text{if } t' \notin [t_a, t_b]. \end{cases} \quad (5.36)$$

We refer the reader to classic texts in Sobolev spaces and distribution theory for a detailed discussion of the function-space implications of modeling with the Dirac delta [104]. In

particular, introducing a point perturbation at $t = t'$ induces a jump discontinuity in the optimal trajectories. By the arguments made in [15, Section 2], Theorems 4 and 5 still hold uniquely but are to be interpreted in the almost everywhere sense.

Let us define $v_p : [0, T] \rightarrow \mathbb{R}^{n_x}$ to be the vector equation (5.18) associated with perturbed parameters d_p . In particular, we have

$$\begin{aligned} \dot{v}_p(t) &= - \left\{ A(t) - B(t)R^{-1}(t)H(t) - B(t)R^{-1}(t)B^\top(t)S(t) \right\}^\top v_p(t) \\ &\quad + \left\{ W(t)R^{-1}(t) \left[B^\top(t)S(t) + H(t) \right] - (G(t) + C^\top(t)S(t)) \right\}^\top d_p(t) \\ &\stackrel{(5.32)}{=} -Z^\top(t)v_p(t) + Y^\top(t)d_p(t), \quad t \in [0, T), \\ v_p(T) &= G_T^\top d_T, \end{aligned}$$

where in the second equality we also define

$$Y(t) := W(t)R^{-1}(t) \left[B^\top(t)S(t) + H(t) \right] - (G(t) + C^\top(t)S(t)), \quad t \in [0, T].$$

Let v denote the original solution that satisfies the unperturbed equation (5.18), and define $\delta v := v_p - v$. Then, the equation of δv is obtained by taking the difference between the above display and (5.18), and is given by

$$\delta \dot{v}(t) = -Z^\top(t)\delta v(t) + Y^\top(t)l_{t'}\delta_{t'}(t), \quad t \in [0, T), \quad \delta v(T) = 0.$$

For consistency of using the solution evolution operator as in (5.2), we transfer the above backward equation to forward equation by performing the change of variable $r := T - t$, $t \in [0, T]$, and defining $\delta \tilde{v}(r) := -\delta v(T - r)$. Thus, we have $\delta \dot{\tilde{v}}(r) = \delta \dot{v}(T - r)$. Within the

r -coordinate system and defining $\tilde{Z}(r) := Z(T - r)$, we obtain

$$\begin{aligned}
\delta\dot{v}(r) &= \delta\dot{v}(T - r) = -Z^\top(T - r)\delta v(T - r) + Y^\top(T - r)l_{t'}\delta_{t'}(T - r) \\
&= \tilde{Z}^\top(r)\delta\tilde{v}(r) + Y^\top(T - r)l_{t'}\delta(T - r - t') \\
&= \tilde{Z}^\top(r)\delta\tilde{v}(r) + Y^\top(T - r)l_{t'}\delta(r - (T - t')) \\
&= \tilde{Z}^\top(r)\delta\tilde{v}(r) + Y^\top(T - r)l_{t'}\delta_{T-t'}(r), \quad r \in (0, T], \\
\delta\tilde{v}(0) &= 0.
\end{aligned}$$

Applying the mild solution in (5.2), we have

$$\delta\tilde{v}(r) = \int_0^r \Phi_{\tilde{Z}^\top}(r, s)Y^\top(T - s)l_{t'}\delta_{T-t'}(s)ds \stackrel{(5.36)}{=} \begin{cases} 0, & r < T - t', \\ \Phi_{\tilde{Z}^\top}(r, T - t')Y^\top(t')l_{t'}, & r \geq T - t'. \end{cases}$$

The above implies that

$$\delta v(t) = -\delta\tilde{v}(T - t) = \begin{cases} -\Phi_{\tilde{Z}^\top}(T - t, T - t')Y^\top(t')l_{t'}, & t \leq t' \\ 0, & t > t'. \end{cases} \quad (5.38)$$

We now relate $\Phi_{\tilde{Z}^\top}$ to Φ_Z . For fixed $r \in [0, T]$ and any $s \in [r, T]$, we know from [165, Theorem 5.2(iv)] that $\Phi_{\tilde{Z}^\top}(s, r)$ as a function of s satisfies

$$\frac{\partial \Phi_{\tilde{Z}^\top}(s, r)}{\partial s} = \tilde{Z}^\top(s)\Phi_{\tilde{Z}^\top}(s, r) = Z^\top(T - s)\Phi_{\tilde{Z}^\top}(s, r), \quad s \in (r, T], \quad \Phi_{\tilde{Z}^\top}(r, r) = I.$$

On the other hand, the operator $\Phi_Z^\top(T - r, T - s)$ as a function of s satisfies

$$\frac{\partial \Phi_Z^\top(T - r, T - s)}{\partial s} = (\Phi_Z(T - r, T - s)Z(T - s))^\top = Z^\top(T - s)\Phi_Z^\top(T - r, T - s), \quad s \in (r, T],$$

$$\Phi_Z(T - r, T - r) = I.$$

By the uniqueness of the solution [165, Theorem 5.1], we finally obtain

$$\Phi_{\tilde{Z}^\top}(s, r) = \Phi_Z^\top(T - r, T - s), \quad 0 \leq r \leq s \leq T. \quad (5.39)$$

Plugging the above display into (5.38), and defining $\mathbf{1}_{t \leq t'} = 1$ if $t \leq t'$ and 0 otherwise, we have

$$\begin{aligned} \|\delta v(t)\| &\leq \|\Phi_{\tilde{Z}^\top}(T - t, T - t')\| \cdot \|Y(t')\| \cdot \|l_{t'}\| \mathbf{1}_{t \leq t'} = \|\Phi_Z(t', t)\| \cdot \|Y(t')\| \cdot \|l_{t'}\| \mathbf{1}_{t \leq t'} \\ &\leq \|l_{t'}\| c_Z(\sigma) e^{-\rho_Z(\sigma)(t' - t)} \left\{ \frac{\lambda_W}{\gamma_R} (\lambda_B c_1(\sigma) + \lambda_H) + \lambda_C c_1(\sigma) + \lambda_G \right\} \mathbf{1}_{t \leq t'} \\ &= \|l_{t'}\| c_Z(\sigma) \left\{ \left(\frac{\lambda_W \lambda_B}{\gamma_R} + \lambda_C \right) c_1(\sigma) + \frac{\lambda_W \lambda_H}{\gamma_R} + \lambda_G \right\} e^{-\rho_Z(\sigma)(t' - t)} \mathbf{1}_{t \leq t'} \\ &=: \|l_{t'}\| c_v(\sigma) e^{-\rho_Z(\sigma)(t' - t)} \mathbf{1}_{t \leq t'}, \end{aligned} \quad (5.40)$$

where the third inequality is due to Lemmas 6 and 7 and Assumptions 1 and 4.

With the above result (5.40), we now establish the decay of δx . By Theorem 6, we have for any $t' \neq t \in [0, T]$,

$$x^*(t) = \Phi_Z(t, 0) d_0 + \int_0^t \Phi_Z(t, s) g(s; d, v) ds,$$

where we define

$$g(t; d, v) := -B(t)R^{-1}(t)B^\top(t)v(t) + \left\{ C(t) - B(t)R^{-1}(t)W^\top(t) \right\} d(t).$$

Applying Theorem 6 to the perturbed problem and taking the difference, we have

$$\begin{aligned}
\delta x(t) &= \int_0^t \Phi_Z(t, s) (g(s; d_p, v_p) - g(s; d, v)) ds \\
&= \int_0^t \Phi_Z(t, s) \left(-B(s)R^{-1}(s)B^\top(s)\delta v(s) + \left\{ C(s) - B(s)R^{-1}(s)W^\top(s) \right\} l_{t'}\delta_{t'}(s) \right) ds \\
&\stackrel{(5.36)}{=} - \int_0^t \Phi_Z(t, s)B(s)R^{-1}(s)B^\top(s)\delta v(s)ds \\
&\quad + \Phi_Z(t, t') \left\{ C(t') - B(t')R^{-1}(t')W^\top(t') \right\} l_{t'} \cdot \mathbf{1}_{t>t'}.
\end{aligned}$$

Thus, we apply Lemma 7 and obtain

$$\begin{aligned}
\|\delta x(t)\| &\leq \frac{\lambda_B^2}{\gamma_R} \int_0^t \|\Phi_Z(t, s)\| \|\delta v(s)\| ds + \left(\lambda_C + \frac{\lambda_W \lambda_B}{\gamma_R} \right) \|\Phi_Z(t, t')\| \|l_{t'}\| \mathbf{1}_{t>t'} \\
&\stackrel{(5.40)}{\leq} \|l_{t'}\| c_v(\sigma) \frac{\lambda_B^2}{\gamma_R} \int_0^{\min\{t, t'\}} \|\Phi_Z(t, s)\| e^{-\rho_Z(\sigma)(t'-s)} ds \\
&\quad + \|l_{t'}\| \left(\lambda_C + \frac{\lambda_W \lambda_B}{\gamma_R} \right) \|\Phi_Z(t, t')\| \mathbf{1}_{t>t'} \\
&\leq \|l_{t'}\| c_v(\sigma) c_Z(\sigma) \frac{\lambda_B^2}{\gamma_R} \int_0^{\min\{t, t'\}} e^{-\rho_Z(\sigma)(t-s)} e^{-\rho_Z(\sigma)(t'-s)} ds \\
&\quad + \|l_{t'}\| c_Z(\sigma) \left(\lambda_C + \frac{\lambda_W \lambda_B}{\gamma_R} \right) e^{-\rho_Z(\sigma)(t-t')} \mathbf{1}_{t>t'} \\
&= \|l_{t'}\| c_v(\sigma) c_Z(\sigma) \frac{\lambda_B^2}{\gamma_R} \cdot e^{-\rho_Z(\sigma)(t+t')} \frac{e^{2\rho_Z(\sigma)\min\{t, t'\}} - 1}{2\rho_Z(\sigma)} \\
&\quad + \|l_{t'}\| c_Z(\sigma) \left(\lambda_C + \frac{\lambda_W \lambda_B}{\gamma_R} \right) e^{-\rho_Z(\sigma)(t-t')} \mathbf{1}_{t>t'} \\
&\leq \|l_{t'}\| \frac{c_v(\sigma) c_Z(\sigma) \lambda_B^2}{2\rho_Z(\sigma) \gamma_R} e^{-\rho_Z(\sigma)|t-t'|} + \|l_{t'}\| c_Z(\sigma) \left(\lambda_C + \frac{\lambda_W \lambda_B}{\gamma_R} \right) e^{-\rho_Z(\sigma)(t-t')} \mathbf{1}_{t>t'} \\
&\leq \|l_{t'}\| c_Z(\sigma) \left\{ \frac{c_v(\sigma) \lambda_B^2}{2\rho_Z(\sigma) \gamma_R} + \frac{\lambda_W \lambda_B}{\gamma_R} + \lambda_C \right\} e^{-\rho_Z(\sigma)|t-t'|} =: \|l_{t'}\| \Lambda_x(\sigma) e^{-\rho_Z(\sigma)|t-t'|}, \quad (5.41)
\end{aligned}$$

where $c_v(\sigma)$ is defined in (5.40). Next, we study the adjoint trajectory. By Theorem 6, we

have for any $t' \neq t \in [0, T]$,

$$\begin{aligned}
\|\delta\lambda(t)\| &= \|\lambda_p^*(t) - \lambda^*(t)\| = \|S(t)(x_p^*(t) - x^*(t)) + (v_p(t) - v(t))\| \\
&\leq c_1(\sigma)\|\delta x(t)\| + \|\delta v(t)\| \leq \|l_{t'}\|c_1(\sigma)\Lambda_x(\sigma)e^{-\rho_Z(\sigma)|t-t'|} + \|l_{t'}\|c_v(\sigma)e^{-\rho_Z(\sigma)(t'-t)}\mathbf{1}_{t \leq t'} \\
&\leq \|l_{t'}\| \{c_1(\sigma)\Lambda_x(\sigma) + c_v(\sigma)\} e^{-\rho_Z(\sigma)|t-t'|} =: \|l_{t'}\|\Lambda_\lambda(\sigma)e^{-\rho_Z(\sigma)|t-t'|}, \tag{5.42}
\end{aligned}$$

where the fourth inequality is due to (5.40) and (5.41). Finally, we consider the control trajectory. By (5.19) in Theorem 6, we know for $t' \neq t \in [0, T]$ that

$$\begin{aligned}
\|\delta u(t)\| &:= \|u_p^*(t) - u^*(t)\| \\
&= \|-R^{-1}(t)\{H(t)\delta x(t) + B^\top(t)\delta\lambda(t) + W^\top(t)l_{t'}\delta_{t'}(t)\}\| \tag{5.43}
\end{aligned}$$

$$\begin{aligned}
&= \|-R^{-1}(t)\{H(t)\delta x(t) + B^\top(t)\delta\lambda(t)\}\| \\
&\leq \frac{\lambda_H}{\gamma_R}\|\delta x(t)\| + \frac{\lambda_B}{\gamma_R}\|\delta\lambda(t)\| \stackrel{(5.41), (5.42)}{\leq} \|l_{t'}\| \cdot \frac{\Lambda_x(\sigma)\lambda_H + \Lambda_\lambda(\sigma)\lambda_B}{\gamma_R} e^{-\rho_Z(\sigma)|t-t'|} \\
&=: \|l_{t'}\|\Lambda_u(\sigma)e^{-\rho_Z(\sigma)|t-t'|}. \tag{5.44}
\end{aligned}$$

Combining (5.41), (5.42), (5.43), we complete the proof by defining

$$\begin{aligned}
\Lambda &= \Lambda(\sigma) := \Lambda_x(\sigma) + \Lambda_u(\sigma) + \Lambda_\lambda(\sigma) \\
&= \left(1 + \frac{\lambda_H}{\gamma_R}\right)\Lambda_x(\sigma) + \left(1 + \frac{\lambda_B}{\gamma_R}\right)\Lambda_\lambda(\sigma) \tag{5.45}
\end{aligned}$$

$$\begin{aligned}
&= \left(1 + \frac{\lambda_H}{\gamma_R}\right)\Lambda_x(\sigma) + \left(1 + \frac{\lambda_B}{\gamma_R}\right)(c_1(\sigma)\Lambda_x(\sigma) + c_v(\sigma)) \\
&= \left\{1 + \frac{\lambda_H}{\gamma_R} + c_1(\sigma)\left(1 + \frac{\lambda_B}{\gamma_R}\right)\right\}\Lambda_x(\sigma) + c_v(\sigma)\left(1 + \frac{\lambda_B}{\gamma_R}\right), \tag{5.46}
\end{aligned}$$

where $c_1(\sigma)$ is in Lemma 6 (cf. (5.81)), $\Lambda_x(\sigma)$ is in (5.41), and $c_v(\sigma)$ is in (5.40). $\square$

Theorem 7 provides the EDS result with respect to the general parameters d that Problem (5.15) has linear dependence on in the integral term in the cost functional. To complete

the section, we analyze the effects of perturbing the boundary parameterizations d_0 and d_T . Following the setup as in Theorem 7, we consider the boundary perturbations of Problem (5.15):

$$d_{p,0} = d_0 + l_0, \quad d_{p,T} = d_T + l_T, \quad (5.47)$$

where $l_0 \in \mathbb{R}^{n_x}$, $l_T \in \mathbb{R}^{n_d}$ are perturbation directions. We also use $(x_p^*, u_p^*, \lambda_p^*)(d, d_{p,0}, d_{p,T})$ to denote the perturbed solution and investigate the solution perturbations measured by $\delta x = x_p^* - x^*$, $\delta u = u_p^* - u^*$, and $\delta \lambda = \lambda_p^* - \lambda^*$.

Theorem 8. *Consider the perturbation of d_0, d_T in (5.47) and suppose Assumptions 1–4 hold. Then, there exists a constant $\Lambda_b = \Lambda_b(\sigma) > 0$, independent of T , such that*

$$\|\delta x(t)\| + \|\delta u(t)\| + \|\delta \lambda(t)\| \leq \Lambda_b(\sigma) \left(\|l_0\| e^{-\rho_Z(\sigma)t} + \|l_T\| e^{-\rho_Z(\sigma)(T-t)} \right), \quad t \in [0, T],$$

where $\rho_Z(\sigma)$ is defined in (5.34). (The expression of $\Lambda_b(\sigma)$ is provided in (5.53).)

Proof. The proof follows from carrying out a similar analysis to that of Theorem 7, where we consider the mild solution of the optimal state (see (5.2)) before and after introducing the perturbations l_0, l_T . Define $v_p = v + \delta v$ as the perturbed solution of (5.18) with terminal condition $v_p(T) = G_T^\top(d_T + l_T)$. Then, $\delta v = v_p - v$ satisfies

$$\delta \dot{v}(t) = -Z(t)^\top \delta v(t), \quad t \in [0, T], \quad \delta v(T) = G_T^\top l_T,$$

where Z is defined in (5.32). By using the reparameterization $r := T - t$ and defining $\delta \tilde{v}(r) := -\delta v(T - r)$ and $\tilde{Z}(r) := Z(T - r)$, we obtain the forward in time dynamics

$$\delta \dot{\tilde{v}}(r) = \tilde{Z}^\top(r) \delta \tilde{v}(r), \quad r \in (0, T], \quad \delta \tilde{v}(0) = -G_T^\top l_T.$$

Applying (5.2), we have for $t \in [0, T]$ that

$$\|\delta v(t)\| = \|\delta \tilde{v}(T-t)\| = \|\Phi_{\tilde{Z}^\top}(T-t, 0)G_T^\top l_T\| \quad (5.48)$$

$$\stackrel{(5.39)}{\leq} \lambda_G \|l_T\| \|\Phi_Z(T, t)\| \leq \|l_T\| \lambda_G c_Z(\sigma) e^{-\rho_Z(\sigma)(T-t)}, \quad (5.49)$$

where the last inequality is due to Lemma 7. Furthermore, by (5.16a), we know

$$\begin{aligned} \dot{x}_p^*(t) &= Z(t)x_p^*(t) - B(t)R^{-1}(t)B^\top(t)v_p(t) + \left\{ C(t) - B(t)R^{-1}(t)W^\top(t) \right\} d(t), \quad t \in (0, T], \\ x_p^*(0) &= d_0 + l_0, \end{aligned}$$

which leads to

$$\delta \dot{x}(t) = Z(t)\delta x(t) - B(t)R^{-1}(t)B^\top(t)\delta v(t), \quad t \in (0, T], \quad \delta x(0) = l_0.$$

Thus, we obtain from (5.2) that for any $t \in [0, T]$,

$$\begin{aligned} \|\delta x(t)\| &= \left\| \Phi_Z(t, 0)l_0 - \int_0^t \Phi_Z(t, s)B(s)R^{-1}(s)B^\top(s)\delta v(s)ds \right\| \\ &\leq \|l_0\| \cdot c_Z(\sigma) e^{-\rho_Z(\sigma)t} + \frac{\lambda_B^2}{\gamma_R} \int_0^t \|\Phi_Z(t, s)\| \|\delta v(s)\| ds \\ &\stackrel{(5.48)}{\leq} \|l_0\| \cdot c_Z(\sigma) e^{-\rho_Z(\sigma)t} + \|l_T\| \cdot c_Z^2(\sigma) \frac{\lambda_B^2 \lambda_G}{\gamma_R} \int_0^t e^{-\rho_Z(\sigma)(t-s)} e^{-\rho_Z(\sigma)(T-s)} ds \\ &= \|l_0\| \cdot c_Z(\sigma) e^{-\rho_Z(\sigma)t} + \|l_T\| \cdot \frac{c_Z^2(\sigma)}{2\rho_Z(\sigma)} \frac{\lambda_B^2 \lambda_G}{\gamma_R} \left(e^{-\rho_Z(\sigma)(T-t)} - e^{-\rho_Z(\sigma)(T+t)} \right) \\ &\leq \|l_0\| \cdot c_Z(\sigma) e^{-\rho_Z(\sigma)t} + \|l_T\| \cdot \frac{c_Z^2(\sigma)}{2\rho_Z(\sigma)} \frac{\lambda_B^2 \lambda_G}{\gamma_R} e^{-\rho_Z(\sigma)(T-t)} \\ &\leq \Lambda_{x,b}(\sigma) \left(\|l_0\| e^{-\rho_Z(\sigma)t} + \|l_T\| e^{-\rho_Z(\sigma)(T-t)} \right), \end{aligned} \quad (5.50)$$

where we define

$$\Lambda_{x,b}(\sigma) := \max \left\{ c_Z(\sigma), \frac{c_Z^2(\sigma) \lambda_B^2 \lambda_G}{2\rho_Z(\sigma) \gamma_R} \right\}.$$

For the adjoint trajectory, we apply Theorem 6, Lemma 6, and have for any $t \in [0, T]$,

$$\begin{aligned}
\|\delta\lambda(t)\| &= \|\lambda_p^*(t) - \lambda^*(t)\| = \|S(t)(x_p^*(t) - x^*(t)) + (v_p(t) - v(t))\| \leq \|S(t)\| \|\delta x(t)\| + \|\delta v(t)\| \\
&\stackrel{(5.48), (5.50)}{\leq} c_1(\sigma) \Lambda_{x,b}(\sigma) \left(\|l_0\| e^{-\rho_Z(\sigma)t} + \|l_T\| e^{-\rho_Z(\sigma)(T-t)} \right) + \|l_T\| \lambda_G c_Z(\sigma) e^{-\rho_Z(\sigma)(T-t)} \\
&\leq \Lambda_{x,b}(\sigma) (c_1(\sigma) + \lambda_G) \left(\|l_0\| e^{-\rho_Z(\sigma)t} + \|l_T\| e^{-\rho_Z(\sigma)(T-t)} \right) \\
&=: \Lambda_{\lambda,b}(\sigma) \left(\|l_0\| e^{-\rho_Z(\sigma)t} + \|l_T\| e^{-\rho_Z(\sigma)(T-t)} \right). \tag{5.51}
\end{aligned}$$

Finally, we consider the control trajectory in (5.19) and have for $t \in [0, T]$,

$$\begin{aligned}
\|\delta u(t)\| &:= \|u_p^*(t) - u^*(t)\| \\
&\stackrel{(5.19)}{=} \left\| -R^{-1}(t) \{ H(t) \delta x(t) + B^\top(t) \delta \lambda(t) \} \right\| \\
&\leq \frac{\lambda_H}{\gamma_R} \|\delta x(t)\| + \frac{\lambda_B}{\gamma_R} \|\delta \lambda(t)\| \\
&\stackrel{(5.50), (5.51)}{\leq} \frac{\Lambda_{x,b}(\sigma) \lambda_H + \Lambda_{\lambda,b}(\sigma) \lambda_B}{\gamma_R} \left(\|l_0\| e^{-\rho_Z(\sigma)t} + \|l_T\| e^{-\rho_Z(\sigma)(T-t)} \right) \\
&=: \Lambda_{u,b}(\sigma) \left(\|l_0\| e^{-\rho_Z(\sigma)t} + \|l_T\| e^{-\rho_Z(\sigma)(T-t)} \right). \tag{5.52}
\end{aligned}$$

Finally, we combine (5.50), (5.51), (5.52), and complete the proof by defining $\Lambda_b(\sigma)$ as

$$\begin{aligned}
\Lambda_b &= \Lambda_b(\sigma) = \Lambda_{x,b}(\sigma) + \Lambda_{u,b}(\sigma) + \Lambda_{\lambda,b}(\sigma) \\
&= \left(1 + \frac{\lambda_H}{\gamma_R} \right) \Lambda_{x,b}(\sigma) + \left(1 + \frac{\lambda_B}{\gamma_R} \right) \Lambda_{\lambda,b}(\sigma) = \left\{ 1 + \frac{\lambda_H}{\gamma_R} + \left(1 + \frac{\lambda_B}{\gamma_R} \right) (c_1(\sigma) + \lambda_G) \right\} \Lambda_{x,b} \\
&= \left\{ 1 + \frac{\lambda_H}{\gamma_R} + \left(1 + \frac{\lambda_B}{\gamma_R} \right) (c_1(\sigma) + \lambda_G) \right\} \max \left\{ c_Z(\sigma), \frac{c_Z^2(\sigma) \lambda_B^2 \lambda_G}{2\rho_Z(\sigma) \gamma_R} \right\}, \tag{5.53}
\end{aligned}$$

where $c_1(\sigma)$ is in Lemma 6 (cf. (5.81)) and $c_Z(\sigma)$ and $\rho_Z(\sigma)$ are in (5.34). $\square$

Up to this point, we have established the sensitivity of the solution of the parameterized linear-quadratic problem (5.15) to the perturbations in the input mapping d and two boundary vectors d_0, d_T . We recall Proposition 4 that the overlapping Schwarz scheme introduces

boundary vector perturbations due to imperfect boundary conditions. In the next section, we leverage Theorem 8 to show that, by recursively updating the boundary conditions, the overlapping Schwarz scheme achieves linear convergence with the linear rate decaying exponentially fast in terms of the size of the overlap.

5.5 Convergence of Continuous-Time OSD

In this section, we present the convergence analysis as a consequence of EDS established in Section 5.4, which provides a rigorous justification for the use of overlapping subdomains in the proposed Algorithm 4. Specifically, we set $d = 0$ in (5.15) and analyze the subproblems in (5.11), parameterized by p_j, q_j . Since perturbations arising from inaccuracies in the initial and terminal parameters of subproblems are damped exponentially as they propagate into the interior of each subinterval, we demonstrate that even when the boundary parameters $d_0 := p_j$ and $d_T := q_j$ are not initialized exactly at p_j^*, q_j^* – the optimal boundary values of the global problem (5.1) (in particular, those satisfying the condition in Proposition 4) – the resulting subproblem solutions remain close to the global optimal trajectory on each non-overlapping interval $[t_{j-1}, t_j]$, provided that the overlaps between subdomains are sufficiently large.

We now present a key result that quantifies perturbations within the domain to errors introduced at the subdomain boundaries. We recall that (x^*, u^*, λ^*) denotes the (unique) optimal state, control, and adjoint trajectories of the full problem $\mathcal{P}([0, T]; x_0)$ in (5.1).

Lemma 8 (Stagewise improvement). *Suppose Assumption 1–3 hold for Problem (5.1), and consider applying Algorithm 4 to solve Problem (5.1). In particular, $(x^{(k)}, u^{(k)}, \lambda^{(k)})$ is the k -th solution iterate aggregating m pieces $\{(x_j^{(k)}, u_j^{(k)}, \lambda_j^{(k)})\}_{j=1}^m$, which are solutions of $\mathcal{P}_j([t_j^0, t_j^1]; p_j^{(k)}, q_j^{(k)})$ in (5.11) with boundary parameters $p_j^{(k)} = x^{(k-1)}(t_j^0)$ and $q_j^{(k)} = x^{(k-1)}(t_j^1) - Q^{-1}(t_j^1)\lambda^{(k-1)}(t_j^1)$. Let $\tau := \min_{1 \leq j \leq m} \min\{\tau_j^0, \tau_j^1\}$ denote the smallest overlap size between subdomains. Then, we have for any $k \geq 1$, any $1 \leq j \leq m$, and any $t \in$*

$[t_{j-1}, t_j]$,

$$\begin{aligned} & \|x^{(k)}(t) - x^*(t)\| + \|u^{(k)}(t) - u^*(t)\| + \|\lambda^{(k)}(t) - \lambda^*(t)\| \\ & \leq c(\sigma)e^{-\rho_Z(\sigma)\tau} \left(\|x^{(k-1)}(t_j^0) - x^*(t_j^0)\| \right. \\ & \quad \left. + (\|x^{(k-1)}(t_j^1) - x^*(t_j^1)\| + \|\lambda^{(k-1)}(t_j^1) - \lambda^*(t_j^1)\|) \mathbf{1}_{j \neq m} \right). \end{aligned}$$

where $\rho_Z(\sigma)$ is defined in (5.34) and $c(\sigma) := \Lambda_b(\sigma)(1 + 1/\gamma_Q)$ with $\Lambda_b(\sigma)$ in (5.53). Here, for both $\rho_Z(\sigma)$ and $\Lambda_b(\sigma)$, we use λ_Q in place of λ_G and set $\lambda_C = \lambda_W = 0$.

Proof. By Proposition 4, the solution (x^*, u^*, λ^*) of $\mathcal{P}([0, T]; x_0)$ restricted to $[t_{j-1}, t_j]$ is equivalent to the solution of $\mathcal{P}_j([t_j^0, t_j^1]; p_j^*, q_j^*)$ (see (5.13) for definitions of p_j^*, q_j^*) with the overlaps discarded. We quantify how boundary errors propagate into the interior of each subdomain. For any $k \geq 1$ and $1 \leq j \leq m$, we define two perturbation vectors on the two ends (only initial perturbation vector when $j = m$):

$$\begin{aligned} l_{j,0}^{(k)} &:= p_j^{(k)} - p_j^* \stackrel{\text{Alg. 4}}{=} x^{(k-1)}(t_j^0) - x^*(t_j^0), \\ l_{j,T}^{(k)} &:= q_j^{(k)} - q_j^* \stackrel{\text{Alg. 4}}{=} (x^{(k-1)}(t_j^1) - x^*(t_j^1)) - Q^{-1}(t_j^1)(\lambda^{(k-1)}(t_j^1) - \lambda^*(t_j^1)). \end{aligned}$$

Note that $\mathcal{P}_j([t_j^0, t_j^1]; p_j, q_j)$ is in the form of (5.15) with $d(t) = 0$, $C(t) = 0$, $G(t) = 0$, $W(t) = 0$, $G_T = -Q(t_j^1)$, $d_0 = p_j$, $d_T = q_j$. Therefore, by Theorem 8 and Proposition 4, for

any $t \in [t_{j-1}, t_j]$, we have

$$\begin{aligned}
& \|x^{(k)}(t) - x^*(t)\| + \|u^{(k)}(t) - u^*(t)\| + \|\lambda^{(k)}(t) - \lambda^*(t)\| \\
&= \|x_j^{(k)}(t) - x^*(t)\| + \|u_j^{(k)}(t) - u^*(t)\| + \|\lambda_j^{(k)}(t) - \lambda^*(t)\| \\
&\leq \Lambda_b(\sigma) \left(\|l_{j,0}^{(k)}\| e^{-\rho_Z(\sigma)(t-t_j^0)} + \|l_{j,T}^{(k)}\| e^{-\rho_Z(\sigma)(t_j^1-t)} \mathbf{1}_{j \neq m} \right) \\
&\leq \Lambda_b(\sigma) e^{-\rho_Z(\sigma)\tau} \left(\|l_{j,0}^{(k)}\| + \|l_{j,T}^{(k)}\| \mathbf{1}_{j \neq m} \right) \\
&\leq \Lambda_b(\sigma) (1 + 1/\gamma_Q) e^{-\rho_Z(\sigma)\tau} \left(\|x^{(k-1)}(t_j^0) - x^*(t_j^0)\| \right. \\
&\quad \left. + \left\{ \|x^{(k-1)}(t_j^1) - x^*(t_j^1)\| + \|\lambda^{(k-1)}(t_j^1) - \lambda^*(t_j^1)\| \right\} \mathbf{1}_{j \neq m} \right),
\end{aligned}$$

where the third inequality is due to the fact that $t - t_j^0 \geq \tau$ and $t_j^1 - t \geq \tau$ since $t \in [t_{j-1}, t_j]$.

We complete the proof. $\square$

We are now ready to present our main convergence result of the overlapping Schwarz scheme.

Theorem 9 (Linear convergence of overlapping Schwarz). *Suppose Assumptions 1–3 hold for Problem (5.1), and consider applying Algorithm 4 to solve Problem (5.1). For any $k \geq 0$, we define*

$$\omega^{(k)} = \max\{\|x^{(k)} - x^*\|_\infty, \|u^{(k)} - u^*\|_\infty, \|\lambda^{(k)} - \lambda^*\|_\infty\}.$$

Then, Algorithm 4 exhibits the linear convergence

$$\omega^{(k+1)} \leq 3c(\sigma) e^{-\rho_Z(\sigma)\tau} \cdot \omega^{(k)},$$

where $c(\sigma)$, $\rho_Z(\sigma)$, and τ are defined in Lemma 8.

Proof. The result directly follows by applying Lemma 8:

$$\begin{aligned}
\omega^{(k+1)} &\leq \sup_{t \in [0, T]} \|x^{(k+1)}(t) - x^*(t)\| + \|u^{(k+1)}(t) - u^*(t)\| + \|\lambda^{(k+1)}(t) - \lambda^*(t)\| \\
&\leq c(\sigma)e^{-\rho_Z(\sigma)\tau} \left(2\|x^{(k)}(t) - x^*(t)\|_\infty + \|\lambda^{(k)}(t) - \lambda^*(t)\|_\infty \right) \\
&\leq 3c(\sigma)e^{-\rho_Z(\sigma)\tau} \cdot \omega^{(k)},
\end{aligned}$$

which concludes that the convergence is linear. $\square$

We see from Theorem 9 that the linear rate $3c(\sigma)e^{-\rho_Z(\sigma)\tau} < 1$ holds as long as

$$\tau > \log(3c(\sigma))/\rho_Z(\sigma),$$

and the above threshold depends only on the controllability parameter $\sigma \ll T$ in Assumption 2, but is independent of the full horizon length T . This makes our algorithm particularly promising for long-horizon problems. Furthermore, the linear rate $3c(\sigma)e^{-\rho_Z(\sigma)\tau}$ improves exponentially with respect to the overlap size τ . Overall, the results in Theorem 9 demonstrate the advantage of incorporating domain overlaps during optimization in Algorithm 4, as it accelerates convergence in the interior regions. The solutions in the overlapped regions can then be discarded while concatenating the remainder. In the numerical experiment section, we provide an example to support this insight.

5.6 Numerical Experiment

To illustrate the effectiveness of the proposed overlapping Schwarz scheme, we draw inspiration from a nonlinear optimal control problem with a final-time constraint studied in [185], and consider a linearized version of it around the origin. We aim to demonstrate three promising properties of our Schwarz scheme (Algorithm 4) via the experiment:

1. The overlapping Schwarz scheme converges linearly, with the linear convergence rate improving exponentially fast in terms of the overlap size.
2. Higher-order time integrators can be easily incorporated when solving the forward and backward equations (5.8a) and (5.8b) for the truncated subproblems.
3. In contrast to “discretize-then-optimize” approaches that fix the time step in advance, our continuous-time Schwarz scheme supports adaptive time-stepping, allowing efficient resolution of stiff dynamics.

To this end, we consider the original nonlinear OCP in [185] formulated in $\mathbb{R}^2$ as:

$$\min_{u(\cdot), x(\cdot)} \quad \frac{1}{2} \int_0^T \|x(t)\|^2 + \|u(t)\|^2 dt, \quad (5.54a)$$

$$\text{s.t.} \quad \dot{x}(t) = b_\xi(x(t)) + Nu(t), \quad t \in (0, T], \quad (5.54b)$$

$$x(0) = \begin{bmatrix} 0 & 0 \end{bmatrix}^\top, \quad \phi(x(T)) = \theta, \quad (5.54c)$$

where $\theta \in \mathbb{R}$ is an adjustable parameter, $N \in \mathbb{R}^{2 \times 2}$ is the control coefficient matrix, and the state dynamics map $b_\xi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and the terminal constraint map $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$ are defined by

$$b_\xi(x) := \begin{bmatrix} -x_1 - x_1 x_2 \\ -\xi x_2 + x_1^2 \end{bmatrix}, \quad \phi(x) := x_1 + 2x_2.$$

Here, $\xi > 0$ is a stiffness parameter that characterizes the relative decay rate of the two state components, and the terminal constraint imposes a scalar condition on the final state. We linearize the nonlinear dynamics (5.54b) at the equilibrium point $x = 0$ and incorporate the terminal constraint (5.54c) into the objective via a soft penalty, resulting in the following

linear-quadratic OCP:

$$\min_{u(\cdot), x(\cdot)} \quad \frac{1}{2} \int_0^T \|x(t)\|^2 + \|u(t)\|^2 dt + \frac{\alpha}{2} (\phi(x(T)) - \theta)^2, \quad (5.55a)$$

$$\text{s.t.} \quad \dot{x}(t) = M_\xi x(t) + Nu(t), \quad t \in (0, T], \quad (5.55b)$$

$$x(0) = 0, \quad (5.55c)$$

where $\alpha > 0$ is the penalty parameter (chosen sufficiently large), and $M_\xi \in \mathbb{R}^{2 \times 2}$ is the Jacobian of b_ξ at origin given by $M_\xi := \nabla b(0; \xi) = \text{diag}(-1, -\xi)$.

In our experiment, we first follow [185] and set $N = \text{diag}(1, 0.25)$, $T = 5$, $\theta = 3$, $\xi = 4$, and $\alpha = 10^2$. We will also vary the parameters by increasing ξ when testing the Schwarz method on stiff dynamics. Throughout the experiment, the true reference solution (x^*, u^*, λ^*) is computed by applying the forward Euler discretization to the full problem (5.55) on a fine uniform grid with the time step $\Delta t = 10^{-3}$, and solving the resulting discrete problem using MATLAB's built-in function `fmincon`. For the overlapping Schwarz method, we partition the time interval $[0, T]$ into $m = 3$ uniform subdomains, each corresponding to a subproblem as defined in (5.11). We express the overlap parameters $\tau_j^0 = \tau_j^1 = \tau$ as a fraction of the subdomain length, and will vary different τ when we illustrate the improvement of the linear rate in τ (i.e., the goal **(a)**). Each subproblem is solved in parallel using Algorithm 5, with the step size $\eta = 10^{-2}$ and the initial control sampled as $u(t) \sim \mathcal{N}(0, I)$. Note that Algorithm 5 (Lines 3 and 4) requires integrating the state and adjoint equations forward and backward, respectively (cf. (5.8a), (5.8b)). We will test the robustness of the Schwarz method to different time integrators (i.e., the goal **(b)**), as well as the benefits of enabling adaptive time-stepping when the dynamics is stiff (i.e., the goal **(c)**). We terminate Algorithm 5 when the norm of the local objective gradient falls below 10^{-6} . Lastly, let $(x^{(k)}, u^{(k)}, \lambda^{(k)})$ denote the k -th Schwarz iterate; the convergence is measured by




Figure 5.1: (Left) Convergence of the overlapping Schwarz method under varying overlap sizes ($\{1\%, 5\%, 10\%, 20\%, 30\%, 60\%\}$ of the subdomain length). Larger overlaps yield faster convergence. (Right) Observed convergence rate and its exponential fit. For each overlap size, we estimate the convergence rate by averaging the ratios of two consecutive errors, and fit an exponential curve of the theoretical form $c \cdot \exp(-\rho_Z \tau)$. The fitted values are $\hat{c} = 0.98$ and $\hat{\rho}_Z = 0.05$; and we have $\hat{\rho}_Z < 0.02$ at the significance level of 5%. We observe that the convergence rate improves exponentially in the overlap size, which is consistent with Theorem 9.

the maximum pointwise deviation from the reference solution, where $t_j = j \cdot \Delta t$,

$$e^{(k)} := \max_{0 \leq j \leq \frac{T}{\Delta t}} \left(\|x^{(k)}(t_j) - x^*(t_j)\| + \|u^{(k)}(t_j) - u^*(t_j)\| + \|\lambda^{(k)}(t_j) - \lambda^*(t_j)\| \right).$$

• **Goal (a).** We verify Theorem 9 by illustrating the convergence behavior of the overlapping Schwarz method. We vary the overlap proportion $\tau \in \{1\%, 5\%, 10\%, 20\%, 30\%, 60\%\}$ relative to the subdomain length, and the forward and backward equations of the subproblems (cf. Algorithm 5, Lines 3 and 4) are simply integrated using the forward Euler method.

The convergence behavior is illustrated in Figure 5.1 (left). From the figure, we observe that the errors decay linearly over the Schwarz iterations, with larger overlaps yielding faster convergence, which indicates that the linear rate improves as the overlap size increases. To examine the linear rate more closely, we estimate it by averaging the ratios of two consecutive errors for each τ and plot the resulting rates versus the overlap proportion in Figure 5.1

(right). We find that the observed convergence rate aligns with a fitted model of the form $\hat{c}_Z \exp(-\hat{\rho}_Z \tau)$, suggesting that the rate improves precisely exponentially with the overlap size. This further supports the theoretical insight of the exponential improvement of the rate with the overlap size in Theorem 9. Overall, the empirical results validate our analytic rate bound and highlight the reasoning of using large overlaps to accelerate convergence.

- **Goal (b).** We now investigate the effect of applying different time integration schemes to the subproblems on the convergence behavior of the overlapping Schwarz method. In this experiment, we fix the overlap size at 5% of the subdomain length. To clearly isolate the effect of the solver choice on convergence, we employ a coarser temporal discretization when applying the integration methods, with a step size of $\Delta t = 0.05$ (recall that the reference solution is computed on a much finer grid with $\Delta t = 10^{-3}$).

Figure 5.2 (left) compares the explicit forward Euler (FE), implicit backward Euler (BE), and explicit Runge-Kutta (RK45) methods for solving the subproblem dynamics (5.8a) and (5.8b). We observe that higher-order solvers attain improved solution accuracy, roughly half an additional order of accuracy per upgrade in method fidelity. While higher-order solvers are indeed more computationally expensive, they are applied only to the subproblems within the Schwarz scheme. Furthermore, despite the differences in numerical integration error, the overall convergence with respect to Schwarz iterations remains linear, consistent with Theorem 9.

- **Goal (c).** A key advantage of our continuous-time formulation is its compatibility with a wide range of time-stepping schemes. In contrast, high-order solvers such as RK45 are not easily integrated into discretize-then-optimize methods [193, 152]. Our approach enables adaptive and high-fidelity solvers to be used without compromising the structure of the optimality system.

To demonstrate the benefits of adaptive time-stepping, we compare the behavior of solvers in a stiff setting with $\xi = 15$ and $T = \theta = 10$, where the forward Euler method becomes




Figure 5.2: (Left) Comparison of numerical solvers in resolving the subproblem Hamiltonian dynamics. The Runge-Kutta method (RK45) yields the highest accuracy per iteration. (Right) Comparison of first 30 Schwarz iterations for different solvers and time step sizes Δt . Adaptive solvers maintain low errors even on coarse grids, while FE accuracy degrades significantly in the stiff regime.

unstable. In this experiment, we compare the fixed-step FE method with the adaptive solver `ode23` [188], applied to the Hamiltonian systems of the local subproblems. As expected, we see from Figure 5.2 (right) that forward Euler fails to converge unless the time step is reduced significantly. In contrast, `ode23` maintains stability and consistently achieves lower error across Schwarz iterations, demonstrating the robustness of our method when paired with adaptive integration.

5.7 Conclusion and Future Work

In this work, we introduced a continuous-time overlapping Schwarz decomposition method for linear-quadratic optimal control problems and established its convergence properties. The method decomposes the global control problem into smaller subproblems defined on overlapping temporal subdomains, with appropriate boundary conditions specified at the interfaces to enforce consistency between adjacent subdomains. Our analysis reveals that the method leverages the intrinsic exponential decay of sensitivity in continuous-time optimal control problems, such that the impact of imperfect boundary conditions on the solution

trajectory decays exponentially as one moves away from the subdomain boundaries. As a result, the Schwarz method achieves linear convergence by recursively updating the boundary conditions. Notably, the convergence rate improves exponentially with the size of the overlap. The continuous-time formulation offers several advantages, including compatibility with higher-order and adaptive time-stepping schemes, as well as improved handling of stiff system dynamics. Numerical experiments corroborate our theoretical findings and demonstrate the practical effectiveness of the method on a stiff linear-quadratic optimal control problem.

Several promising extensions can be pursued. First, we believe it is possible to extend the overlapping Schwarz decomposition from linear-quadratic program to the nonlinear case, as a continuous-time generalization of [152, 193]. However, doing so would require a functional analog of the convexification procedure [207], which we leave for future exploration. Second, in this work we have adopted a gradient-based optimization approach in solving continuous-time subproblems (Algorithm 5). The local convergence of the subproblems can be further enhanced using second-order optimization solvers [99], which can be implemented by solving second-order adjoint equations. Other practical directions include developing robust solvers compatible with the Schwarz scheme. Possible extensions include directly solving matrix Riccati equations [185] or applying symplectic Runge-Kutta methods [183]. Finally, there are opportunities to apply the Schwarz method in deep learning, particularly in settings where models can be framed as locally linear-quadratic optimal control problems of the form (5.1). The established convergence properties of the Schwarz method serve as a foundation for practical applications of broad control methods inspired by turnpike analysis [78], which aligns well with deep learning architectures that can be interpreted as discrete dynamical systems, such as neural ODEs [69]. Future work can explore these intersections, leading to model- and data-parallel distributed learning algorithms.

5.8 Appendix

5.8.1 Gradient Method for Subproblems

To solve the subproblems $\mathcal{P}_j([t_j^0, t_j^1]; p_j, q_j)$, we can apply gradient-based optimization method in infinite-dimensional space summarized in Algorithm 5.

Algorithm 5 Gradient Descent for Subproblem $\mathcal{P}_j([t_j^0, t_j^1]; p_j, q_j)$

Input: Subproblem control $u_j^{(0)}$; initial and terminal state parameters (p_j, q_j) ; step size $\eta > 0$.

1. For $\ell = 0, 1, \dots$:

- (a) Integrate state equation forward from t_j^0 to t_j^1 to obtain $x_j^{(\ell)}$ (see (5.2));
- (b) Integrate adjoint equation backward from t_j^1 to t_j^0 to obtain $\lambda_j^{(\ell)}$;
- (c) Compute the (total) gradient of the local objective:

$$\left(\frac{\delta \tilde{\mathcal{J}}_j}{\delta u_j} \right)^{(\ell)}(t) \leftarrow H(t)x_j^{(\ell)}(t) + R(t)u_j^{(\ell)}(t) + B^\top(t)\lambda_j^{(\ell)}(t), \quad t \in [t_j^0, t_j^1];$$

- (d) Update control: $u_j^{(\ell+1)} \leftarrow u_j^{(\ell)} - \eta \cdot \left(\frac{\delta \tilde{\mathcal{J}}_j}{\delta u_j} \right)^{(\ell)}$;
-

We provide a few clarifications regarding the derivation of the gradient (i.e., descent direction). In particular, the total gradient of the subproblem cost functional with respect to the control u_j is derived using first-order perturbation analysis [168]. The first-order optimality condition given by the PMP is numerically solved in an open-loop fashion: a control u_j is input, the states x_j and adjoint states λ_j are first subsequently solved, and the control is then updated using information of the states and adjoint states in a direction that minimizes the cost functional. In this formulation, we account for the fact that the states $x_j = x_j[u_j; p_j]$ are fully determined by both the control u_j and the initial boundary parameter p_j through the solution (5.2).

We begin by reviewing the subproblem (5.11) as a cost functional in u_j only:

$$\min_{u_j(\cdot), x_j(\cdot)} \quad \tilde{\mathcal{J}}_j[u_j; p_j, q_j] = \mathcal{J}_j[u_j, x_j[u_j; p_j]; q_j], \quad (5.56a)$$

$$\text{s.t.} \quad \dot{x}_j(t) = A(t)x_j(t) + B(t)u_j(t), \quad t \in (t_j^0, t_j^1], \quad (5.56b)$$

$$x_j(t_j^0) = p_j. \quad (5.56c)$$

Let u_j denote any control function for (5.56). Consider a perturbation at u_j in the direction of a function δu_j . The total variation (i.e. upon accounting for the variations of x_j as a result of variations in u_j) of the cost functional can be decomposed as the following:

$$\begin{aligned} \left\langle \frac{\delta \tilde{\mathcal{J}}_j[u_j; p_j, q_j]}{\delta u_j}, \delta u_j \right\rangle &:= \lim_{\epsilon \rightarrow 0} \frac{\tilde{\mathcal{J}}_j[u_j + \epsilon \cdot \delta u_j; p_j, q_j] - \tilde{\mathcal{J}}_j[u_j; p_j, q_j]}{\epsilon} \\ &= \lim_{\epsilon \rightarrow 0} \frac{\mathcal{J}_j[u_j + \epsilon \cdot \delta u_j, x_j[u_j + \epsilon \cdot \delta u_j; p_j]; q_j] - \mathcal{J}_j[u_j; x_j[u_j; p_j]; q_j]}{\epsilon} \\ &= \left\langle \frac{\partial \mathcal{J}_j[u_j, x_j; q_j]}{\partial u_j}, \delta u_j \right\rangle + \left\langle \frac{\partial \mathcal{J}_j[u_j, x_j; q_j]}{\partial x_j}, \delta x_j \right\rangle, \end{aligned} \quad (5.57)$$

where we define

$$\delta x_j := \lim_{\epsilon \rightarrow 0} \frac{x_j[u_j + \epsilon \cdot \delta u_j; p_j] - x_j[u_j; p_j]}{\epsilon}.$$

To obtain the formula of $\delta \tilde{\mathcal{J}}_j / \delta u_j$, we would like to convert the term involving δx_j into an expression involving δu_j by exploiting the structure of the PMP (5.8). We first derive a differential equation that δx_j satisfies. Let $x_j^\epsilon := x_j[u_j + \epsilon \cdot \delta u_j; p_j]$. Then, x_j^ϵ is the solution to the following system on $[t_j^0, t_j^1]$:

$$\dot{x}_j^\epsilon(t) = A(t)x_j^\epsilon(t) + B(t)(u_j(t) + \epsilon \cdot \delta u_j(t)), \quad t \in (t_j^0, t_j^1], \quad x_j^\epsilon(t_j^0) = p_j,$$

which is compared to the original system on $[t_j^0, t_j^1]$:

$$\dot{x}_j(t) = A(t)x_j(t) + B(t)u_j(t), \quad t \in (t_j^0, t_j^1], \quad x_j(t_j^0) = p_j.$$

Subtracting the above equations, dividing by ϵ , and taking the limit as $\epsilon \rightarrow 0$, we have

$$\delta \dot{x}_j(t) = A(t)\delta x_j(t) + B(t)\delta u_j(t), \quad t \in (t_j^0, t_j^1], \quad (5.58a)$$

$$\delta x_j(t_j^0) = 0. \quad (5.58b)$$

On the other hand, let $\lambda_j(t)$ denote the adjoint states satisfying the differential equation (cf. Theorem 4):

$$\begin{aligned} \dot{\lambda}_j(t) &= -A^\top(t)\lambda_j(t) - \nabla_x \left(\frac{1}{2} \begin{bmatrix} x_j(t) \\ u_j(t) \end{bmatrix}^\top \begin{bmatrix} Q(t) & H^\top(t) \\ H(t) & R(t) \end{bmatrix} \begin{bmatrix} x_j(t) \\ u_j(t) \end{bmatrix} \right) \\ &= -A^\top(t)\lambda_j(t) - Q(t)x_j(t) - H^\top(t)u_j(t), \quad t \in [t_j^0, t_j^1), \end{aligned} \quad (5.59a)$$

$$\lambda_j(t_j^1) = \nabla_x L_j(x_j(t_j^1); q_j), \quad (5.59b)$$

where L_j is as defined in (5.12). Then we take the inner product between the functions λ_j and $\delta \dot{x}_j$, and apply integration by parts to obtain

$$\langle \lambda_j, \delta \dot{x}_j \rangle = \int_{t_j^0}^{t_j^1} \lambda_j^\top(t) \delta \dot{x}_j(t) dt = \lambda_j^\top(t_j^1) \delta x_j(t_j^1) - \underbrace{\lambda_j^\top(t_j^0) \delta x_j(t_j^0)}_{=0, (5.58b)} - \langle \dot{\lambda}_j, \delta x_j \rangle. \quad (5.60)$$

Furthermore, we have

$$\begin{aligned} \langle \lambda_j, \delta \dot{x}_j \rangle &\stackrel{(5.58a)}{=} \int_{t_j^0}^{t_j^1} \lambda_j^\top(t) (A(t)\delta x_j(t) + B(t)\delta u_j(t)) dt \\ &= \int_{t_j^0}^{t_j^1} (A^\top(t)\lambda_j(t))^\top \delta x_j(t) dt + \int_{t_j^0}^{t_j^1} (B^\top(t)\lambda_j(t))^\top \delta u_j(t) dt. \end{aligned} \quad (5.61)$$

Combining (5.60) and (5.61), we obtain

$$\lambda_j^\top(t_j^1)\delta x_j(t_j^1) - \int_{t_j^0}^{t_j^1} \left(\dot{\lambda}_j(t) + A^\top(t)\lambda_j(t) \right)^\top \delta x_j(t) dt = \int_{t_j^0}^{t_j^1} \left(B^\top(t)\lambda_j(t) \right)^\top \delta u_j(t) dt. \quad (5.62)$$

Now, we revisit our main objective to evaluate (5.57). We have

$$\begin{aligned} \left\langle \frac{\partial \mathcal{J}_j[u_j, x_j; q_j]}{\partial u_j}, \delta u_j \right\rangle &:= \lim_{\epsilon \rightarrow 0} \frac{\mathcal{J}_j[u_j + \epsilon \cdot \delta u_j, x_j[u_j]; p_j] - \mathcal{J}_j[u_j; x_j[u_j]; p_j]}{\epsilon} \\ &= \int_{t_j^0}^{t_j^1} (H(t)x_j(t) + R(t)u_j(t))^\top \delta u_j(t) dt \\ \left\langle \frac{\partial \mathcal{J}_j[u_j, x_j; q_j]}{\partial x_j}, \delta x_j \right\rangle &:= \lim_{\epsilon \rightarrow 0} \frac{\mathcal{J}_j[u_j, x_j[u_j] + \epsilon \cdot \delta x_j; p_j] - \mathcal{J}_j[u_j; x_j[u_j]; p_j]}{\epsilon} \\ &= \int_{t_j^0}^{t_j^1} (Q(t)x_j(t) + H^\top(t)u_j(t))^\top \delta x_j(t) dt + \nabla_x L_j^\top(x_j(t_j^1); q_j) \delta x_j(t_j^1) \\ &\stackrel{(5.59)}{=} - \int_{t_j^0}^{t_j^1} (\dot{\lambda}_j(t) + A^\top(t)\lambda_j(t))^\top \delta x_j(t) dt + \lambda_j^\top(t_j^1) \delta x_j(t_j^1) \\ &\stackrel{(5.62)}{=} \int_{t_j^0}^{t_j^1} \left(B^\top(t)\lambda_j(t) \right)^\top \delta u_j(t) dt. \end{aligned}$$

Plugging the above display back to (5.57), we know the total variation with respect to control u_j can be evaluated in closed-form as:

$$\left\langle \frac{\delta \tilde{\mathcal{J}}_j[u_j; p_j, q_j]}{\delta u_j}, \delta u_j \right\rangle = \int_{t_j^0}^{t_j^1} \left(H(t)x_j(t) + R(t)u_j(t) + B^\top(t)\lambda_j(t) \right)^\top \delta u_j(t) dt.$$

Therefore, the gradient is given by:

$$\left(\frac{\delta \tilde{\mathcal{J}}_j[u_j; p_j, q_j]}{\delta u_j} \right) (t) = H(t)x_j(t) + R(t)u_j(t) + B^\top(t)\lambda_j(t), \quad t \in [t_j^0, t_j^1]. \quad (5.63)$$

We note that (5.63) can be computed efficiently using a forward-backward-in-time procedure (cf. Algorithm 5) for numerical integration, which supports adaptive time-stepping

and is compatible with numerous standard solvers. Other gradient-based algorithms, such as steepest descent, conjugate gradient, or BFGS in infinite-dimensional space can be similarly defined through variational analysis [122, 157].

5.8.2 Appendix: Proof of Theorem 6

The proof proceeds by directly applying Hamilton–Jacobi–Bellman (HJB) equations (cf. Theorem 5) to the parameterized optimal control problem (5.15). We define the Hamiltonian function

$$\mathcal{H}(t, x, u, \lambda; d) := \frac{1}{2} \begin{bmatrix} x \\ u \\ d \end{bmatrix}^\top \begin{bmatrix} Q(t) & H^\top(t) & G^\top(t) \\ H(t) & R(t) & W^\top(t) \\ G(t) & W(t) & 0 \end{bmatrix} \begin{bmatrix} x \\ u \\ d \end{bmatrix} + \lambda^\top (A(t)x + B(t)u + C(t)d).$$

Since the terminal cost in (5.15) is quadratic in $x(T)$, we propose a quadratic ansatz for the value function

$$J^*(t, x) := \frac{1}{2} x^\top S(t)x + x^\top v(t) + \frac{1}{2} z(t), \quad (5.64)$$

where the symmetric matrix $S : [0, T] \rightarrow \mathbb{R}^{n_x \times n_x}$ and the vectors $v : [0, T] \rightarrow \mathbb{R}^{n_x}$ and $z : [0, T] \rightarrow \mathbb{R}$ are to be determined. We know that the gradient of the value function with respect to the state yields the optimal adjoint trajectory [22, equation (4)]

$$\lambda^*(t) = \nabla_x J^*(t, x) = S(t)x + v(t). \quad (5.65)$$

Substituting the above expression into the Hamiltonian, we obtain

$$\begin{aligned} \mathcal{H}(t, x, u, \nabla_x J^*(t, x); d) &= \frac{1}{2} x^\top Q(t)x + x^\top H^\top(t)u + x^\top G^\top(t)d + \frac{1}{2} u^\top R(t)u + u^\top W^\top(t)d \\ &\quad + (S(t)x + v(t))^\top (A(t)x + B(t)u + C(t)d). \end{aligned}$$

We now proceed to compute the minimization in (5.10a) with respect to u for each $t \in [0, T]$, derive the associated equations, and validate the resulting closed-loop dynamics as stated in Theorem 6. From Assumption 3, the Hamiltonian is strongly convex in u , so the minimizer exists and is unique for every $t \in [0, T]$. Taking the gradient of the Hamiltonian with respect to u and setting it to 0 gives the following condition that the optimal control u^* must satisfy:

$$\begin{aligned} \nabla_u \mathcal{H}(t, x^*, u^*, \nabla_x J^*(t, x^*); d) &= H(t)x^*(t) + R(t)u^*(t) + W^\top(t)d(t) \\ &\quad + B^\top(t)S(t)x^*(t) + B^\top(t)v(t) = 0. \end{aligned}$$

Upon solving this equation, we obtain

$$u^*(t) = -R^{-1}(t) \left\{ H(t)x^*(t) + B^\top(t)S(t)x^*(t) + B^\top(t)v(t) + W^\top(t)d(t) \right\}. \quad (5.66)$$

Substituting (5.66) into the Hamiltonian, we obtain the minimal value. In particular, we let $\gamma(t) := H(t)x^* + B^\top(t)S(t)x^* + B^\top(t)v(t) + W^\top(t)d(t)$ and have for every $t \in [0, T]$ that

$$\begin{aligned} \mathcal{H}(t, x^*, u^*, \nabla_x J^*(t, x^*); d) &= \frac{1}{2}(x^*)^\top Q(t)x^* - (x^*)^\top H^\top(t)R^{-1}(t)\gamma(t) \\ &\quad + (x^*)^\top G^\top(t)d(t) + \frac{1}{2}\gamma^\top(t)R^{-1}(t)\gamma(t) \\ &\quad - \gamma^\top(t)R^{-1}(t)W^\top(t)d(t) + (x^*)^\top S(t)A(t)x^* \\ &\quad - (x^*)^\top S(t)B(t)R^{-1}(t)\gamma(t) + (x^*)^\top S(t)C(t)d(t) \\ &\quad + (x^*)^\top A^\top(t)v(t) - \gamma^\top(t)R^{-1}(t)B^\top(t)v(t) + d^\top C^\top(t)v. \end{aligned}$$

To clearly isolate the above terms and match them with those defined in (5.64), we let

$$\mathcal{H}(t, x^*, u^*, \nabla_x J^*(t, x^*); d) = \mathcal{Q}(t, x^*; d) + \mathcal{L}(t, x^*; d) + \mathcal{C}(t, x^*; d),$$

where $\mathcal{Q}, \mathcal{L}, \mathcal{C}$ respectively denote terms that are quadratic, linear, and constant in optimal states x^* , given by

$$\begin{aligned}
\mathcal{Q}(t, x^*; d) &= (x^*)^\top \left[\frac{1}{2} Q(t) - H^\top(t) R^{-1}(t) \left\{ H(t) + B^\top(t) S(t) \right\} \right. \\
&\quad + \frac{1}{2} \left\{ H(t) + B^\top(t) S(t) \right\}^\top R^{-1}(t) \left\{ H(t) + B^\top(t) S(t) \right\} + S(t) A(t) \\
&\quad \left. - S(t) B(t) R^{-1}(t) \left\{ H(t) + B^\top(t) S(t) \right\} \right] x^*, \\
\mathcal{L}(t, x^*; d) &= (x^*)^\top \left[-H^\top(t) R^{-1}(t) \left\{ B^\top(t) v(t) + W^\top(t) d(t) \right\} + G^\top(t) d(t) \right. \\
&\quad - \left\{ H(t) + B^\top(t) S(t) \right\}^\top R^{-1}(t) W^\top(t) d(t) \\
&\quad + \left\{ H(t) + B^\top(t) S(t) \right\}^\top R^{-1}(t) \left\{ B^\top(t) v(t) + W^\top(t) d(t) \right\} \\
&\quad - S(t) B(t) R^{-1}(t) \left\{ B^\top(t) v(t) + W^\top(t) d(t) \right\} + S(t) C(t) d(t) \\
&\quad \left. + A^\top(t) v(t) - \left\{ H(t) + B^\top(t) S(t) \right\}^\top R^{-1}(t) B^\top(t) v(t) \right], \\
\mathcal{C}(t, x^*; d) &= \frac{1}{2} \left\{ W^\top(t) d(t) + B^\top(t) v(t) \right\}^\top R^{-1}(t) \left\{ W^\top(t) d(t) + B^\top(t) v(t) \right\} \\
&\quad + d^\top(t) C^\top(t) v(t) - \left\{ W^\top(t) d(t) + B^\top(t) v(t) \right\}^\top R^{-1}(t) W^\top(t) d(t) \\
&\quad - \left\{ W^\top(t) d(t) + B^\top(t) v(t) \right\}^\top R^{-1}(t) B^\top(t) v(t).
\end{aligned}$$

By (5.64), we compute the time derivative of the value function as

$$\frac{\partial J^*(t, x)}{\partial t} = \frac{1}{2} x^\top \dot{S}(t) x + x^\top \dot{v}(t) + \frac{1}{2} \dot{z}(t).$$

Now, we substitute the above expression into the HJB equation and obtain

$$\frac{\partial J^*(t, x^*)}{\partial t} + \mathcal{Q}(t, x^*; d) + \mathcal{L}(t, x^*; d) + \mathcal{C}(t, x^*; d) = 0.$$

By matching terms according to their degree in x^* , a sufficient condition for this identity to hold for all x^* is that each group of terms is 0 independently. This yields a system of

differential equations for S , v , and z . We now present the equations for S and v , which are responsible for determining the optimal control and closed-loop dynamics. The equation for z is omitted (i.e. $0.5\dot{z}(t) + \mathcal{C}(t, x^*; d) = 0$), as it does not influence the state x^* , the adjoint λ^* , or the optimal control u^* . In particular, by requiring $0.5(x^*)^\top \dot{S}(t)x^* + \mathcal{Q}(t, x^*; d) = 0$, we obtain

$$(x^*)^\top \left(\frac{1}{2}\dot{S} + \frac{1}{2}Q - H^\top R^{-1}H - H^\top R^{-1}B^\top S + \frac{1}{2}H^\top R^{-1}H + \frac{1}{2}H^\top R^{-1}B^\top S \right. \\ \left. + \frac{1}{2}SBR^{-1}H + \frac{1}{2}SBR^{-1}B^\top S + \frac{1}{2}A^\top S + \frac{1}{2}SA - SBR^{-1}H - SBR^{-1}B^\top S \right) x^* = 0,$$

where we have used the symmetry of S . Simplifying this expression and canceling the terms, applying (5.10b), we have $S(T) = Q_T$ and

$$\dot{S} + Q - H^\top R^{-1}H - H^\top R^{-1}B^\top S - SBR^{-1}H - SBR^{-1}B^\top S + A^\top S + SA = 0, \quad t \in [0, T]. \quad (5.67)$$

Similarly, for the linear terms in x^* , by requiring $(x^*)^\top \dot{v}(t) + \mathcal{L}(t, x^*; d) = 0$, we obtain

$$(x^*)^\top \left(\dot{v}(t) + \left\{ A(t) - B(t)R^{-1}(t)H(t) - B(t)R^{-1}(t)B^\top(t)S(t) \right\}^\top v(t) \right. \\ \left. + \left\{ G(t) + C^\top(t)S(t) - W(t)R^{-1}(t)(B^\top(t)S(t) + H(t)) \right\}^\top d(t) \right) = 0.$$

Therefore, we apply (5.10b), let $v(T) = G_T^\top d_T$, and have

$$\dot{v}(t) + \left\{ A(t) - B(t)R^{-1}(t)H(t) - B(t)R^{-1}(t)B^\top(t)S(t) \right\}^\top v(t) \\ + \left\{ G(t) + C^\top(t)S(t) - W(t)R^{-1}(t)(B^\top(t)S(t) + H(t)) \right\}^\top d(t) = 0, \quad t \in [0, T]. \quad (5.68)$$

Combining (5.65), (5.66), (5.67), and (5.68) together, we complete the proof.

5.8.3 Proof of Lemma 6

The existence and uniqueness of the solution S is known from Lemma 4. As preparation, we review a controllability result for a closely related system, whose properties will become relevant in establishing the lower bound of the matrix Riccati solution.

Lemma 9. *Let A, B, H, R be the matrices satisfying Assumption 1–3, then the pair $-(A - BR^{-1}H)^\top, I$ satisfies UCC on $[0, T - \sigma]$ in the sense of Assumption 2. In particular, there exist constants $\tilde{\alpha}_0(\sigma), \tilde{\alpha}_1(\sigma), \tilde{\beta}_0(\sigma), \tilde{\beta}_1(\sigma) > 0$, depending on σ from Assumption 2 but independent of T , such that for every $t_0 \in [0, T - \sigma]$, the choice $t_1 = t_0 + \sigma \in [0, T]$ ensures that the following bounds hold:*

$$\begin{aligned}\tilde{\alpha}_0(\sigma)I &\preceq W_{-(A+BR^{-1}H)^\top, I}(t_0, t_1) \preceq \tilde{\alpha}_1(\sigma)I, \\ \tilde{\beta}_0(\sigma)I &\preceq \Phi_{-(A+BR^{-1}H)^\top}(t_1, t_0)W_{-(A+BR^{-1}H)^\top, I}(t_0, t_1)\Phi_{-(A+BR^{-1}H)^\top}^\top(t_1, t_0) \preceq \tilde{\beta}_1(\sigma)I.\end{aligned}$$

(The expressions of $\tilde{\alpha}_0(\sigma), \tilde{\alpha}_1(\sigma), \tilde{\beta}_0(\sigma), \tilde{\beta}_1(\sigma) > 0$ are provided in (5.69) and (5.70).)

Proof. By the definition of controllability Gramian in (5.3), we have

$$\begin{aligned}W_{-(A+BR^{-1}H)^\top, I}(t_0, t_1) &= \int_{t_0}^{t_1} \Phi_{-(A+BR^{-1}H)^\top}(t_0, s)\Phi_{-(A+BR^{-1}H)^\top}^\top(t_0, s)ds \\ &= \int_{t_0}^{t_1} \Phi_{A+BR^{-1}H}^\top(s, t_0)\Phi_{A+BR^{-1}H}(s, t_0)ds.\end{aligned}$$

Let $0 \neq v \in \mathbb{R}^{n_x}$ be any vector, we first consider the upper bound and have

$$\begin{aligned}
v^\top W_{-(A+BR^{-1}H)^\top, I}(t_0, t_1)v &= \int_{t_0}^{t_1} v^\top \Phi_{A+BR^{-1}H}^\top(s, t_0) \Phi_{A+BR^{-1}H}(s, t_0) v ds \\
&= \int_{t_0}^{t_1} \|\Phi_{A+BR^{-1}H}(s, t_0)v\|^2 ds \\
&= \int_{t_0}^{t_0+\sigma} \|\Phi_{A+BR^{-1}H}(s, t_0)v\|^2 ds \\
&\leq \int_{t_0}^{t_0+\sigma} e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(s-t_0)} \|v\|^2 ds \\
&= \frac{\|v\|^2}{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)} \left(e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)\sigma} - 1 \right),
\end{aligned}$$

where the inequality is due to (5.23). For the lower bound, for any $0 \neq v \in \mathbb{R}^{n_x}$, we have by the bound derived in (5.24) that

$$\begin{aligned}
v^\top W_{-(A+BR^{-1}H)^\top, I}(t_0, t_1)v &= \int_{t_0}^{t_1} \|\Phi_{A+BR^{-1}H}(s, t_0)v\|^2 ds \\
&\geq \int_{t_0}^{t_1} e^{-2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(s-t_0)} \|v\|^2 ds \\
&= \int_{t_0}^{t_0+\sigma} e^{-2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(s-t_0)} \|v\|^2 ds \\
&= \frac{\|v\|^2}{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)} \left(1 - e^{-2(\lambda_A + \lambda_B \lambda_H / \gamma_R)\sigma} \right).
\end{aligned}$$

Combining the above two displays and defining

$$\tilde{\alpha}_0(\sigma) := \frac{1}{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)} \left(1 - e^{-2(\lambda_A + \lambda_B \lambda_H / \gamma_R)\sigma} \right), \quad (5.69a)$$

$$\tilde{\alpha}_1(\sigma) := \frac{1}{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)} \left(e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)\sigma} - 1 \right), \quad (5.69b)$$

we prove the first part of the statement. Furthermore, we note from (5.23) and (5.24) that

for any $0 \neq v \in \mathbb{R}^{n_x}$,

$$\begin{aligned} v^\top \Phi_{-(A+BR^{-1}H)^\top}(t_1, t_0) W_{-(A+BR^{-1}H)^\top, I}(t_0, t_1) \Phi_{-(A+BR^{-1}H)^\top}^\top(t_1, t_0) v \\ \leq \tilde{\alpha}_1(\sigma) \|\Phi_{-(A+BR^{-1}H)^\top}^\top(t_1, t_0) v\|^2 \leq \tilde{\alpha}_1(\sigma) e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R) \sigma} \|v\|^2, \end{aligned}$$

and

$$\begin{aligned} v^\top \Phi_{-(A+BR^{-1}H)^\top}(t_1, t_0) W_{-(A+BR^{-1}H)^\top, I}(t_0, t_1) \Phi_{-(A+BR^{-1}H)^\top}^\top(t_1, t_0) v \\ \geq \tilde{\alpha}_0(\sigma) \|\Phi_{-(A+BR^{-1}H)^\top}^\top(t_1, t_0) v\|^2 \geq \tilde{\alpha}_0(\sigma) e^{-2(\lambda_A + \lambda_B \lambda_H / \gamma_R) \sigma} \|v\|^2. \end{aligned}$$

Thus, we define

$$\tilde{\beta}_0(\sigma) := \tilde{\alpha}_0(\sigma) e^{-2(\lambda_A + \lambda_B \lambda_H / \gamma_R) \sigma} \quad \text{and} \quad \tilde{\beta}_1(\sigma) := \tilde{\alpha}_1(\sigma) e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R) \sigma}, \quad (5.70)$$

and conclude the proof. $\square$

We note that the statement of Lemma 9 is slightly weaker than those of Assumption 2 and Lemma 5 in that at any time $t_0 \in [0, T - \sigma]$, any state $v \in \mathbb{R}^{n_x}$ is steerable to 0 precisely at time $t_0 + \sigma$ for the system $(-(A + BF)^\top, I)$ rather than some other time $t_1 < t_0 + \sigma$. However, since we only require the existence of a time in $[t_0, t_0 + \sigma]$ for any $t_0 \in [0, T - \sigma]$, the choice $t_1 = t_0 + \sigma$ remains valid and does not impact the proof of Lemma 6.

We now present the main proof of Lemma 6. In what follows, we first establish the desired bounds on the controllable segment $[0, T - \sigma]$.

• **Case 1:** $t_0 \in [0, T - \sigma]$. By Lemmas 2 and 5, we know for any initial time $t_0 \in [0, T - \sigma]$ and any initial state $0 \neq x_{t_0} \in \mathbb{R}^{n_x}$, there exists a control input that steers the system $(A - BR^{-1}H, B)$ to the origin by some final time $t_1 \in [t_0, t_0 + \sigma] \subseteq [0, T]$. An explicit

control that accomplishes this is given by the following

$$u(t; x_{t_0}) = \begin{cases} -B^\top(t)\Phi_{A-BR^{-1}H}^\top(t_0, t)W_{A-BR^{-1}H, B}^{-1}(t_0, t_1)x_{t_0}, & t \in [t_0, t_1], \\ 0, & t \in (t_1, T]. \end{cases} \quad (5.71)$$

Let us denote $x(t; x_{t_0})$, $t \in [t_0, T]$, to be state trajectory of the system $(A - BR^{-1}H, B)$ with the above control input (5.71); and denote $x^* : [t_0, T] \rightarrow \mathbb{R}^{n_x}$ and $u^* : [t_0, T] \rightarrow \mathbb{R}^{n_u}$ to be the optimal state and control for the truncated, shifted linear-quadratic problem (5.20) on $[t_0, T]$ with initial state x_{t_0} . Then, by definition of optimal cost-to-go (an analogy of (5.9) with different system; and an analogy of (5.64) with $v(t) = 0$, $z(t) = 0$), we have

$$\begin{aligned} \frac{1}{2}x_{t_0}^\top S(t_0)x_{t_0} &= \frac{1}{2} \int_{t_0}^T \left(x^{*\top}(t)(Q(t) - H^\top(t)R^{-1}(t)H(t))x^*(t) + u^{*\top}(t)R(t)u^*(t) \right) dt \\ &\quad + \frac{1}{2}x^{*\top}(T)Q_Tx^*(T). \end{aligned}$$

Thus, by the optimality of (x^*, u^*) on $[t_0, T]$, we have

$$\begin{aligned} \frac{1}{2}x_{t_0}^\top S(t_0)x_{t_0} &\leq \frac{1}{2} \int_{t_0}^T \begin{bmatrix} x(t; x_{t_0}) \\ u(t; x_{t_0}) \end{bmatrix}^\top \begin{bmatrix} Q(t) - H^\top(t)R^{-1}(t)H(t) & 0 \\ 0 & R(t) \end{bmatrix} \begin{bmatrix} x(t; x_{t_0}) \\ u(t; x_{t_0}) \end{bmatrix} dt \\ &\quad + \frac{1}{2}x^\top(T; x_{t_0})Q_Tx(T; x_{t_0}) \\ &= \frac{1}{2} \int_{t_0}^{t_1} x^\top(t; x_{t_0})(Q(t) - H^\top(t)R^{-1}(t)H(t))x(t; x_{t_0}) \\ &\quad + u^\top(t; x_{t_0})R(t)u(t; x_{t_0})dt, \end{aligned}$$

where we use the fact that $u(t, x_{t_0}) = 0$ and $x(t; x_{t_0}) = 0$ for all $t > t_1$. For the second term

on the right hand side, we have for $t \in [t_0, t_1]$,

$$\begin{aligned}
& u^\top(t; x_{t_0})R(t)u(t; x_{t_0}) \\
& \leq \lambda_R \|u(t; x_{t_0})\|^2 \stackrel{(5.71)}{\leq} \lambda_R \|B(t)\|^2 \|\Phi_{A-BR^{-1}H}(t_0, t)\|^2 \|W_{A-BR^{-1}H, B}^{-1}(t_0, t_1)\|^2 \|x_{t_0}\|^2 \\
& \leq \frac{\lambda_R \lambda_B^2}{(\alpha'_0(\sigma))^2} \|x_{t_0}\|^2 \|\Phi_{(A-BR^{-1}H)^\top}(t, t_0)\|^2 \stackrel{(5.23)}{\leq} \frac{\lambda_R \lambda_B^2}{(\alpha'_0(\sigma))^2} e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(t-t_0)} \|x_{t_0}\|^2,
\end{aligned} \tag{5.72}$$

where the third inequality is due to Lemma 5. For the first term on the right hand side, we substitute the control (5.71) into (5.2), and have for $t \in [t_0, t_1]$,

$$\begin{aligned}
\|x(t; x_{t_0})\| &= \left\| \Phi_{A-BR^{-1}H}(t, t_0) \left(I - W_{A-BR^{-1}H, B}(t_0, t) W_{A-BR^{-1}H, B}^{-1}(t_0, t_1) \right) x_{t_0} \right\| \\
&\stackrel{(5.29)}{\leq} \left(1 + \frac{\alpha'_1(\sigma)}{\alpha'_0(\sigma)} \right) \|x_{t_0}\| \cdot \|\Phi_{A-BR^{-1}H}(t, t_0)\| \\
&\stackrel{(5.23)}{\leq} \left(1 + \frac{\alpha'_1(\sigma)}{\alpha'_0(\sigma)} \right) e^{(\lambda_A + \lambda_B \lambda_H / \gamma_R)(t-t_0)} \|x_{t_0}\|,
\end{aligned} \tag{5.73}$$

where the second inequality is also due to Lemma 5. Combining the above three displays together, and noting that $\|Q(t) - H^\top(t)R^{-1}(t)H(t)\| \leq \lambda_Q + \lambda_H^2/\gamma_R$, we obtain

$$x_{t_0}^\top S(t_0) x_{t_0} \leq \left\{ \left(\lambda_Q + \frac{\lambda_H^2}{\gamma_R} \right) \left(1 + \frac{\alpha'_1(\sigma)}{\alpha'_0(\sigma)} \right)^2 + \frac{\lambda_R \lambda_B^2}{(\alpha'_0(\sigma))^2} \right\} \|x_{t_0}\|^2 \times \dots \tag{5.74}$$

$$\begin{aligned}
& \int_{t_0}^{t_1} e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(t-t_0)} dt \\
& \leq \left\{ \left(\lambda_Q + \frac{\lambda_H^2}{\gamma_R} \right) \left(1 + \frac{\alpha'_1(\sigma)}{\alpha'_0(\sigma)} \right)^2 + \frac{\lambda_R \lambda_B^2}{(\alpha'_0(\sigma))^2} \right\} \|x_{t_0}\|^2 \times \dots \\
& \frac{e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)\sigma} - 1}{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)}.
\end{aligned} \tag{5.75}$$

To establish the lower bound, we know from Lemma 4 that the Riccati matrix S remains positive definite for all $t \in [0, T]$. Thus, let us define $U(t) := S^{-1}(t)$ and have $\dot{U}(t) =$

$-S^{-1}(t)\dot{S}(t)S^{-1}(t)$. Substituting this relation into the Riccati equation (5.17), we have

$$\begin{aligned}\dot{U}(t) &= (A(t) - B(t)R^{-1}(t)H(t))U(t) + U(t)(A(t) - B(t)R^{-1}(t)H(t))^{\top} \\ &\quad + U(t)(Q(t) - H^{\top}(t)R^{-1}(t)H(t))U(t) - B(t)R^{-1}(t)B^{\top}(t), \quad t \in [0, T],\end{aligned}\quad (5.76a)$$

$$U(T) = Q_T^{-1}. \quad (5.76b)$$

Note that (5.76) itself is a matrix Riccati equation. The existence and uniqueness follow from that of S . This Riccati equation corresponds to a dual linear-quadratic optimal control problem on the interval $[0, T]$, formulated as [111, (3.12-14)]

$$\begin{aligned}\min_{y(\cdot), v(\cdot)} \quad & \frac{1}{2} \int_0^T \begin{bmatrix} y(t) \\ v(t) \end{bmatrix}^{\top} \begin{bmatrix} B(t)R^{-1}(t)B^{\top}(t) & 0 \\ 0 & (Q(t) - H^{\top}(t)R^{-1}(t)H(t))^{-1} \end{bmatrix} \begin{bmatrix} y(t) \\ v(t) \end{bmatrix} dt \\ & + \frac{1}{2} y^{\top}(T)Q_T^{-1}y(T),\end{aligned}$$

$$\text{s.t.} \quad \dot{y}(t) = -(A(t) - B(t)R^{-1}(t)H(t))^{\top}y(t) + v(t), \quad t \in (0, T], \quad (5.77a)$$

$$y(t_0) = y_0 \in \mathbb{R}^{n_x}. \quad (5.77b)$$

The state dynamics are now given by the pair $(-(A - BR^{-1}H)^{\top}, I)$, and the analysis of this system mirrors that of the shifted OCP in Lemma 3. In particular, for any $t_0 \in [0, T - \sigma]$, we apply Lemma 9 to construct an explicit zero-steering control over the horizon $[t_0, t_0 + \sigma]$ that drives any initial state $0 \neq y_{t_0} \in \mathbb{R}^{n_x}$ to the origin within time σ from t_0 , given by

$$v(t; y_{t_0}) := \begin{cases} -\Phi_{-(A-BR^{-1}H)^{\top}}^{\top}(t_0, t) W_{-(A-BR^{-1}H)^{\top}, I}^{-1}(t_0, t_0 + \sigma) y_{t_0}, & t \in [t_0, t_0 + \sigma], \\ 0, & t \in (t_0 + \sigma, T]. \end{cases}$$

The corresponding state with the above control trajectory is given by (5.2): for $t \in [t_0, t_0 + \sigma]$,

$$y(t; y_{t_0}) = \Phi_{-(A-BR^{-1}H)^{\top}}(t, t_0) \left(I - W_{-(A-BR^{-1}H)^{\top}, I}^{-1}(t_0, t) W_{-(A-BR^{-1}H)^{\top}, I}^{-1}(t_0, t_0 + \sigma) \right) y_{t_0},$$

and for $t \in (t_0 + \sigma, T]$, $y(t; y_{t_0}) = 0$. With the above specific state-control trajectories, we have

$$\frac{1}{2} y_{t_0}^\top U(t_0) y_{t_0} \leq \frac{1}{2} \int_{t_0}^{t_0+\sigma} \begin{bmatrix} y \\ v \end{bmatrix}^\top \begin{bmatrix} B(t)R^{-1}(t)B^\top(t) & 0 \\ 0 & (Q(t) - H^\top(t)R^{-1}(t)H(t))^{-1} \end{bmatrix} \begin{bmatrix} y \\ v \end{bmatrix} dt.$$

where for notational simplicity above, we have $v = v(t; y_{t_0})$, $y = y(t; t_0)$. Following the same derivation as in (5.72), we have from Assumption 3 and Lemma 9 that

$$v^\top(t; y_{t_0})(Q(t) - H^\top(t)R^{-1}(t)H(t))^{-1}v(t; y_{t_0}) \leq \frac{1}{\gamma_Q(\tilde{\alpha}_0(\sigma))^2} e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(t-t_0)} \|y_{t_0}\|^2.$$

Following the same derivation as in (5.73), we have

$$\|y(t; y_{t_0})\| \leq \left(1 + \frac{\tilde{\alpha}_1(\sigma)}{\tilde{\alpha}_0(\sigma)}\right) \|y_{t_0}\| e^{(\lambda_A + \lambda_B \lambda_H / \gamma_R)(t-t_0)}.$$

Combining the above three displays, we obtain

$$\begin{aligned} y_{t_0}^\top S^{-1}(t_0) y_{t_0} &= y_{t_0}^\top U(t_0) y_{t_0} \\ &\leq \left\{ \frac{\lambda_B^2}{\gamma_R} \left(1 + \frac{\tilde{\alpha}_1(\sigma)}{\tilde{\alpha}_0(\sigma)}\right)^2 + \frac{1}{\gamma_Q(\tilde{\alpha}_0(\sigma))^2} \right\} \|y_{t_0}\|^2 \int_{t_0}^{t_0+\sigma} e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(t-t_0)} dt \\ &= \left\{ \frac{\lambda_B^2}{\gamma_R} \left(1 + \frac{\tilde{\alpha}_1(\sigma)}{\tilde{\alpha}_0(\sigma)}\right)^2 + \frac{1}{\gamma_Q(\tilde{\alpha}_0(\sigma))^2} \right\} \|y_{t_0}\|^2 \cdot \frac{e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)\sigma} - 1}{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)}. \end{aligned} \tag{5.78}$$

• **Case 2:** $t_0 \in (T - \sigma, T]$. We now discuss the case without uniform controllability. The bounds will similarly be derived from the relationship between the Riccati matrix and optimal cost, and without the ability to steer an arbitrary initial state at t_0 to 0 due to the lack of complete controllability, we simply apply the zero control trajectory $\tilde{u}(t) = 0$, $t \in [t_0, T]$. With the zero control, the state trajectory evolves according to the homogeneous dynamics,

i.e., $\tilde{x}(t) = \Phi_{A-BR^{-1}H}(t, t_0)x_{t_0}$, $t \in [t_0, T]$. With the above state-control trajectories $(\tilde{x}, \tilde{u})$, we have

$$\begin{aligned}
x_{t_0}^\top S(t_0)x_{t_0} &\leq \int_{t_0}^T \begin{bmatrix} \tilde{x}(t) \\ \tilde{u}(t) \end{bmatrix}^\top \begin{bmatrix} Q(t) - H^\top(t)R^{-1}(t)H(t) & 0 \\ 0 & R(t) \end{bmatrix} \begin{bmatrix} \tilde{x}(t) \\ \tilde{u}(t) \end{bmatrix} dt + \tilde{x}(T)^\top Q_T \tilde{x}(T) \\
&\leq \left(\lambda_Q + \frac{\lambda_H^2}{\gamma_R} \right) \int_{t_0}^T \|\tilde{x}(t)\|^2 dt + \lambda_Q \|\tilde{x}(T)\|^2 \\
&\stackrel{(5.23)}{\leq} \left(\lambda_Q + \frac{\lambda_H^2}{\gamma_R} \right) \|x_{t_0}\|^2 \int_{t_0}^T e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(t-t_0)} dt \\
&\quad + \lambda_Q \|x_{t_0}\|^2 e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(T-t_0)} \\
&\leq \left(\lambda_Q + \frac{\lambda_H^2}{\gamma_R} \right) \|x_{t_0}\|^2 \frac{e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)\sigma} - 1}{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)} + \lambda_Q \|x_{t_0}\|^2 e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)\sigma} \\
&\leq \left(\lambda_Q + \frac{\lambda_H^2}{\gamma_R} \right) \left(1 + \frac{1}{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)} \right) e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)\sigma} \cdot \|x_{t_0}\|^2. \quad (5.79)
\end{aligned}$$

To obtain the lower bound, we consider again the inverse Riccati equation (5.76) and corresponding optimal control problem (5.77). In particular, for any $t_0 \in (T - \sigma, T]$ and any $y_{t_0} \in \mathbb{R}^{n_x}$, we apply the zero control $\tilde{v}(t) = 0$, $t \in [t_0, T]$, and the state trajectory becomes $\tilde{y}(t) = \Phi_{-(A-BR^{-1}H)^\top}(t, t_0)y_{t_0}$. Therefore, the bound of $U(t) = S^{-1}(t)$ on $(T - \sigma, T]$

becomes

$$\begin{aligned}
& y_{t_0}^\top S^{-1}(t_0) y_{t_0} = y_{t_0}^\top U(t_0) y_{t_0} \\
& \leq \int_{t_0}^T \begin{bmatrix} \tilde{y}(t) \\ \tilde{v}(t) \end{bmatrix}^\top \begin{bmatrix} B(t)R^{-1}(t)B^\top(t) & 0 \\ 0 & (Q(t) - H^\top(t)R^{-1}(t)H(t))^{-1} \end{bmatrix} \begin{bmatrix} \tilde{y}(t) \\ \tilde{v}(t) \end{bmatrix} dt \\
& + \tilde{y}^\top(T) Q_T^{-1} \tilde{y}(T) \\
& \leq \frac{\lambda_B^2}{\gamma_R} \int_{t_0}^T \|\tilde{y}(t)\|^2 dt + \frac{1}{\gamma_Q} \|\tilde{y}(T)\|^2 \stackrel{(5.23)}{\leq} \frac{\lambda_B^2}{\gamma_R} \|y_{t_0}\|^2 \int_{t_0}^T e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(t-t_0)} dt \\
& + \frac{\|y_{t_0}\|^2}{\gamma_Q} e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)(T-t_0)} \\
& \leq \left\{ \frac{\lambda_B^2}{2\gamma_R (\lambda_A + \lambda_B \lambda_H / \gamma_R)} + \frac{1}{\gamma_Q} \right\} e^{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)\sigma} \cdot \|y_{t_0}\|^2. \tag{5.80}
\end{aligned}$$

Finally, we combine (5.74), (5.78), (5.79), and (5.80), and define

$$\begin{aligned}
c_0(\sigma) := \min & \left\{ 2 \left(\lambda_A + \frac{\lambda_B \lambda_H}{\gamma_R} \right) \left(\frac{\lambda_B^2}{\gamma_R} \left(1 + \frac{\tilde{\alpha}_1(\sigma)}{\tilde{\alpha}_0(\sigma)} \right)^2 + \frac{1}{\gamma_Q (\tilde{\alpha}_0(\sigma))^2} \right)^{-1}, \right. \\
& \left. \left(\frac{\lambda_B^2}{2\gamma_R (\lambda_A + \lambda_B \lambda_H / \gamma_R)} + \frac{1}{\gamma_Q} \right)^{-1} \right\} \cdot \exp(-2(\lambda_A + \lambda_B \lambda_H / \gamma_R) \sigma); \tag{5.81a}
\end{aligned}$$

$$\begin{aligned}
c_1(\sigma) := \max & \left\{ \left(\left(\lambda_Q + \frac{\lambda_H^2}{\gamma_R} \right) \left(1 + \frac{\alpha'_1(\sigma)}{\alpha'_0(\sigma)} \right)^2 + \frac{\lambda_R \lambda_B^2}{(\alpha'_0(\sigma))^2} \right) \frac{1}{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)}, \right. \\
& \left. \left(\lambda_Q + \frac{\lambda_H^2}{\gamma_R} \right) \left(1 + \frac{1}{2(\lambda_A + \lambda_B \lambda_H / \gamma_R)} \right) \right\} \exp(2(\lambda_A + \lambda_B \lambda_H / \gamma_R) \sigma), \tag{5.81b}
\end{aligned}$$

where $\alpha'_0(\sigma)$, $\alpha'_1(\sigma)$ are from Lemma 5, and $\tilde{\alpha}_0(\sigma)$, $\tilde{\alpha}_1(\sigma)$ are from Lemma 9. This completes the proof.

CHAPTER 6

FUTURE WORK

This dissertation opens up several avenues for further explorations. In Chapter 2, a natural future direction lies in the development of principled closures for quantities of interest that evolve under the dynamics of other systems. While the present framework captures macroscopic behavior, nonlinear dependence on states and spatial variables requires more sophisticated parameterizations. Future work will consider simultaneous identification of advection and diffusion coefficients within constrained optimization or Bayesian inversion formulations. To address the curse of dimensionality inherent in high-dimensional advection problems, tensor-network representations offer a promising path to reduce computational cost while preserving essential dynamics.

Chapter 3 and Chapter 4 together suggest a broader research direction for uncertainty quantification in multiscale and stochastic networked systems. The RoPDF approach raises theoretical questions regarding convergence rates and train-test error decomposition, which invites a joint analysis from both the perspective of statistical learning theory and stability of dynamical systems. A principled study of how assumptions on conditional expectations, noise processes, and regression methods affect approximation quality would clarify the method’s expected convergence guarantees. On the algorithmic side, deep neural architectures with Fourier features or long-short temporal memory could provide robust extrapolation in regimes with sparse observations, while structured low-complexity representations such as tensor-networks or generative flows may counteract dimensionality in high-dimensional observables. These directions are particularly relevant for probabilistic characterization of extreme events in power systems, where extending RoPDF to vector-valued quantities of interest could enable realistic prediction of cascading failures in larger grids.

Finally, Chapter 5 highlights opportunities in distributed optimal control. Extending the

overlapping Schwarz framework from linear-quadratic settings to nonlinear optimal control would be the natural next step, though it requires new convexification strategies in infinite-dimensions, which is left for future work. Beyond gradient-based local solvers, incorporating second-order methods or direct Riccati-based techniques could improve robustness and efficiency of the subproblem solutions. Symplectic time-stepping methods are also attractive and could offer further parallelization due to preserved subproblem structures. More broadly, the overlapping Schwarz method could inspire distributed learning algorithms by leveraging its convergence guarantees in locally linear-quadratic control formulations, including neural ODE [40] and continuous-in-depth neural networks [175]. These connections point toward interactions between domain decomposition and control theory interpretations of deep learning.

CHAPTER 7

REFERENCES

- [1] Assyr Abdulle, E Weinan, Björn Engquist, and Eric Vanden-Eijnden. The heterogeneous multiscale method. *Acta Numerica*, 21:1–87, April 2012. ISSN 1474-0508. doi:[10.1017/s0962492912000025](https://doi.org/10.1017/s0962492912000025).
- [2] Muhammad Adeen and Federico Milano. Modeling of correlated stochastic processes for the transient stability analysis of power systems. *IEEE Transactions on Power Systems*, 36(5):4445–4456, September 2021. ISSN 1558-0679. doi:[10.1109/tpwrs.2021.3066453](https://doi.org/10.1109/tpwrs.2021.3066453).
- [3] Hiroshi Akima. A method of bivariate interpolation and smooth surface fitting based on local procedures. *Communications of the ACM*, 17(1):18–20, January 1974. ISSN 1557-7317. doi:[10.1145/360767.360779](https://doi.org/10.1145/360767.360779).
- [4] Abdulrahman Alawadhi, Francesca Boso, and Daniel M. Tartakovsky. Method of distributions for water hammer equations with uncertain parameters. *Water Resources Research*, 54(11):9398–9411, November 2018. ISSN 1944-7973. doi:[10.1029/2018wr023383](https://doi.org/10.1029/2018wr023383).
- [5] Mustaffa Alfatlawi and Vaibhav Srivastava. An incremental approach to online dynamic mode decomposition for time-varying systems with applications to eeg data modeling. *Journal of Computational Dynamics*, 7(2):209–241, 2020. ISSN 2158-2505. doi:[10.3934/jcd.2020009](https://doi.org/10.3934/jcd.2020009).
- [6] Brian DO Anderson and John B Moore. *Optimal control: linear quadratic methods*. Courier Corporation, 2007.
- [7] Edward Anderson, Michael Ferris, Andrew Philpott, Mihai Anitescu, Peter Cramton, Sijia Geng, Richard Green, Tito Homem-de Mello, Olivier Huber, Vincent Leclère, and

- Ramteen Sioshansi. Ten challenges for mathematical modeling of the green-energy transition. *Current Sustainable/Renewable Energy Reports*, 12(1), September 2025. ISSN 2196-3010. doi:[10.1007/s40518-025-00274-9](https://doi.org/10.1007/s40518-025-00274-9).
- [8] Paul M. Anderson and A. A. Fouad. *Multimachine Systems with Constant Impedance Loads*, chapter 3, pages 368–397. IEEE, 2009. ISBN 0470545577. doi:[10.1109/9780470545577.ch9](https://doi.org/10.1109/9780470545577.ch9).
- [9] Harbir Antil and Dmitriy Leykekhman. *A Brief Introduction to PDE-Constrained Optimization*, pages 3–40. Springer New York, 2018. ISBN 9781493986361. doi:[10.1007/978-1-4939-8636-1_1](https://doi.org/10.1007/978-1-4939-8636-1_1).
- [10] A.C. Antoulas. Approximation of large-scale dynamical systems: An overview. *IFAC Proceedings Volumes*, 37(11):19–28, July 2004. ISSN 1474-6670. doi:[10.1016/s1474-6670\(17\)31584-7](https://doi.org/10.1016/s1474-6670(17)31584-7).
- [11] Özgür Asar, Ozlem Ilk, and Osman Dag. Estimating box-cox power transformation parameter via goodness-of-fit tests. *Communications in Statistics - Simulation and Computation*, 46(1):91–105, December 2014. ISSN 1532-4141. doi:[10.1080/03610918.2014.957839](https://doi.org/10.1080/03610918.2014.957839).
- [12] Michael Athans and Peter L Falb. *Optimal control: an introduction to the theory and its applications*. Courier Corporation, 2013.
- [13] Giles Auchmuty. Variational principles for eigenvalues of nonsymmetric matrices. *SIAM Journal on Matrix Analysis and Applications*, 10(1):105–117, January 1989. ISSN 1095-7162. doi:[10.1137/0610008](https://doi.org/10.1137/0610008).
- [14] Peter J. Baddoo, Benjamin Herrmann, Beverley J. McKeon, and Steven L. Brunton. Kernel learning for robust dynamic mode decomposition: linear and nonlinear disambiguation optimization. *Proceedings of the Royal Society A: Mathematical*

- ical, Physical and Engineering Sciences*, 478(2260), April 2022. ISSN 1471-2946. doi:[10.1098/rspa.2021.0830](https://doi.org/10.1098/rspa.2021.0830).
- [15] Emmanuel Nicholas Barron and Robert Jensen. The pontryagin maximum principle from dynamic programming and viscosity solutions to first-order partial differential equations. *Transactions of the American Mathematical Society*, 298(2):635–641, 1986. ISSN 1088-6850. doi:[10.1090/s0002-9947-1986-0860384-4](https://doi.org/10.1090/s0002-9947-1986-0860384-4).
- [16] Peter Benner, Serkan Gugercin, and Karen Willcox. A survey of projection-based model reduction methods for parametric dynamical systems. *SIAM Review*, 57(4):483–531, January 2015. ISSN 1095-7200. doi:[10.1137/130932715](https://doi.org/10.1137/130932715).
- [17] G Berkooz, P Holmes, and J L Lumley. The proper orthogonal decomposition in the analysis of turbulent flows. *Annual Review of Fluid Mechanics*, 25(1):539–575, January 1993. ISSN 1545-4479. doi:[10.1146/annurev.fl.25.010193.002543](https://doi.org/10.1146/annurev.fl.25.010193.002543).
- [18] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*, volume I. Athena Scientific, Belmont, MA, USA, 3rd edition, 2005.
- [19] John T. Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. Society for Industrial and Applied Mathematics, January 2010. ISBN 9780898718577. doi:[10.1137/1.9780898718577](https://doi.org/10.1137/1.9780898718577).
- [20] R Byron Bird. Transport phenomena. *Applied Mechanics Reviews*, 55(1):R1–R4, January 2002. ISSN 2379-0407. doi:[10.1115/1.1424298](https://doi.org/10.1115/1.1424298).
- [21] Arnout M.P. Boelens, Daniele Venturi, and Daniel M. Tartakovsky. Parallel tensor methods for high-dimensional linear pdes. *Journal of Computational Physics*, 375:519–539, December 2018. ISSN 0021-9991. doi:[10.1016/j.jcp.2018.08.057](https://doi.org/10.1016/j.jcp.2018.08.057).
- [22] Olivier Bokanowski, Anya Désilles, and Hasnaa Zidani. Relationship between maximum principle and dynamic programming in presence of intermediate and final state

- constraints. *ESAIM: Control, Optimisation and Calculus of Variations*, 27:91, 2021. ISSN 1262-3377. doi:[10.1051/cocv/2021084](https://doi.org/10.1051/cocv/2021084).
- [23] Vivek S. Borkar. Contributions to the theory of optimal control. *Resonance*, 29(9): 1309–1327, September 2024. ISSN 0973-712X. doi:[10.1007/s12045-024-1309-5](https://doi.org/10.1007/s12045-024-1309-5).
- [24] Francesca Boso and Daniel M. Tartakovsky. The method of distributions for dispersive transport in porous media with uncertain hydraulic properties. *Water Resources Research*, 52(6):4700–4712, June 2016. ISSN 1944-7973. doi:[10.1002/2016wr018745](https://doi.org/10.1002/2016wr018745).
- [25] Francesca Boso and Daniel M. Tartakovsky. Data-informed method of distributions for hyperbolic conservation laws. *SIAM Journal on Scientific Computing*, 42(1):A559–A583, January 2020. ISSN 1095-7197. doi:[10.1137/19m1260773](https://doi.org/10.1137/19m1260773).
- [26] Z. I. Botev, J. F. Grotowski, and D. P. Kroese. Kernel density estimation via diffusion. *The Annals of Statistics*, 38(5), October 2010. ISSN 0090-5364. doi:[10.1214/10-aos799](https://doi.org/10.1214/10-aos799).
- [27] Anne-Celine Boulanger, Philippe Moireau, Benoît Perthame, and Jacques Sainte-Marie. Data assimilation for hyperbolic conservation laws: A luenberger observer approach based on a kinetic description. *Communications in Mathematical Sciences*, 13(3):587–622, 2015. ISSN 1945-0796. doi:[10.4310/cms.2015.v13.n3.a1](https://doi.org/10.4310/cms.2015.v13.n3.a1).
- [28] D.C.C Bover. Moment equation methods for nonlinear stochastic systems. *Journal of Mathematical Analysis and Applications*, 65(2):306–320, September 1978. ISSN 0022-247X. doi:[10.1016/0022-247x\(78\)90182-8](https://doi.org/10.1016/0022-247x(78)90182-8).
- [29] Phelim P. Boyle, Jeremy Evnine, and Stephen Gibbs. Numerical evaluation of multivariate contingent claims. *Review of Financial Studies*, 2(2):241–250, April 1989. ISSN 1465-7368. doi:[10.1093/rfs/2.2.241](https://doi.org/10.1093/rfs/2.2.241).
- [30] Matthew Brand. *Incremental Singular Value Decomposition of Uncertain Data*

- with Missing Values*, pages 707–720. Springer Berlin Heidelberg, 2002. ISBN 9783540479697. doi:[10.1007/3-540-47969-4_47](https://doi.org/10.1007/3-540-47969-4_47).
- [31] Catherine Brennan and Daniele Venturi. Data-driven closures for stochastic dynamical systems. *Journal of Computational Physics*, 372:281–298, November 2018. ISSN 0021-9991. doi:[10.1016/j.jcp.2018.06.038](https://doi.org/10.1016/j.jcp.2018.06.038).
- [32] Susanne C. Brenner and L. Ridgway Scott. *The Mathematical Theory of Finite Element Methods*. Springer New York, 2008. ISBN 9780387759340. doi:[10.1007/978-0-387-75934-0](https://doi.org/10.1007/978-0-387-75934-0).
- [33] J. Nathan Kutz Steven L. Brunton Bingni W. Brunton and Joshua L. Proctor. *Dynamic Mode Decomposition*, pages 25–38. Society for Industrial and Applied Mathematics, November 2016. ISBN 9781611974508. doi:[10.1137/1.9781611974508.ch2](https://doi.org/10.1137/1.9781611974508.ch2).
- [34] Steven L. Brunton and J. Nathan Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, January 2019. ISBN 9781108422093. doi:[10.1017/9781108380690](https://doi.org/10.1017/9781108380690).
- [35] Steven L. Brunton, Marko Budišić, Eurika Kaiser, and J. Nathan Kutz. Modern koopman theory for dynamical systems. *SIAM Review*, 64(2):229–340, May 2022. ISSN 1095-7200. doi:[10.1137/21m1401243](https://doi.org/10.1137/21m1401243).
- [36] Arthur Earl Bryson. *Applied optimal control: optimization, estimation and control*. Routledge, 2018.
- [37] Benoit Chachuat. *Nonlinear and dynamic optimization: From theory to practice*, 2007.
- [38] Abhishek Chandra, Bram Daniels, Mitrofan Curti, Koen Tiels, Elena A. Lomonova, and Daniel M. Tartakovsky. Discovery of sparse hysteresis models for piezoelectric materials. *Applied Physics Letters*, 122(21), May 2023. ISSN 1077-3118. doi:[10.1063/5.0146134](https://doi.org/10.1063/5.0146134).

- [39] Kevin K. Chen, Jonathan H. Tu, and Clarence W. Rowley. Variants of dynamic mode decomposition: Boundary condition, koopman, and fourier analyses. *Journal of Nonlinear Science*, 22(6):887–915, April 2012. ISSN 1432-1467. doi:[10.1007/s00332-012-9130-9](https://doi.org/10.1007/s00332-012-9130-9).
- [40] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18, page 6572–6583, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [41] Zhen Chen, Victor Churchill, Kailiang Wu, and Dongbin Xiu. Deep neural network modeling of unknown partial differential equations in nodal space. *Journal of Computational Physics*, 449:110782, January 2022. ISSN 0021-9991. doi:[10.1016/j.jcp.2021.110782](https://doi.org/10.1016/j.jcp.2021.110782).
- [42] Alexandre J. Chorin, Ole H. Hald, and Raz Kupferman. Optimal prediction with memory. *Physica D: Nonlinear Phenomena*, 166(3–4):239–257, June 2002. ISSN 0167-2789. doi:[10.1016/s0167-2789\(02\)00446-3](https://doi.org/10.1016/s0167-2789(02)00446-3).
- [43] Alexandre Joel Chorin. Gaussian fields and random flow. *Journal of Fluid Mechanics*, 63(1):21–32, March 1974. ISSN 1469-7645. doi:[10.1017/s0022112074000991](https://doi.org/10.1017/s0022112074000991).
- [44] I-Chun Chou and Eberhard O. Voit. Recent developments in parameter estimation and structure identification of biochemical and genomic systems. *Mathematical Biosciences*, 219(2):57–83, June 2009. ISSN 0025-5564. doi:[10.1016/j.mbs.2009.03.002](https://doi.org/10.1016/j.mbs.2009.03.002).
- [45] Walter Craig, editor. *Hamiltonian Dynamical Systems and Applications*. Springer Netherlands, 2008. ISBN 9781402069642. doi:[10.1007/978-1-4020-6964-2](https://doi.org/10.1007/978-1-4020-6964-2).
- [46] Christopher W. Curtis and Daniel Jay Alford-Lago. Dynamic-mode decomposition and

- optimal prediction. *Physical Review E*, 103(1):012201, January 2021. ISSN 2470-0053. doi:[10.1103/physreve.103.012201](https://doi.org/10.1103/physreve.103.012201).
- [47] Frank E. Curtis, Jorge Nocedal, and Andreas Wächter. A matrix-free algorithm for equality constrained optimization problems with rank-deficient jacobians. *SIAM Journal on Optimization*, 20(3):1224–1249, January 2010. ISSN 1095-7189. doi:[10.1137/08072471x](https://doi.org/10.1137/08072471x).
- [48] J. H. Curtiss. A note on the theory of moment generating functions. *The Annals of Mathematical Statistics*, 13(4):430–433, December 1942. ISSN 0003-4851. doi:[10.1214/aoms/1177731541](https://doi.org/10.1214/aoms/1177731541).
- [49] Tobias Damm, Lars Grüne, Marleen Stieler, and Karl Worthmann. An exponential turnpike theorem for dissipative discrete time optimal control problems. *SIAM Journal on Control and Optimization*, 52(3):1935–1957, January 2014. ISSN 1095-7138. doi:[10.1137/120888934](https://doi.org/10.1137/120888934).
- [50] Maria Yuliani Danggo and Sudi Mungkasi. A staggered grid finite difference method for solving the elastic wave equations. *Journal of Physics: Conference Series*, 909: 012047, November 2017. ISSN 1742-6596. doi:[10.1088/1742-6596/909/1/012047](https://doi.org/10.1088/1742-6596/909/1/012047).
- [51] Juan Carlos De los Reyes. *Numerical PDE-Constrained Optimization*. Springer International Publishing, 2015. ISBN 9783319133959. doi:[10.1007/978-3-319-13395-9](https://doi.org/10.1007/978-3-319-13395-9).
- [52] C. De Marco and A. Bergen. A security measure for random load disturbances in nonlinear power system models. *IEEE Transactions on Circuits and Systems*, 34(12): 1546–1557, December 1987. ISSN 0098-4094. doi:[10.1109/tcs.1987.1086092](https://doi.org/10.1109/tcs.1987.1086092).
- [53] Alec Dektor, Abram Rodgers, and Daniele Venturi. Rank-adaptive tensor methods for high-dimensional nonlinear pdes. *Journal of Scientific Computing*, 88(2), June 2021. ISSN 1573-7691. doi:[10.1007/s10915-021-01539-3](https://doi.org/10.1007/s10915-021-01539-3).

- [54] James W. Demmel. *Applied Numerical Linear Algebra*, pages 265–360. Society for Industrial and Applied Mathematics, January 1997. ISBN 9781611971446. doi:[10.1137/1.9781611971446.ch6](https://doi.org/10.1137/1.9781611971446.ch6).
- [55] L. Dieci and C. Elia. The singular value decomposition to approximate spectra of dynamical systems. theoretical aspects. *Journal of Differential Equations*, 230(2):502–531, November 2006. ISSN 0022-0396. doi:[10.1016/j.jde.2006.08.007](https://doi.org/10.1016/j.jde.2006.08.007).
- [56] Ian Dobson, Benjamin A. Carreras, Vickie E. Lynch, and David E. Newman. Complex systems analysis of series of blackouts: Cascading failure, critical points, and self-organization. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 17(2), June 2007. ISSN 1089-7682. doi:[10.1063/1.2737822](https://doi.org/10.1063/1.2737822).
- [57] Victorita Dolean, Pierre Jolivet, and Frédéric Nataf. *An Introduction to Domain Decomposition Methods: Algorithms, Theory, and Parallel Implementation*. Society for Industrial and Applied Mathematics, November 2015. ISBN 9781611974065. doi:[10.1137/1.9781611974065](https://doi.org/10.1137/1.9781611974065).
- [58] Weinan E and Bing Yu. The deep ritz method: A deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, February 2018. ISSN 2194-671X. doi:[10.1007/s40304-018-0127-z](https://doi.org/10.1007/s40304-018-0127-z).
- [59] Weinan E, Jiequn Han, and Qianxiao Li. A mean-field optimal control formulation of deep learning. *Research in the Mathematical Sciences*, 6(1), December 2018. ISSN 2197-9847. doi:[10.1007/s40687-018-0172-y](https://doi.org/10.1007/s40687-018-0172-y).
- [60] Weinan E, Martin Hutzenthaler, Arnulf Jentzen, and Thomas Kruse. On multilevel picard numerical approximations for high-dimensional nonlinear parabolic partial differential equations and high-dimensional nonlinear backward stochastic differential equa-

- tions. *Journal of Scientific Computing*, 79(3):1534–1571, March 2019. ISSN 1573-7691. doi:[10.1007/s10915-018-00903-0](https://doi.org/10.1007/s10915-018-00903-0).
- [61] Philipp Eichmeir, Thomas Lauß, Stefan Oberpeilsteiner, Karin Nachbagauer, and Wolfgang Steiner. The adjoint method for time-optimal control problems. *Journal of Computational and Nonlinear Dynamics*, 16(2), November 2020. ISSN 1555-1423. doi:[10.1115/1.4048808](https://doi.org/10.1115/1.4048808).
- [62] Lawrence Evans. *Partial Differential Equations*. American Mathematical Society, March 2010. ISBN 9781470411442. doi:[10.1090/gsm/019](https://doi.org/10.1090/gsm/019).
- [63] Timm Faulwasser, Milan Korda, Colin N. Jones, and Dominique Bonvin. On turnpike and dissipativity properties of continuous-time optimal control problems. *Automatica*, 81:297–304, July 2017. ISSN 0005-1098. doi:[10.1016/j.automatica.2017.03.012](https://doi.org/10.1016/j.automatica.2017.03.012).
- [64] Carl Folkestad, Daniel Pastor, Igor Mezic, Ryan Mohr, Maria Fonoberova, and Joel Burdick. Extended dynamic mode decomposition with learned koopman eigenfunctions for prediction and control. In *2020 American Control Conference (ACC)*, pages 3906–3913. IEEE, July 2020. doi:[10.23919/acc45564.2020.9147729](https://doi.org/10.23919/acc45564.2020.9147729).
- [65] Mattia Frasca and Lucia Valentina Gambuzza. Control of cascading failures in dynamical models of power grids. *Chaos, Solitons and Fractals*, 153:111460, December 2021. ISSN 0960-0779. doi:[10.1016/j.chaos.2021.111460](https://doi.org/10.1016/j.chaos.2021.111460).
- [66] Mark I. Freidlin and Alexander D. Wentzell. *The Multidimensional Case*, pages 355–389. Springer Berlin Heidelberg, 2012. ISBN 9783642258473. doi:[10.1007/978-3-642-25847-3_9](https://doi.org/10.1007/978-3-642-25847-3_9).
- [67] Mark I. Freidlin and Alexander D. Wentzell. *Random Perturbations of Dynamical Systems*. Springer Berlin Heidelberg, 2012. ISBN 9783642258473. doi:[10.1007/978-3-642-25847-3](https://doi.org/10.1007/978-3-642-25847-3).

- [68] Xiaohan Fu, Lo-Bin Chang, and Dongbin Xiu. Learning reduced systems via deep neural networks with memory. *Journal of Machine Learning for Modeling and Computing*, 1(2):97–118, 2020. ISSN 2689-3967. doi:[10.1615/.2020034232](https://doi.org/10.1615/2020034232).
- [69] Borjan Geshkovski and Enrique Zuazua. Turnpike in optimal control of pdes, resnets, and beyond. *Acta Numerica*, 31:135–263, May 2022. ISSN 1474-0508. doi:[10.1017/s0962492922000046](https://doi.org/10.1017/s0962492922000046).
- [70] Roger G. Ghanem and Pol D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer New York, 1991. ISBN 9781461230946. doi:[10.1007/978-1-4612-3094-6](https://doi.org/10.1007/978-1-4612-3094-6).
- [71] Michael B. Giles. Multilevel monte carlo path simulation. *Operations Research*, 56(3):607–617, June 2008. ISSN 1526-5463. doi:[10.1287/opre.1070.0496](https://doi.org/10.1287/opre.1070.0496).
- [72] Philip E. Gill and Elizabeth Wong. *Sequential Quadratic Programming Methods*, pages 147–224. Springer New York, November 2011. ISBN 9781461419273. doi:[10.1007/978-1-4614-1927-3_6](https://doi.org/10.1007/978-1-4614-1927-3_6).
- [73] Daniel T. Gillespie. Stochastic simulation of chemical kinetics. *Annual Review of Physical Chemistry*, 58(1):35–55, May 2007. ISSN 1545-1593. doi:[10.1146/annurev.physchem.58.032806.104637](https://doi.org/10.1146/annurev.physchem.58.032806.104637).
- [74] Jacek Gondzio. Matrix-free interior point method. *Computational Optimization and Applications*, 51(2):457–480, October 2010. ISSN 1573-2894. doi:[10.1007/s10589-010-9361-3](https://doi.org/10.1007/s10589-010-9361-3).
- [75] Tobias Grafke and Eric Vanden-Eijnden. Numerical computation of rare events via large deviation theory. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29(6), June 2019. ISSN 1089-7682. doi:[10.1063/1.5084025](https://doi.org/10.1063/1.5084025).

- [76] Muzio Grilli, Peter J. Schmid, Stefan Hickel, and Nikolaus A. Adams. Analysis of unsteady behaviour in shockwave turbulent boundary layer interaction. *Journal of Fluid Mechanics*, 700:16–28, February 2012. ISSN 1469-7645. doi:[10.1017/jfm.2012.37](https://doi.org/10.1017/jfm.2012.37).
- [77] Lars Grüne. Economic receding horizon control without terminal constraints. *Automatica*, 49(3):725–734, March 2013. ISSN 0005-1098. doi:[10.1016/j.automatica.2012.12.003](https://doi.org/10.1016/j.automatica.2012.12.003).
- [78] Lars Grüne, Manuel Schaller, and Anton Schiela. Abstract nonlinear sensitivity and turnpike analysis and an application to semilinear parabolic pdes. *ESAIM: Control, Optimisation and Calculus of Variations*, 27:56, 2021. ISSN 1262-3377. doi:[10.1051/cocv/2021030](https://doi.org/10.1051/cocv/2021030).
- [79] Damien Gueho, Puneet Singla, and Manoranjan Majji. Time-varying koopman operator theory for nonlinear systems prediction. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 2294–2299. IEEE, December 2021. doi:[10.1109/cdc45484.2021.9683082](https://doi.org/10.1109/cdc45484.2021.9683082).
- [80] Serkan Gugercin and Athanasios C. Antoulas. A survey of model reduction by balanced truncation and some new results. *International Journal of Control*, 77(8):748–766, May 2004. ISSN 1366-5820. doi:[10.1080/00207170410001713448](https://doi.org/10.1080/00207170410001713448).
- [81] Max D. Gunzburger. *Perspectives in Flow Control and Optimization*, pages 11–62. Society for Industrial and Applied Mathematics, January 2002. ISBN 9780898718720. doi:[10.1137/1.9780898718720.ch2](https://doi.org/10.1137/1.9780898718720.ch2).
- [82] Mengwu Guo and Jan S. Hesthaven. Data-driven reduced order modeling for time-dependent problems. *Computer Methods in Applied Mechanics and Engineering*, 345: 75–99, March 2019. ISSN 0045-7825. doi:[10.1016/j.cma.2018.10.029](https://doi.org/10.1016/j.cma.2018.10.029).

- [83] William W. Hager. Updating the inverse of a matrix. *SIAM Review*, 31(2):221–239, June 1989. ISSN 1095-7200. doi:[10.1137/1031049](https://doi.org/10.1137/1031049).
- [84] William W. Hager. Updating the inverse of a matrix. *SIAM Review*, 31(2):221–239, June 1989. ISSN 1095-7200. doi:[10.1137/1031049](https://doi.org/10.1137/1031049).
- [85] William W. Hager. Runge-kutta methods in optimal control and the transformed adjoint system. *Numerische Mathematik*, 87(2):247–282, December 2000. ISSN 0945-3245. doi:[10.1007/s002110000178](https://doi.org/10.1007/s002110000178).
- [86] Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, August 2018. ISSN 1091-6490. doi:[10.1073/pnas.1718942115](https://doi.org/10.1073/pnas.1718942115).
- [87] Xiaoying Han and Peter E. Kloeden. *Random Ordinary Differential Equations and Their Numerical Solution*. Springer Singapore, 2017. ISBN 9789811062650. doi:[10.1007/978-981-10-6265-0](https://doi.org/10.1007/978-981-10-6265-0).
- [88] Yue Hao and Valeria Simoncini. The sherman–morrison–woodbury formula for generalized linear matrix equations and applications. *Numerical Linear Algebra with Applications*, 28(5), May 2021. ISSN 1099-1506. doi:[10.1002/nla.2384](https://doi.org/10.1002/nla.2384).
- [89] Moritz Hardt, Ben Recht, and Yoram Singer. Train faster, generalize better: Stability of stochastic gradient descent. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1225–1234, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <https://proceedings.mlr.press/v48/hardt16.html>.
- [90] Ramin Hasani, Mathias Lechner, Alexander Amini, Lucas Liebenwein, Aaron Ray, Max Tschaikowski, Gerald Teschl, and Daniela Rus. Closed-form continuous-time

- neural networks. *Nature Machine Intelligence*, 4(11):992–1003, November 2022. ISSN 2522-5839. doi:[10.1038/s42256-022-00556-7](https://doi.org/10.1038/s42256-022-00556-7).
- [91] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer New York, 2009. ISBN 9780387848587. doi:[10.1007/978-0-387-84858-7](https://doi.org/10.1007/978-0-387-84858-7).
- [92] Maziar S. Hemati, Matthew O. Williams, and Clarence W. Rowley. Dynamic mode decomposition for large and streaming datasets. *Physics of Fluids*, 26(11), November 2014. ISSN 1089-7666. doi:[10.1063/1.4901016](https://doi.org/10.1063/1.4901016).
- [93] Monica Hernandez. Newton-krylov pde-constrained lddmm in the space of band-limited vector fields. *arXiv preprint arXiv:1807.05117*, 2018.
- [94] Monica Hernandez. A comparative study of different variants of newton–krylov pde-constrained stokes-lddmm parameterized in the space of band-limited vector fields. *SIAM Journal on Imaging Sciences*, 12(2):1038–1070, January 2019. ISSN 1936-4954. doi:[10.1137/18m1195310](https://doi.org/10.1137/18m1195310).
- [95] Roland Herzog and Karl Kunisch. Algorithms for pde-constrained optimization. *GAMM-Mitteilungen*, 33(2):163–176, October 2010. ISSN 1522-2608. doi:[10.1002/gamm.201010013](https://doi.org/10.1002/gamm.201010013).
- [96] Martin W. Hess, Annalisa Quaini, and Gianluigi Rozza. A data-driven surrogate modeling approach for time-dependent incompressible navier-stokes equations with dynamic mode decomposition and manifold interpolation. *Advances in Computational Mathematics*, 49(2), March 2023. ISSN 1572-9044. doi:[10.1007/s10444-023-10016-4](https://doi.org/10.1007/s10444-023-10016-4).
- [97] Paul Hines, Ian Dobson, and Pooya Rezaei. Cascading power outages propagate locally in an influence graph that is not the actual grid topology. *IEEE Transactions on Power Systems*, pages 1–1, 2016. ISSN 1558-0679. doi:[10.1109/tpwrs.2016.2578259](https://doi.org/10.1109/tpwrs.2016.2578259).

- [98] Michael Hinze. *Mathematical Modelling: Theory and Applications*. Springer Netherlands, 2009. ISBN 9781402088391. doi:[10.1007/978-1-4020-8839-1](https://doi.org/10.1007/978-1-4020-8839-1).
- [99] Michael Hinze and Karl Kunisch. Second order methods for optimal control of time-dependent fluid flow. *SIAM Journal on Control and Optimization*, 40(3):925–946, January 2001. ISSN 1095-7138. doi:[10.1137/s0363012999361810](https://doi.org/10.1137/s0363012999361810).
- [100] Michael Hinze, Andrea Walther, and Julia Sternberg. An optimal memory-reduced procedure for calculating adjoints of the instationary navier-stokes equations. *Optimal Control Applications and Methods*, 27(1):19–40, 2006. doi:<https://doi.org/10.1002/oca.771>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/oca.771>.
- [101] Philip Holmes, John L. Lumley, and Gal Berkooz. *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press, October 1996. ISBN 9780511622700. doi:[10.1017/cbo9780511622700](https://doi.org/10.1017/cbo9780511622700).
- [102] Ignacio Huerta, Pablo Monzón, and Gonzalo Robledo. Controllability and feedback stabilizability in a nonuniform framework. *Comptes Rendus. Mathématique*, 362(G12):1667–1692, November 2024. ISSN 1778-3569. doi:[10.5802/crmath.672](https://doi.org/10.5802/crmath.672).
- [103] G D Hutomo, J Kusuma, A Ribal, A G Mahie, and N Aris. Numerical solution of 2-d advection-diffusion equation with variable coefficient using du-fort frankel method. *Journal of Physics: Conference Series*, 1180:012009, February 2019. ISSN 1742-6596. doi:[10.1088/1742-6596/1180/1/012009](https://doi.org/10.1088/1742-6596/1180/1/012009).
- [104] Michael Holst Ivar Stakgold, editor. *Green's Functions and Boundary Value Problems*. Wiley, January 2011. ISBN 9780470906538. doi:[10.1002/9780470906538](https://doi.org/10.1002/9780470906538).
- [105] Andrew H. Jazwinski, editor. *Stochastic Processes and Filtering Theory*, pages 194–265. Elsevier, 1970. doi:[10.1016/s0076-5392\(09\)60376-3](https://doi.org/10.1016/s0076-5392(09)60376-3).

- [106] Weinan E Jiequn Han and Qianxiao Li. *Dynamical Systems and Optimal Control Approach to Deep Learning*, pages 422–438. Cambridge University Press, December 2022. ISBN 9781316516782. doi:[10.1017/9781009025096.011](https://doi.org/10.1017/9781009025096.011).
- [107] Karen Kafadar, Adrian W. Bowman, and Adelchi Azzalini. Applied smoothing techniques for data analysis: The kernel approach with s-plus illustrations. *Journal of the American Statistical Association*, 94(447):982, September 1999. ISSN 0162-1459. doi:[10.2307/2670015](https://doi.org/10.2307/2670015).
- [108] A. A. Kaptanoglu, K. D. Morgan, C. J. Hansen, and S. L. Brunton. Characterizing magnetized plasmas with dynamic mode decomposition. *Physics of Plasmas*, 27(3), March 2020. ISSN 1089-7674. doi:[10.1063/1.5138932](https://doi.org/10.1063/1.5138932).
- [109] Takahiro Kawashima and Hideitsu Hino. Gaussian process koopman mode decomposition. *Neural Computation*, 35(1):82–103, January 2023. ISSN 1530-888X. doi:[10.1162/neco_a_01555](https://doi.org/10.1162/neco_a_01555).
- [110] Swasti R. Khuntia, Jose L. Rueda, and Mart A. M. M. van der Meijden. Mutual information based bayesian analysis of power system reliability. In *2015 IEEE Eindhoven PowerTech*, pages 1–6. IEEE, June 2015. doi:[10.1109/ptc.2015.7232592](https://doi.org/10.1109/ptc.2015.7232592).
- [111] Donald E Kirk. *Optimal control theory: an introduction*. Courier Corporation, 2004.
- [112] Peter E. Kloeden and Eckhard Platen. *Selected Applications of Strong Approximations*, pages 427–456. Springer Berlin Heidelberg, 1992. ISBN 9783662126165. doi:[10.1007/978-3-662-12616-5_13](https://doi.org/10.1007/978-3-662-12616-5_13).
- [113] Stefan Klus, Patrick Gelß, Sebastian Peitz, and Christof Schütte. Tensor-based dynamic mode decomposition. *Nonlinearity*, 31(7):3359–3380, June 2018. ISSN 1361-6544. doi:[10.1088/1361-6544/aabc8f](https://doi.org/10.1088/1361-6544/aabc8f).

- [114] Milan Korda and Igor Mezić. On convergence of extended dynamic mode decomposition to the koopman operator. *Journal of Nonlinear Science*, 28(2):687–710, November 2017. ISSN 1432-1467. doi:[10.1007/s00332-017-9423-0](https://doi.org/10.1007/s00332-017-9423-0).
- [115] Robert H. Kraichnan. Dynamics of nonlinear stochastic systems. *Journal of Mathematical Physics*, 2(1):124–148, January 1961. ISSN 1089-7658. doi:[10.1063/1.1724206](https://doi.org/10.1063/1.1724206).
- [116] Mario Krenn, Robert Pollice, Si Yue Guo, Matteo Aldeghi, Alba Cervera-Lierta, Pascal Friederich, Gabriel dos Passos Gomes, Florian Häse, Adrian Jinich, AkshatKumar Nigam, Zhenpeng Yao, and Alán Aspuru-Guzik. On scientific understanding with artificial intelligence. *Nature Reviews Physics*, 4(12):761–769, October 2022. ISSN 2522-5820. doi:[10.1038/s42254-022-00518-3](https://doi.org/10.1038/s42254-022-00518-3).
- [117] Aditi S. Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W Mahoney. Characterizing possible failure modes in physics-informed neural networks. *Advances in Neural Information Processing Systems*, 34, 2021.
- [118] Ryogo Kubo. Generalized cumulant expansion method. *Journal of the Physical Society of Japan*, 17(7):1100–1120, July 1962. ISSN 1347-4073. doi:[10.1143/jpsj.17.1100](https://doi.org/10.1143/jpsj.17.1100).
- [119] K. Kunisch and S. Volkwein. Galerkin proper orthogonal decomposition methods for a general equation in fluid dynamics. *SIAM Journal on Numerical Analysis*, 40(2):492–515, January 2002. ISSN 1095-7170. doi:[10.1137/s0036142900382612](https://doi.org/10.1137/s0036142900382612).
- [120] J. Nathan Kutz, Xing Fu, and Steven L. Brunton. Multiresolution dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 15(2):713–735, January 2016. ISSN 1536-0040. doi:[10.1137/15m1023543](https://doi.org/10.1137/15m1023543).
- [121] John H. Lagergren, John T. Nardini, Ruth E. Baker, Matthew J. Simpson, and Kevin B. Flores. Biologically-informed neural networks guide mechanistic modeling

- from sparse experimental data. *PLOS Computational Biology*, 16(12):e1008462, December 2020. ISSN 1553-7358. doi:[10.1371/journal.pcbi.1008462](https://doi.org/10.1371/journal.pcbi.1008462).
- [122] L. Lasdon, S. Mitter, and A. Waren. The conjugate gradient method for optimal control problems. *IEEE Transactions on Automatic Control*, 12(2):132–138, 1967. doi:[10.1109/TAC.1967.1098538](https://doi.org/10.1109/TAC.1967.1098538).
- [123] Soledad Le Clainche and José M. Vega. Higher order dynamic mode decomposition to identify and extrapolate flow patterns. *Physics of Fluids*, 29(8), August 2017. ISSN 1089-7666. doi:[10.1063/1.4997206](https://doi.org/10.1063/1.4997206).
- [124] François-Xavier Le Dimet and Olivier Talagrand. Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects. *Tellus A: Dynamic Meteorology and Oceanography*, 38(2):97, January 1986. ISSN 1600-0870. doi:[10.3402/tellusa.v38i2.11706](https://doi.org/10.3402/tellusa.v38i2.11706).
- [125] Lili Lei and Joshua P. Hacker. Nudging, ensemble, and nudging ensembles for data assimilation in the presence of model error. *Monthly Weather Review*, 143(7):2600–2610, July 2015. ISSN 1520-0493. doi:[10.1175/mwr-d-14-00295.1](https://doi.org/10.1175/mwr-d-14-00295.1).
- [126] Amanda Lenzi and Mihai Anitescu. How can statistics help to prevent blackouts? *Significance*, 20(1):24–27, February 2023. ISSN 1740-9713. doi:[10.1093/jrssig/qmad009](https://doi.org/10.1093/jrssig/qmad009). URL <http://dx.doi.org/10.1093/jrssig/qmad009>.
- [127] Randall J. LeVeque. *Variable-Coefficient Linear Equations*, pages 158–187. Cambridge University Press, August 2002. doi:[10.1017/cbo9780511791253.010](https://doi.org/10.1017/cbo9780511791253.010).
- [128] Randall J. LeVeque. *Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems*, pages 113–135. Society for Industrial and Applied Mathematics, January 2007. ISBN 9780898717839. doi:[10.1137/1.9780898717839.ch5](https://doi.org/10.1137/1.9780898717839.ch5).

- [129] Jing Li and Panos Stinis. Model reduction for a power grid model. *Journal of Computational Dynamics*, 9(1):1, 2022. ISSN 2158-2505. doi:[10.3934/jcd.2021019](https://doi.org/10.3934/jcd.2021019).
- [130] Hannah Lu and Daniel M. Tartakovsky. Lagrangian dynamic mode decomposition for construction of reduced-order models of advection-dominated phenomena. *Journal of Computational Physics*, 407:109229, April 2020. ISSN 0021-9991. doi:[10.1016/j.jcp.2020.109229](https://doi.org/10.1016/j.jcp.2020.109229).
- [131] Hannah Lu and Daniel M. Tartakovsky. Prediction accuracy of dynamic mode decomposition. *SIAM Journal on Scientific Computing*, 42(3):A1639–A1662, January 2020. ISSN 1095-7197. doi:[10.1137/19m1259948](https://doi.org/10.1137/19m1259948).
- [132] Hannah Lu and Daniel M. Tartakovsky. Dynamic mode decomposition for construction of reduced-order models of hyperbolic problems with shocks. *Journal of Machine Learning for Modeling and Computing*, 2(1):1–29, 2021. ISSN 2689-3967. doi:[10.1615/jmachlearnmodelcomput.2021036132](https://doi.org/10.1615/jmachlearnmodelcomput.2021036132).
- [133] Hannah Lu and Daniel M. Tartakovsky. Extended dynamic mode decomposition for inhomogeneous problems. *Journal of Computational Physics*, 444:110550, November 2021. ISSN 0021-9991. doi:[10.1016/j.jcp.2021.110550](https://doi.org/10.1016/j.jcp.2021.110550).
- [134] Hannah Lu and Daniel M. Tartakovsky. Extended dynamic mode decomposition for inhomogeneous problems. *Journal of Computational Physics*, 444:110550, November 2021. ISSN 0021-9991. doi:[10.1016/j.jcp.2021.110550](https://doi.org/10.1016/j.jcp.2021.110550).
- [135] Hannah Lu and Daniel M. Tartakovsky. Drips: A framework for dimension reduction and interpolation in parameter space. *Journal of Computational Physics*, 493:112455, November 2023. ISSN 0021-9991. doi:[10.1016/j.jcp.2023.112455](https://doi.org/10.1016/j.jcp.2023.112455).
- [136] Qiugang Lu, Sungho Shin, and Victor M. Zavala. Characterizing the predictive accu-

- racy of dynamic mode decomposition for data-driven control. *IFAC-PapersOnLine*, 53(2):11289–11294, 2020. ISSN 2405-8963. doi:[10.1016/j.ifacol.2020.12.373](https://doi.org/10.1016/j.ifacol.2020.12.373).
- [137] S. H. Lui. On schwarz alternating methods for nonlinear elliptic pdes. *SIAM Journal on Scientific Computing*, 21(4):1506–1523, January 1999. ISSN 1095-7197. doi:[10.1137/s1064827597327553](https://doi.org/10.1137/s1064827597327553).
- [138] Lina Ma, Xiantao Li, and Chun Liu. Coarse-graining langevin dynamics using reduced-order techniques. *Journal of Computational Physics*, 380:170–190, March 2019. ISSN 0021-9991. doi:[10.1016/j.jcp.2018.11.035](https://doi.org/10.1016/j.jcp.2018.11.035).
- [139] Lina Ma, Xiantao Li, and Chun Liu. Coarse-graining langevin dynamics using reduced-order techniques. *Journal of Computational Physics*, 380:170–190, March 2019. ISSN 0021-9991. doi:[10.1016/j.jcp.2018.11.035](https://doi.org/10.1016/j.jcp.2018.11.035).
- [140] Andrew Majda and Xiaoming Wang. *Non-linear dynamics and statistical theories for basic geophysical flows*. Cambridge University Press, 2006. ISBN 9780511616778. doi:[10.1017/cbo9780511616778](https://doi.org/10.1017/cbo9780511616778).
- [141] Daniel A. Maldonado, Michel Schanen, and Mihai Anitescu. Uncertainty propagation in power system dynamics with the method of moments. In *2018 IEEE Power and Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, August 2018. doi:[10.1109/pesgm.2018.8586023](https://doi.org/10.1109/pesgm.2018.8586023).
- [142] Tyler Maltba, Pierre A. Gremaud, and Daniel M. Tartakovsky. Nonlocal pdf methods for langevin equations with colored noise. *Journal of Computational Physics*, 367:87–101, August 2018. ISSN 0021-9991. doi:[10.1016/j.jcp.2018.04.023](https://doi.org/10.1016/j.jcp.2018.04.023).
- [143] Tyler E. Maltba, Hongli Zhao, and Daniel M. Tartakovsky. Autonomous learning of nonlocal stochastic neuron dynamics. *Cognitive Neurodynamics*, 16(3):683–705, November 2021. ISSN 1871-4099. doi:[10.1007/s11571-021-09731-9](https://doi.org/10.1007/s11571-021-09731-9).

- [144] Tyler E. Maltba, Vishwas Rao, and Daniel Adrian Maldonado. Learning the evolution of correlated stochastic power system dynamics. In *2022 IEEE Power and Energy Society General Meeting (PESGM)*, pages 01–05. IEEE, July 2022. doi:[10.1109/pesgm48719.2022.9916982](https://doi.org/10.1109/pesgm48719.2022.9916982).
- [145] Tyler E. Maltba, Hongli Zhao, and D. Adrian Maldonado. Data-driven closures and assimilation for stiff multiscale random dynamics. *SIAM Journal on Scientific Computing*, 47(1):C52–C76, January 2025. ISSN 1095-7197. doi:[10.1137/23m1625366](https://doi.org/10.1137/23m1625366).
- [146] Karel Matouš, Marc G.D. Geers, Varvara G. Kouznetsova, and Andrew Gillman. A review of predictive nonlinear theories for multiscale modeling of heterogeneous materials. *Journal of Computational Physics*, 330:192–220, February 2017. ISSN 0021-9991. doi:[10.1016/j.jcp.2016.10.070](https://doi.org/10.1016/j.jcp.2016.10.070).
- [147] Igor Mezic and Amit Surana. Koopman mode decomposition for periodic/quasi-periodic time dependence**the funding provided by utrc is greatly appreciated. *IFAC-PapersOnLine*, 49(18):690–697, 2016. ISSN 2405-8963. doi:[10.1016/j.ifacol.2016.10.246](https://doi.org/10.1016/j.ifacol.2016.10.246).
- [148] Igor Mezic and Amit Surana. Koopman mode decomposition for periodic/quasi-periodic time dependence**the funding provided by utrc is greatly appreciated. *IFAC-PapersOnLine*, 49(18):690–697, 2016. ISSN 2405-8963. doi:[10.1016/j.ifacol.2016.10.246](https://doi.org/10.1016/j.ifacol.2016.10.246).
- [149] Igor Mezić. Spectrum of the koopman operator, spectral expansions in functional spaces, and state-space geometry. *Journal of Nonlinear Science*, 30(5):2091–2145, December 2019. ISSN 1432-1467. doi:[10.1007/s00332-019-09598-5](https://doi.org/10.1007/s00332-019-09598-5).
- [150] Federico Milano and Rafael Zarate-Minano. A systematic method to model power systems as stochastic differential algebraic equations. *IEEE Transac-*

- tions on Power Systems*, 28(4):4537–4544, November 2013. ISSN 1558-0679. doi:[10.1109/tpwrs.2013.2266441](https://doi.org/10.1109/tpwrs.2013.2266441).
- [151] Sen Na and Mihai Anitescu. Exponential decay in the sensitivity analysis of nonlinear dynamic programming. *SIAM Journal on Optimization*, 30(2):1527–1554, January 2020. ISSN 1095-7189. doi:[10.1137/19m1265065](https://doi.org/10.1137/19m1265065).
- [152] Sen Na, Sungho Shin, Mihai Anitescu, and Victor M. Zavala. On the convergence of overlapping schwarz decomposition for nonlinear optimal control. *IEEE Transactions on Automatic Control*, 67(11):5996–6011, November 2022. ISSN 2334-3303. doi:[10.1109/tac.2022.3194087](https://doi.org/10.1109/tac.2022.3194087).
- [153] Sen Na, Mihai Anitescu, and Mladen Kolar. A fast temporal decomposition procedure for long-horizon nonlinear dynamic programming. *Mathematics of Operations Research*, 49(2):1012–1044, May 2024. ISSN 1526-5471. doi:[10.1287/moor.2023.1378](https://doi.org/10.1287/moor.2023.1378).
- [154] Nirmal J. Nair and Maciej Balajewicz. Transported snapshot model order reduction approach for parametric, steady-state fluid flows containing parameter-dependent shocks. *International Journal for Numerical Methods in Engineering*, 117(12):1234–1262, December 2018. ISSN 1097-0207. doi:[10.1002/nme.5998](https://doi.org/10.1002/nme.5998).
- [155] Hiroya Nakao and Igor Mezić. Spectral analysis of the koopman operator for partial differential equations. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 30(11), November 2020. ISSN 1089-7682. doi:[10.1063/5.0011470](https://doi.org/10.1063/5.0011470).
- [156] Tommaso Nesti, Jayakrishnan Nair, and Bert Zwart. Temperature overloads in power grids under uncertainty: A large deviations approach. *IEEE Transactions on Control of Network Systems*, 6(3):1161–1173, September 2019. ISSN 2372-2533. doi:[10.1109/tcms.2019.2922492](https://doi.org/10.1109/tcms.2019.2922492).

- [157] J. W. Neuberger. *Continuous Steepest Descent in Hilbert Space: Nonlinear Case*, pages 19–34. Springer Berlin Heidelberg, September 2009. ISBN 9783642040412. doi:[10.1007/978-3-642-04041-2_4](https://doi.org/10.1007/978-3-642-04041-2_4).
- [158] Martin P. Neuenhofen and Eric C. Kerrigan. An integral penalty-barrier direct transcription method for optimal control. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 456–463. IEEE, December 2020. doi:[10.1109/cdc42340.2020.9304216](https://doi.org/10.1109/cdc42340.2020.9304216).
- [159] BERND R. NOACK, KONSTANTIN AFANASIEV, MAREK MORZYŃSKI, GILEAD TADMOR, and FRANK THIELE. A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *Journal of Fluid Mechanics*, 497:335–363, December 2003. ISSN 1469-7645. doi:[10.1017/s0022112003006694](https://doi.org/10.1017/s0022112003006694).
- [160] Matthew O. Williams, Clarence W. Rowley, and Ioannis G. Kevrekidis. A kernel-based method for data-driven koopman spectral analysis. *Journal of Computational Dynamics*, 2(2):247–265, 2015. ISSN 2158-2505. doi:[10.3934/jcd.2015005](https://doi.org/10.3934/jcd.2015005).
- [161] Said Ouala, Bertrand Chapron, Fabrice Collard, Lucile Gaultier, and Ronan Fablet. Extending the extended dynamic mode decomposition with latent observables: the latent edmd framework. *Machine Learning: Science and Technology*, 4(2):025018, May 2023. ISSN 2632-2153. doi:[10.1088/2632-2153/acccd6](https://doi.org/10.1088/2632-2153/acccd6).
- [162] J. Owedyk and A. Kociszewski. On the fokker-planck equation with time-dependent drift and diffusion coefficients and its exponential solutions. *Zeitschrift für Physik B Condensed Matter*, 59(1):69–74, 1985. doi:[10.1007/BF01325384](https://doi.org/10.1007/BF01325384). URL <https://doi.org/10.1007/BF01325384>.
- [163] T.N Palmer and P.D Williams. Introduction. stochastic physics and climate modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physi-*

- cal and Engineering Sciences*, 366(1875):2419–2425, April 2008. ISSN 1471-2962. doi:[10.1098/rsta.2008.0059](https://doi.org/10.1098/rsta.2008.0059).
- [164] Suraj Pawar, Omer San, Adil Rasheed, and Prakash Vedula. Frame invariant neural network closures for kraichnan turbulence. *Physica A: Statistical Mechanics and its Applications*, 609:128327, January 2023. ISSN 0378-4371. doi:[10.1016/j.physa.2022.128327](https://doi.org/10.1016/j.physa.2022.128327).
- [165] A. Pazy. *Semigroups of Linear Operators and Applications to Partial Differential Equations*. Springer New York, 1983. ISBN 9781461255611. doi:[10.1007/978-1-4612-5561-1](https://doi.org/10.1007/978-1-4612-5561-1).
- [166] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review*, 60(3):550–591, January 2018. ISSN 1095-7200. doi:[10.1137/16m1082469](https://doi.org/10.1137/16m1082469).
- [167] J. R. Philip. Dynamics of fluids in porous media. *Journal of Fluid Mechanics*, 61(1):206–208, October 1973. ISSN 1469-7645. doi:[10.1017/s0022112073210662](https://doi.org/10.1017/s0022112073210662).
- [168] R.-E. Plessix. A review of the adjoint-state method for computing the gradient of a functional with geophysical applications. *Geophysical Journal International*, 167(2):495–503, November 2006. ISSN 1365-246X. doi:[10.1111/j.1365-246x.2006.02978.x](https://doi.org/10.1111/j.1365-246x.2006.02978.x).
- [169] Stephen B. Pope. *Turbulent Flows*. Cambridge University Press, August 2000. ISBN 9780511840531. doi:[10.1017/cbo9780511840531](https://doi.org/10.1017/cbo9780511840531).
- [170] Ernesto Prudencio, Richard Byrd, and Xiao-Chuan Cai. *Domain Decomposition Methods for PDE Constrained Optimization Problems*, pages 569–582. Springer Berlin Heidelberg, 2005. ISBN 9783540318545. doi:[10.1007/11403937_43](https://doi.org/10.1007/11403937_43).
- [171] Ernesto E. Prudencio, Richard Byrd, and Xiao-Chuan Cai. Parallel full space sqp lagrange-newton-krylov-schwarz algorithms for pde-constrained optimization prob-

- lems. *SIAM Journal on Scientific Computing*, 27(4):1305–1328, January 2006. ISSN 1095-7197. doi:[10.1137/040602997](https://doi.org/10.1137/040602997).
- [172] Tong Qin, Kailiang Wu, and Dongbin Xiu. Data driven governing equations approximation using deep neural networks. *Journal of Computational Physics*, 395:620–635, October 2019. ISSN 0021-9991. doi:[10.1016/j.jcp.2019.06.042](https://doi.org/10.1016/j.jcp.2019.06.042).
- [173] Tong Qin, Zhen Chen, John D. Jakeman, and Dongbin Xiu. Data-driven learning of nonautonomous systems. *SIAM Journal on Scientific Computing*, 43(3):A1607–A1624, January 2021. ISSN 1095-7197. doi:[10.1137/20m1342859](https://doi.org/10.1137/20m1342859).
- [174] Yiwei Qiu, Jin Lin, Xiaoshuang Chen, Feng Liu, and Yonghua Song. Nonintrusive uncertainty quantification of dynamic power systems subject to stochastic excitations. *IEEE Transactions on Power Systems*, 36(1):402–414, January 2021. ISSN 1558-0679. doi:[10.1109/tpwrs.2020.3007746](https://doi.org/10.1109/tpwrs.2020.3007746).
- [175] Alejandro F. Queiruga, N. Benjamin Erichson, Dane Taylor, and Michael W. Mahoney. Continuous-in-depth neural networks, 2020.
- [176] M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, February 2019. ISSN 0021-9991. doi:[10.1016/j.jcp.2018.10.045](https://doi.org/10.1016/j.jcp.2018.10.045).
- [177] Christoph Reisinger and Gabriel Wittum. Efficient hierarchical approximation of high-dimensional option pricing problems. *SIAM Journal on Scientific Computing*, 29(1):440–458, January 2007. ISSN 1095-7197. doi:[10.1137/060649616](https://doi.org/10.1137/060649616).
- [178] Hannes Risken. *Fokker-Planck Equation*, pages 63–95. Springer Berlin Heidelberg, 1996. ISBN 9783642615443. doi:[10.1007/978-3-642-61544-3_4](https://doi.org/10.1007/978-3-642-61544-3_4).

- [179] Hannes Risken. *Fokker-Planck Equation for Several Variables; Methods of Solution*, pages 133–162. Springer Berlin Heidelberg, 1996. ISBN 9783642615443. doi:[10.1007/978-3-642-61544-3_6](https://doi.org/10.1007/978-3-642-61544-3_6).
- [180] Abram Rodgers, Alec Dektor, and Daniele Venturi. Adaptive integration of nonlinear evolution equations on tensor manifolds. *Journal of Scientific Computing*, 92(2), June 2022. ISSN 1573-7691. doi:[10.1007/s10915-022-01868-x](https://doi.org/10.1007/s10915-022-01868-x).
- [181] Houman Owhadi Roger Ghanem, David Higdon. Handbook of uncertainty quantification, 2017.
- [182] Jacob Roth, David A. Barajas-Solano, Panos Stinis, Jonathan Weare, and Mihai Anitescu. A kinetic monte carlo approach for simulating cascading transmission line failure. *Multiscale Modeling and Simulation*, 19(1):208–241, January 2021. ISSN 1540-3467. doi:[10.1137/19m1306865](https://doi.org/10.1137/19m1306865).
- [183] J. M. Sanz-Serna. Symplectic runge–kutta schemes for adjoint equations, automatic differentiation, optimal control, and more. *SIAM Review*, 58(1):3–33, January 2016. ISSN 1095-7200. doi:[10.1137/151002769](https://doi.org/10.1137/151002769).
- [184] PETER J. SCHMID. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, July 2010. ISSN 1469-7645. doi:[10.1017/s0022112010001217](https://doi.org/10.1017/s0022112010001217).
- [185] Timo Schorlepp, Shanyin Tong, Tobias Grafke, and Georg Stadler. Scalable methods for computing sharp extreme event probabilities in infinite-dimensional stochastic systems. *Statistics and Computing*, 33(6), October 2023. ISSN 1573-1375. doi:[10.1007/s11222-023-10307-2](https://doi.org/10.1007/s11222-023-10307-2).
- [186] DAVID W. SCOTT. Feasibility of multivariate density estimates. *Biometrika*, 78(1):197–205, 1991. ISSN 1464-3510. doi:[10.1093/biomet/78.1.197](https://doi.org/10.1093/biomet/78.1.197).

- [187] Liang Xu Seon K. Park, editor. *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications*. Springer Berlin Heidelberg, 2009. ISBN 9783540710561. doi:[10.1007/978-3-540-71056-1](https://doi.org/10.1007/978-3-540-71056-1).
- [188] Lawrence F. Shampine and Mark W. Reichelt. The matlab ode suite. *SIAM Journal on Scientific Computing*, 18(1):1–22, January 1997. ISSN 1095-7197. doi:[10.1137/s1064827594276424](https://doi.org/10.1137/s1064827594276424).
- [189] S. Shapiro. Lagrange and mayer problems in optimal control. *Automatica*, 3(3):219–230, 1966. ISSN 0005-1098. doi:[https://doi.org/10.1016/0005-1098\(66\)90014-8](https://doi.org/10.1016/0005-1098(66)90014-8).
- [190] Jie Shen, Tao Tang, and Li-Lian Wang. *Spectral Methods: Algorithms, Analysis and Applications*. Springer Berlin Heidelberg, 2011. ISBN 9783540710417. doi:[10.1007/978-3-540-71041-7](https://doi.org/10.1007/978-3-540-71041-7).
- [191] Wanan Sheng. A revisit of navier–stokes equation. *European Journal of Mechanics - B/Fluids*, 80:60–71, March 2020. ISSN 0997-7546. doi:[10.1016/j.euromechflu.2019.12.005](https://doi.org/10.1016/j.euromechflu.2019.12.005).
- [192] Sungho Shin, Timm Faulwasser, Mario Zanon, and Victor M. Zavala. A parallel decomposition scheme for solving long-horizon optimal control problems. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 5264–5271. IEEE, December 2019. doi:[10.1109/cdc40024.2019.9030139](https://doi.org/10.1109/cdc40024.2019.9030139).
- [193] Sungho Shin, Mihai Anitescu, and Victor M. Zavala. Overlapping schwarz decomposition for constrained quadratic programs. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 3004–3009. IEEE, December 2020. doi:[10.1109/cdc42340.2020.9304139](https://doi.org/10.1109/cdc42340.2020.9304139).
- [194] Lawrence Sirovich. Turbulence and the dynamics of coherent structures. ii. symmetries

- and transformations. *Quarterly of Applied Mathematics*, 45(3):573–582, October 1987. ISSN 1552-4485. doi:[10.1090/qam/910463](https://doi.org/10.1090/qam/910463).
- [195] Hermann Sohr. *The Stationary Navier-Stokes Equations*, pages 107–184. Birkhäuser Basel, 2001. ISBN 9783034882552. doi:[10.1007/978-3-0348-8255-2_3](https://doi.org/10.1007/978-3-0348-8255-2_3).
- [196] Eduardo D. Sontag. *Mathematical Control Theory*. Springer New York, 1998. ISBN 9781461205777. doi:[10.1007/978-1-4612-0577-7](https://doi.org/10.1007/978-1-4612-0577-7).
- [197] M. Studený and J. Vejnarová. *The Multiinformation Function as a Tool for Measuring Stochastic Dependence*, pages 261–297. Springer Netherlands, 1998. ISBN 9789401150149. doi:[10.1007/978-94-011-5014-9_10](https://doi.org/10.1007/978-94-011-5014-9_10).
- [198] Ellad B. Tadmor, Ronald E. Miller, and Ryan S. Elliott. *Continuum Mechanics and Thermodynamics: From Fundamental Concepts to Governing Equations*. Cambridge University Press, December 2011. ISBN 9781139017657. doi:[10.1017/cbo9781139017657](https://doi.org/10.1017/cbo9781139017657).
- [199] Daniel M. Tartakovsky and Pierre A. Gremaud. *Method of Distributions for Uncertainty Quantification*, pages 763–783. Springer International Publishing, 2017. ISBN 9783319123851. doi:[10.1007/978-3-319-12385-1_27](https://doi.org/10.1007/978-3-319-12385-1_27).
- [200] Søren Taverniers and Daniel M. Tartakovsky. Estimation of distributions via multilevel monte carlo with stratified sampling. *Journal of Computational Physics*, 419:109572, October 2020. ISSN 0021-9991. doi:[10.1016/j.jcp.2020.109572](https://doi.org/10.1016/j.jcp.2020.109572).
- [201] Guilherme Tegoni Goedert and Luca Biferale. Instanton-based importance sampling for extreme fluctuations in a shell model for turbulent energy cascade. *The European Physical Journal E*, 46(12), December 2023. ISSN 1292-895X. doi:[10.1140/epje/s10189-023-00392-6](https://doi.org/10.1140/epje/s10189-023-00392-6).

- [202] Shanyin Tong and Georg Stadler. Large deviation theory-based adaptive importance sampling for rare events in high dimensions. *SIAM/ASA Journal on Uncertainty Quantification*, 11(3):788–813, July 2023. ISSN 2166-2525. doi:[10.1137/22m1524758](https://doi.org/10.1137/22m1524758).
- [203] Fredi Tröltzsch. *Optimal Control of Partial Differential Equations*. American Mathematical Society, April 2010. ISBN 9781470411749. doi:[10.1090/gsm/112](https://doi.org/10.1090/gsm/112).
- [204] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer New York, 2009. ISBN 9780387790527. doi:[10.1007/b13794](https://doi.org/10.1007/b13794).
- [205] Stefan Ulbrich. *Analytical Background and Optimality Theory*, pages 1–95. Springer Netherlands, 2009. ISBN 9781402088391. doi:[10.1007/978-1-4020-8839-1_1](https://doi.org/10.1007/978-1-4020-8839-1_1).
- [206] D. Venturi, T. P. Sapsis, H. Cho, and G. E. Karniadakis. A computable evolution equation for the joint response-excitation probability density function of stochastic dynamical systems. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 468(2139):759–783, November 2011. ISSN 1471-2946. doi:[10.1098/rspa.2011.0186](https://doi.org/10.1098/rspa.2011.0186).
- [207] Robin Verschueren, Mario Zanon, Rien Quirynen, and Moritz Diehl. A sparsity preserving convexification procedure for indefinite quadratic programs arising in direct optimal control. *SIAM Journal on Optimization*, 27(3):2085–2109, January 2017. ISSN 1095-7189. doi:[10.1137/16m1081543](https://doi.org/10.1137/16m1081543).
- [208] Claire Voisin. *Hodge Theory and Complex Algebraic Geometry II*. Cambridge University Press, July 2003. ISBN 9780511615177. doi:[10.1017/cbo9780511615177](https://doi.org/10.1017/cbo9780511615177).
- [209] P. Wang, D. M. Tartakovsky, K. D. Jarman, and A. M. Tartakovsky. Cdf solutions of buckley–leverett equation with uncertain parameters. *Multiscale Modeling and Simulation*, 11(1):118–133, January 2013. ISSN 1540-3467. doi:[10.1137/120865574](https://doi.org/10.1137/120865574).

- [210] Peng Wang, Alexandre M. Tartakovsky, and Daniel M. Tartakovsky. Probability density function method for langevin equations with colored noise. *Physical Review Letters*, 110(14):140602, April 2013. ISSN 1079-7114. doi:[10.1103/physrevlett.110.140602](https://doi.org/10.1103/physrevlett.110.140602).
- [211] Peng Wang, Alexandre M. Tartakovsky, and Daniel M. Tartakovsky. Probability density function method for langevin equations with colored noise. *Physical Review Letters*, 110(14):140602, April 2013. ISSN 1079-7114. doi:[10.1103/physrevlett.110.140602](https://doi.org/10.1103/physrevlett.110.140602).
- [212] Karen E. Willcox, Omar Ghattas, and Patrick Heimbach. The imperative of physics-based modeling and inverse theory in computational science. *Nature Computational Science*, 1(3):166–168, March 2021. ISSN 2662-8457. doi:[10.1038/s43588-021-00040-z](https://doi.org/10.1038/s43588-021-00040-z).
- [213] Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, April 2005. ISSN 1436-4646. doi:[10.1007/s10107-004-0559-y](https://doi.org/10.1007/s10107-004-0559-y).
- [214] Dongbin Xiu. *Numerical Methods for Stochastic Computations: A Spectral Method Approach*. Princeton University Press, July 2010. ISBN 9780691142128. doi:[10.2307/j.ctv7h0skv](https://doi.org/10.2307/j.ctv7h0skv).
- [215] Dongbin Xiu and George Em Karniadakis. The wiener–askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computing*, 24(2):619–644, January 2002. ISSN 1095-7197. doi:[10.1137/s1064827501387826](https://doi.org/10.1137/s1064827501387826).
- [216] Keyulu Xu, Mozhi Zhang, Jingling Li, Simon Shaolei Du, Ken-Ichi Kawarabayashi, and Stefanie Jegelka. How neural networks extrapolate: From feedforward to graph neural networks. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=UH-cmocLJC>.
- [217] Wanting Xu and Mihai Anitescu. Exponentially accurate temporal decomposition for

- long-horizon linear-quadratic dynamic optimization. *SIAM Journal on Optimization*, 28(3):2541–2573, January 2018. ISSN 1095-7189. doi:[10.1137/16m1081993](https://doi.org/10.1137/16m1081993).
- [218] Haruo Yanai, Kei Takeuchi, and Yoshio Takane. *Projection Matrices*, pages 25–54. Springer New York, 2011. ISBN 9781441998873. doi:[10.1007/978-1-4419-9887-3_2](https://doi.org/10.1007/978-1-4419-9887-3_2).
- [219] Hyung Jun Yang, Hamdi A. Tchelepi, and Daniel M. Tartakovsky. Method of distributions for two-phase flow in heterogeneous porous media. *Water Resources Research*, 58(12), December 2022. ISSN 1944-7973. doi:[10.1029/2022wr032607](https://doi.org/10.1029/2022wr032607).
- [220] Yafei Yang, Xiaohong Guan, and Qiaozhu Zhai. Fast grid security assessment with n-k contingencies. In *2017 IEEE Power and Energy Society General Meeting*, pages 1–1. IEEE, July 2017. doi:[10.1109/pesgm.2017.8274108](https://doi.org/10.1109/pesgm.2017.8274108).
- [221] Pei Duo Yu, Chee Wei Tan, and Hung-Lin Fu. Graph algorithms for preventing cascading failures in networks. In *2018 52nd Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE, March 2018. doi:[10.1109/ciss.2018.8362273](https://doi.org/10.1109/ciss.2018.8362273).
- [222] Pei Duo Yu, Chee Wei Tan, and Hung-Lin Fu. Graph algorithms for preventing cascading failures in networks. In *2018 52nd Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE, March 2018. doi:[10.1109/ciss.2018.8362273](https://doi.org/10.1109/ciss.2018.8362273).
- [223] V.A. Zaitsev. Criteria for uniform complete controllability of a linear system. *Vestnik Udmurtskogo Universiteta. Matematika. Mekhanika. Komp'yuternye Nauki*, 25(2):157–179, June 2015. ISSN 2076-5959. doi:[10.20537/vm150202](https://doi.org/10.20537/vm150202).
- [224] Eberhard Zeidler. *Nonlinear Functional Analysis and Its Applications*. Springer New York, 1990. ISBN 9781461209850. doi:[10.1007/978-1-4612-0985-0](https://doi.org/10.1007/978-1-4612-0985-0).
- [225] Fuzhen Zhang. *The Schur Complement and Its Applications*. Springer-Verlag, 2005. ISBN 0387242716. doi:[10.1007/b105056](https://doi.org/10.1007/b105056).

- [226] Hao Zhang, Clarence W. Rowley, Eric A. Deem, and Louis N. Cattafesta. On-line dynamic mode decomposition for time-varying systems. *SIAM Journal on Applied Dynamical Systems*, 18(3):1586–1609, January 2019. ISSN 1536-0040. doi:[10.1137/18m1192329](https://doi.org/10.1137/18m1192329).
- [227] Ting Zhang and Wei Biao Wu. Time-varying nonlinear regression models: Nonparametric estimation and model selection. *The Annals of Statistics*, 43(2), April 2015. ISSN 0090-5364. doi:[10.1214/14-aos1299](https://doi.org/10.1214/14-aos1299).
- [228] Hongli Zhao and Daniel M. Tartakovsky. Discovery of linear representations for nonautonomous translation-invariant problems. *SIAM Journal on Scientific Computing*, 47(4):A2074–A2097, July 2025. ISSN 1095-7197. doi:[10.1137/24m1632905](https://doi.org/10.1137/24m1632905).
- [229] Hongli Zhao, Tyler E. Maltba, D. Adrian Maldonado, Emil Constantinescu, and Mihai Anitescu. Data-driven estimation of failure probabilities in correlated structure-preserving stochastic power system models, 2024.
- [230] Hongli Zhao, Mihai Anitescu, and Sen Na. Overlapping schwarz scheme for linear-quadratic programs in continuous time, 2025.
- [231] Honghao Zheng and Christopher L. DeMarco. A bi-stable branch model for energy-based cascading failure analysis in power systems. In *North American Power Symposium 2010*, pages 1–7. IEEE, September 2010. doi:[10.1109/naps.2010.5618968](https://doi.org/10.1109/naps.2010.5618968).
- [232] Shengyuan Zhong, Jun Zhao, Wenjia Li, Hao Li, Shuai Deng, Yang Li, Sajjad Hussain, Xiaoyuan Wang, and Jiebei Zhu. Quantitative analysis of information interaction in building energy systems based on mutual information. *Energy*, 214:118867, January 2021. ISSN 0360-5442. doi:[10.1016/j.energy.2020.118867](https://doi.org/10.1016/j.energy.2020.118867).
- [233] Min Zhu, Handi Zhang, Anran Jiao, George Em Karniadakis, and Lu Lu. Reliable extrapolation of deep neural operators informed by physics or sparse observations. *Com-*

puter Methods in Applied Mechanics and Engineering, 412:116064, July 2023. ISSN 0045-7825. doi:[10.1016/j.cma.2023.116064](https://doi.org/10.1016/j.cma.2023.116064).

- [234] Ray Daniel Zimmerman, Carlos Edmundo Murillo-Sanchez, and Robert John Thomas. Matpower: Steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Transactions on Power Systems*, 26(1):12–19, February 2011. ISSN 1558-0679. doi:[10.1109/tpwrs.2010.2051168](https://doi.org/10.1109/tpwrs.2010.2051168).