RESEARCH ARTICLE

Rigorous data-driven computation of spectral properties of Koopman operators for dynamical systems

Matthew J. Colbrook¹ | Alex Townsend²

¹Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, UK

²Department of Mathematics, Cornell University, Ithaca, New York, USA

Correspondence

Matthew J. Colbrook, Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, UK. Email: m.colbrook@damtp.cam.ac.uk

Alex Townsend, Department of Mathematics, Cornell University, Ithaca, New York, USA. Email: townsend@cornell.edu

Funding information

National Science Foundation, Grant/Award Numbers: DMS-1818757, DMS-1952757, DMS-2045646; Fondation Sciences Mathématiques de Paris Postdoctoral Fellowship; LMS Cecil King Travel Scholarship

Abstract

Koopman operators are infinite-dimensional operators that globally linearize nonlinear dynamical systems, making their spectral information valuable for understanding dynamics. However, Koopman operators can have continuous spectra and infinite-dimensional invariant subspaces, making computing their spectral information a considerable challenge. This paper describes data-driven algorithms with rigorous convergence guarantees for computing spectral information of Koopman operators from trajectory data. We introduce residual dynamic mode decomposition (ResDMD), which provides the first scheme for computing the spectra and pseudospectra of general Koopman operators from snapshot data without spectral pollution. Using the resolvent operator and ResDMD, we compute smoothed approximations of spectral measures associated with general measure-preserving dynamical systems. We prove explicit convergence theorems for our algorithms (including for general systems that are not measure-preserving), which can achieve high-order convergence even for chaotic systems when computing the density of the continuous spectrum and the discrete

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

^{© 2023} The Authors. Communications on Pure and Applied Mathematics published by Courant Institute of Mathematics and Wiley Periodicals LLC.

spectrum. Since our algorithms have error control, Res-DMD allows aposteri verification of spectral quantities, Koopman mode decompositions, and learned dictionaries. We demonstrate our algorithms on the tent map, circle rotations, Gauss iterated map, nonlinear pendulum, double pendulum, and Lorenz system. Finally, we provide kernelized variants of our algorithms for dynamical systems with a high-dimensional state space. This allows us to compute the spectral measure associated with the dynamics of a protein molecule with a 20,046-dimensional state space and compute nonlinear Koopman modes with error bounds for turbulent flow past aerofoils with Reynolds number > 10^5 that has a 295,122-dimensional state space.

KEYWORDS

dynamical systems, Koopman operator, data-driven discovery, dynamic mode decomposition, spectral theory

Mathematics Subject Classification (MSC) 2020 37M10, 65P99, 65F99, 65T99, 37A30, 47A10, 47B33, 37N10, 37N25

CONTENTS

1.	INTRODUCTION		2		
2.	SPECTRAL MEASURES, SPECTRA AND PSEUDOSPECTRA		8		
3.	COMPUTING SPECTRAL MEASURES FROM TRAJECTORY DATA		12		
4.	RESIDUAL DMD (ResDMD)		26		
5.	COMPUTING SPECTRAL MEASURES USING A ResDMD FRAMEWORK		34		
6.	HIGH-DIMENSIONAL DYNAMICAL SYSTEMS		45		
A	CKNOWLEDGMENTS		51		
REFERENCES					
APPENDIX A: A USEFUL CRITERION FOR <i>m</i> th ORDER KERNELS					
A	PPENDIX B: ResDMD FOR COMPUTING SPECTRA AND PSEUDOSPECTRA		57		
APPENDIX C: COMPUTING RATIONAL KERNEL CONVOLUTIONS WITH ERROR					
	CONTROL	•	61		

1 | INTRODUCTION

Dynamical systems are a mathematical description of time-dependent states or quantities that characterize evolving processes in classical mechanics, electrical circuits, fluid flows, climatology, finance, neuroscience, epidemiology, and many other fields. Throughout the paper, we consider autonomous dynamical systems whose state evolves over a state-space $\Omega \subseteq \mathbb{R}^d$ in discrete time-steps according to a function $F : \Omega \to \Omega$. In other words, we consider dynamical systems of the

form

where \mathbf{x}_0 is a given initial condition. Such a dynamical system forms a trajectory of iterates $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, ...$ in Ω . We want to analyze such trajectories to answer questions about the system's behavior. The interaction between numerical analysis and dynamical systems theory has stimulated remarkable growth in the subject since the 1960s [47, 69, 85, 123]. With the arrival of big data [64], modern statistical learning [61], and machine learning [95], data-driven algorithms are now becoming increasingly important in understanding dynamical systems [19, 111].

A classical viewpoint to analyze dynamical systems that originates in the seminal work of Poincaré [101] is to study fixed points and periodic orbits, as well as stable and unstable manifolds. Two fundamental challenges with Poincaré's geometric state-space viewpoint are:

- Nonlinear dynamics: To understand the stability of fixed points of nonlinear dynamical systems, one typically forms local models centered at these fixed points. Such models allow one to predict long-time dynamics in small neighborhoods of fixed points and attracting manifolds. However, they do not provide reasonable predictions for all initial conditions. A global understanding of nonlinear dynamics in state-space remains largely qualitative [24]. For example, Brunton and Kutz describe the global understanding of nonlinear dynamics in state-space as a "mathematical grand challenge of the 21st century" [19]. This paper focuses on the Koopman operator viewpoint of dynamical systems. Koopman operators provide a global linearization of (1.1) by studying the evolution of observables of the system.
- Unknown dynamics: For many applications, a system's dynamics may be too complicated to describe analytically, or we may have incomplete knowledge of its evolution. Typically, we can only acquire several sequences of iterates of (1.1) starting at different values of \boldsymbol{x}_0 . This constraint means that constructing local models can be impossible. In this paper, we focus on data-driven approaches to learning and analyzing the dynamical system with trajectories of iterates from (1.1).

Koopman operator theory, which dates back to Koopman and von Neumann [77, 78], is an alternative viewpoint to analyze a dynamical system that uses the space of scalar observable functions [93]. Its increasing popularity has led to the term "Koopmanism" [24] and thousands of articles over the last decade. Sparked by [89, 94], a reason for the recent attention is its use in data-driven methods for studying dynamical systems (see [23] for a review and the history). Some popular applications include fluid dynamics [54, 90, 107, 110], epidemiology [102], neuroscience [18], finance [87], robotics [13, 17], power grids [125, 126], and molecular dynamics[74, 97, 113, 114].

Let $g: \Omega \to \mathbb{C}$ be a function that one can use to indirectly measure the dynamical system's state in (1.1). Such a function g is often called an observable. Therefore, $g(\mathbf{x}_k)$ indirectly measures \mathbf{x}_k and $g(\mathbf{x}_{k+1}) = g(F(\mathbf{x}_k))$ measures the state one time-step later. Since we are interested in how the state evolves, we can define an observable $\mathcal{K}_g: \Omega \to \mathbb{C}$ such that $\mathcal{K}_g(\mathbf{x}) = g(F(\mathbf{x})) = (g \circ F)(\mathbf{x})$, which tells us how our measurement evolves from one time-step to the next. In this way, \mathcal{K}_g indirectly monitors the evolution of (1.1). After exhaustively using g to measure the dynamical system at various time steps, we might want to probe the dynamical system using a different observable. As one tries different observables to probe the dynamics, one wonders if there is a principled way to select "good" observables. This motivates studying the Koopman operator given by $\mathcal{K}: g \mapsto \mathcal{K}_g$. One typically works in the Hilbert space $L^2(\Omega, \omega)$ of observables for a positive measure ω on Ω .¹ To be precise, we consider $\mathcal{K} : \mathcal{D}(\mathcal{K}) \to L^2(\Omega, \omega)$, where $\mathcal{D}(\mathcal{K}) \subseteq L^2(\Omega, \omega)$ is a suitable domain of observables. That is,

$$[\mathcal{K}g](\mathbf{x}) = (g \circ F)(\mathbf{x}), \qquad \mathbf{x} \in \Omega, \qquad g \in \mathcal{D}(\mathcal{K}), \tag{1.2}$$

where the equality is understood in the $L^2(\Omega, \omega)$ sense. For a fixed g, $[\mathcal{K}g](\mathbf{x})$ measures the state of the dynamical system after one time-step if it is currently at \mathbf{x} . On the other hand, for a fixed \mathbf{x} , (1.2) is a linear composition operator with F. Since \mathcal{K} is a linear operator, regardless of whether the dynamics are linear or nonlinear, the spectral information of \mathcal{K} determines the behavior of the dynamical system (1.1) (e.g., see (2.5)). However, since \mathcal{K} is an infinite-dimensional operator, its spectral information can be far more complicated than that of a finite matrix and far more difficult to compute [11, 29, 31, 34, 35]. For example, \mathcal{K} can have both discrete² and continuous spectra [89].

Computing the spectral properties of \mathcal{K} is an active area of research, and several popular algorithms exist.³ Generalized Laplace analysis [90], which is related to the power method, uses prior information about the eigenvalues of \mathcal{K} to compute corresponding Koopman modes.⁴ Dynamic mode decomposition (DMD), which is related to the proper orthogonal decomposition, uses a linear model to fit trajectory data [81, 107, 110, 132], and extended DMD (EDMD) [75, 138, 139] is a Galerkin approximation of \mathcal{K} that uses a rich dictionary of observables (see Section 4.1.1). See also measure-preserving EDMD [32]. Other data-driven methods include deep learning [84, 86, 88, 98, 141], reduced-order modeling [12, 58], sparse identification of nonlinear dynamics [20, 108], and kernel analog forecasting [25, 55, 143]. However, remaining challenges include the following:

C1. **Continuous spectra**: Since the Koopman operator is infinite-dimensional, it can have a continuous spectrum. Continuous spectra must be treated with considerable care as discretizing \mathcal{K} destroys its presence. Several data-driven approaches are proposed to handle continuous spectra, such as HAVOK analysis [22], which applies DMD to a vector of time-delayed measurements, and feed-forward neural networks with an auxiliary network to parametrize the continuous spectrum [86]. While these methods are promising, they currently lack convergence guarantees. Most existing nonparametric approaches for computing continuous spectra of \mathcal{K} are restricted to ergodic systems [52], as this allows relevant integrals to be computed using long-time averages. For example, [3] approximates discrete spectra using harmonic averaging, applies an ergodic cleaning process, and then estimates the continuous spectrum using Welch's method [137]. Though these methods are related to well-studied techniques in signal processing [51], they can be challenging to apply in the presence of noise or if a pair of eigenvalues are close together. Moreover, they often rely on heuristic parameter

¹We do not assume that this measure is invariant, and a common choice of ω is the standard Lebesgue measure. This choice is natural for Hamiltonian systems for which the Koopman operator is unitary on $L^2(\Omega, \omega)$. For other systems, we can select ω according to the region where we wish to study the dynamics, such as a Gaussian measure. In some applications, ω corresponds to an unknown ergodic measure on an attractor.

² Throughout this paper, we use the term "discrete spectra" to mean the eigenvalues of \mathcal{K} , also known as the point spectrum. This also includes embedded eigenvalues, in contrast to the usual definition of discrete spectrum.

³ Computing spectral properties of Perron–Frobenius or transfer operators is another area of active research [41, 49, 50, 66]. On appropriate spaces, which are often not Hilbert spaces, the Perron–Frobenius operator is the adjoint of the Koopman operator. One can extend some of our notions to Banach spaces [115]; however, a difficulty is that the approximation of finite matrix norms can be NP-hard [62].

⁴ These are vectors to reconstruct the system's state as a linear combination of the candidate Koopman eigenfunctions.

choices and cleanup procedures. Other approaches for ergodic systems include approximating the functional calculus using integral operator compact regularizations [40], and approximating the moments of the spectral measure using ergodicity followed by separating the atomic and continuous parts of the spectrum using the Christoffel–Darboux kernel [80]. In this paper, we do not assume ergodicity, and we develop a computational framework that deals jointly with continuous and discrete spectra via computing smoothed approximations of spectral measures with explicit high-order convergence rates.

- C2. Invariant subspaces: A finite-dimensional invariant subspace of \mathcal{K} is a space of observables $\mathcal{G} = \operatorname{span}\{g_1, \dots, g_k\}$ such that $\mathcal{K}g \in \mathcal{G}$ for all $g \in \mathcal{G}$. A common assumption in the literature is that \mathcal{K} has a finite-dimensional nontrivial⁵ invariant space, which may not be the case (e.g., when the system is mixing). Even if a (nontrivial) finite-dimensional invariant subspace exists, it can be challenging to compute or may not capture all of the dynamics of interest. Often, one must settle for approximate invariant subspaces, and methods such as EDMD generally result in closure issues [21]. In this paper, we develop methods that directly compute spectral properties of \mathcal{K} instead of restrictions of \mathcal{K} to finite-dimensional subspaces.
- C3. **Spectral pollution**: A well-known difficulty of computing spectra of infinite-dimensional operators is spectral pollution, where discretizations cause spurious eigenvalues that have nothing to do with the operator [36, 37, 83]. Methods such as EDMD suffer from spectral pollution [138], and heuristics are common to reduce a user's concern. One can compare eigenvalues computed using different discretization sizes or verify that a candidate eigenfunction behaves linearly on the trajectories of (1.1) in the way that the corresponding eigenvalue predicts [68]. In some cases, it is possible to approximate the spectral information of a Koopman operator without spectral pollution. For example, when performing a Krylov subspace method on a finite-dimensional invariant subspace of known dimension, the Hankel-DMD algorithm computes the corresponding eigenvalues for ergodic systems [2]. Instead, it is highly desirable to have a principled way of detecting spectral pollution with as few assumptions as possible. In this paper, we develop methods to compute residuals associated with the spectrum with error control, allowing us to compute spectra of general Koopman operators without spectral pollution.
- C4. **Chaotic behavior**: Many dynamical systems exhibit rich, chaotic behavior. If a system is chaotic, small perturbations to the initial state x_0 can lead to large perturbations of the state later. Koopman operators with continuous spectra are a generic feature when the underlying dynamics are chaotic [3, 7]. Chaotic behavior can make data-driven approaches challenging when using long-time trajectories.⁶ In this paper, we use one-step trajectory data, called snapshot data, that can take the form of long or short trajectories depending on the system and available data.
- C5. **Nonlinearity and high-dimensional state-space**: For many dynamical systems, the corresponding dynamical system is strongly nonlinear and has a very large state-space dimension. This can make choosing a good set of observables challenging. Often this choice takes the form of a learned dictionary. In this paper, our algorithms come with verification (i.e., error bounds), which allows us to perform a-posteriori verification that the learned dictionary was reasonable.

⁵ If $\omega(\Omega) < \infty$, the constant function 1 generates an invariant subspace, but this is not interesting for studying dynamics.

⁶ Many dynamical systems, such as hyperbolic systems with chaotic dynamics, satisfy shadowing theorems, meaning that numerical trajectories remain close to true trajectories. Under suitable conditions, we can still apply quadrature rules such as (3.3).

ALGORITHM 1 A computational framework for recovering an approximation of the spectral measure ν_g associated with a Koopman operator that is an isometry.

Input: Trajectory data, a filter φ , and an observable $g \in L^2(\Omega, \omega)$.

- 1: Approximate the autocorrelations $a_n = \frac{1}{2\pi} \langle g, \mathcal{K}^n g \rangle$ for $0 \le n \le N$. (The precise value of *N* and the approach depends on the trajectory data (see Section 3.1).)
- 2: Set $a_{-n} = \overline{a_n}$ for $1 \le n \le N$.

Output: The function $\nu_{g,N}(\theta) = \sum_{n=-N}^{N} \varphi(\frac{n}{N}) a_n e^{in\theta}$ that can be evaluated for any $\theta \in [-\pi, \pi]_{\text{per}}$.

In response to these challenges, we provide rigorously convergent and data-driven algorithms to compute the spectral properties of Koopman operators. Throughout the paper, we only assume that \mathcal{K} is a closed and densely defined operator, which is an essential assumption before talking about spectral information of \mathcal{K} ; otherwise, the spectrum is the whole of \mathbb{C} . In particular, we do not assume that \mathcal{K} has a nontrivial finite-dimensional invariant subspace or that it only has a discrete spectrum.

1.1 | Novel contributions

To tackle **C1**, we give a computational framework for computing the spectral measures of Koopman operators associated with measure-preserving dynamical systems. Our algorithms can achieve arbitrarily high-order convergence. We deal with discrete and continuous spectra of a Koopman operator \mathcal{K} by calculating them together, and we prove rigorous convergence results:

- Pointwise recovery of spectral densities of the continuous part of the spectrum of \mathcal{K} (see Theorem 3.1).
- Weak convergence, involving integration against test functions (see Theorem 3.2).
- The recovery of the eigenvalues of \mathcal{K} (see Theorem 3.3).

Algorithm 1 is based on estimating autocorrelations and filtering, which requires long-time trajectory data from (1.1). Using a connection with the resolvent operator, we develop an alternative rational-based approach that can achieve high-order convergence on short-time trajectory data (see Algorithm 4) while being robust to noise (see Section 5.2.2).

To tackle **C2** and **C3**, we derive a new DMD-type algorithm called residual DMD (ResDMD). ResDMD constructs Galerkin approximations for not only \mathcal{K} but also $\mathcal{K}^*\mathcal{K}$. This key difference (which requires no extra data) allows us to rigorously compute spectra of general Koopman operators (without the measure-preserving assumption) together with residuals. ResDMD goes a long way to overcoming challenges **C2** and **C3** by allowing us to:

- Remove the spectral pollution of extended DMD (see Algorithm 2 and 4.1).
- Compute spectra and pseudospectra with convergence guarantees (see Algorithm 3 and Appendix B).

It is important to stress that ResDMD computes residuals associated with the underlying infinitedimensional operator \mathcal{K} . This is in contrast to works that compute finite-dimensional residuals of observables with respect to a finite DMD discretization matrix [42]. In contrast to ResDMD (which computes residuals with respect to \mathcal{K}), residuals with respect to a finite DMD discretization

ALGORITHM 2 ResDMD for computing eigenpairs without spectral pollution.

Input: Snapshot data $\{\boldsymbol{x}^{(j)}\}_{j=1}^{M}, \{\boldsymbol{y}^{(j)}\}_{j=1}^{M}$ (such that $\boldsymbol{y}^{(j)} = F(\boldsymbol{x}^{(j)})$), quadrature weights $\{w_j\}_{j=1}^{M}$, a dictionary of observables $\{\psi_j\}_{j=1}^{N_K}$ and an accuracy goal $\epsilon > 0$.

- 1: Compute $\Psi_X^* W \Psi_X$, $\Psi_X^* W \Psi_Y$, and $\Psi_Y^* W \Psi_Y$, where Ψ_X and Ψ_Y are given in (4.4).
- 2: Solve $(\Psi_X^* W \Psi_Y) \mathbf{g} = \lambda (\Psi_X^* W \Psi_X) \mathbf{g}$ for eigenpairs $\{(\lambda_j, g_{(j)} = \Psi \mathbf{g}_j)\}$.
- 3: Compute $\operatorname{res}(\lambda_j, g_{(j)})$ for all *j* (see (4.6)) and discard if $\operatorname{res}(\lambda_j, g_{(j)}) > \epsilon$.

Output: A collection of accurate eigenpairs $\{(\lambda_j, \boldsymbol{g}_j) : \operatorname{res}(\lambda_j, \boldsymbol{g}_{(j)}) \le \epsilon\}$.

ALGORITHM 3 ResDMD for estimating ϵ -pseudospectra.

Input: Snapshot data $\{\boldsymbol{x}^{(j)}\}_{j=1}^{M}, \{\boldsymbol{y}^{(j)}\}_{j=1}^{M}$ (such that $\boldsymbol{y}^{(j)} = F(\boldsymbol{x}^{(j)})$), quadrature weights $\{w_j\}_{j=1}^{M}$, a dictionary of observables $\{\psi_j\}_{i=1}^{N_k}$, an accuracy goal $\varepsilon > 0$, and a grid $z_1, \dots, z_k \in \mathbb{C}$ (see (B.5)).

- 1: Compute $\Psi_X^* W \Psi_X, \Psi_X^* W \Psi_Y$, and $\Psi_Y^* W \Psi_Y$, where Ψ_X and Ψ_Y are given in (4.4).
- 2: For each z_j , compute $\tau_j = \min_{\boldsymbol{g} \in \mathbb{C}^{N_K}} \operatorname{res}(z_j, \Psi \boldsymbol{g})$ (see (4.6)), which is a generalized SVD problem, and the corresponding singular vectors \boldsymbol{g}_j .

Output: Estimate of the ϵ -pseudospectrum { $z_i : \tau_i < \epsilon$ } and approximate eigenfunctions { $g_i : \tau_i < \epsilon$ }.

ALGORITHM 4 A computational framework for evaluating an approximate spectral measure with respect to $g \in L^2(\Omega, \omega)$ at $\{\theta_k\}_{k=1}^p \subset [-\pi, \pi]_{per}$ of an isometry \mathcal{K} using snapshot data.

Input: Snapshot data $\{\mathbf{x}^{(j)}\}_{j=1}^{M}, \{\mathbf{y}^{(j)}\}_{j=1}^{M}$ (such that $\mathbf{y}^{(j)} = F(\mathbf{x}^{(j)})$), quadrature weights $\{w_j\}_{j=1}^{M}$, a dictionary of observables $\{\psi_j\}_{j=1}^{N_K}, m \in \mathbb{N}$, smoothing parameter $0 < \epsilon < 1$ (accuracy goal is ϵ^m), distinct points $\{z_j\}_{j=1}^m \subset \mathbb{C}$ with $\operatorname{Re}(z_j) > 0$ (recommended choice is (5.12)), and evaluation points $\{\theta_k\}_{k=1}^{P} \subset [-\pi, \pi]_{\operatorname{per}}$.

- 1: Solve (5.8) and (5.11) for $c_1(\epsilon), \dots, c_m(\epsilon) \in \mathbb{C}$ and $d_1, \dots, d_m \in \mathbb{C}$, respectively.
- 2: Compute $\Psi_X^* W \Psi_X$ and $\Psi_X^* W \Psi_Y$, where Ψ_X and Ψ_Y are given in (4.4).
- 3: Compute a generalized Schur decomposition of $\Psi_X^* W \Psi_Y$ and $\Psi_X^* W \Psi_X$, that is, $\Psi_X^* W \Psi_Y = QSZ^*$ and $\Psi_X^* W \Psi_X = QTZ^*$, where Q, Z are unitary and S, T are upper triangular.
- 4: Compute **a** in (5.13) and $v_1 = TZ^*a$, $v_2 = T^*Q^*a$, and $v_3 = S^*Q^*a$.

5: **for**
$$k = 1, ..., P$$
 do

- 6: Compute $I_j = (S e^{i\theta_k}(1 + \epsilon z_j)T)^{-1}v_1$ for $1 \le j \le m$.
- 7: Compute $\nu_g^{\epsilon}(\theta_k) = \frac{-1}{2\pi} \sum_{j=1}^m \operatorname{Re}[c_j(\epsilon)e^{-i\theta_k}(1+\epsilon\overline{z_j})(I_j^*v_2) + d_j(v_3^*I_j)]$.
- 8: end for

Output: Values of the approximate spectral measure, that is, $\{v_g^{\epsilon}(\theta_k)\}_{k=1}^{p}$.

cannot give error bounds on the spectral information of \mathcal{K} and can suffer from issues such as spectral pollution.

All of our algorithms can be used for chaotic systems, and we do not run into the difficulties of **C4**. We also describe a kernelized variant of ResDMD(see Algorithm 5). We first use the kernelized EDMD [139] to learn a suitable dictionary of observables from a subset of the given data, and then we employ this dictionary in Algorithms 2 to 4. Compared with traditional kernelized approaches, our key advantage is that we have convergence theory and perform a posterior verification that the learned dictionary is suitable, allowing us to tackle **C5**.

Code for ResDMD and the examples of this paper is provided at https://github.com/ MColbrook/Residual-Dynamic-Mode-Decomposition. For fluid dynamical applications of the methods we develop in this paper, the reader is also invited to read [33].

ALGORITHM 5 A computational framework for kernelized versions of Algorithms 2 to 4.

Input: Snapshot data $\{\boldsymbol{x}^{(j)}, \boldsymbol{y}^{(j)}\}_{j=1}^{M'}$ and $\{\hat{\boldsymbol{x}}^{(j)}, \hat{\boldsymbol{y}}^{(j)}\}_{j=1}^{M''}$, positive-definite kernel function $S : \Omega \times \Omega \to \mathbb{R}$, and positive integer $N_{K}^{''} \leq M'$.

- 1: Apply kernel EDMD to $\{\mathbf{x}^{(j)}, \mathbf{y}^{(j)}\}_{j=1}^{M'}$ with kernel *S* to compute the matrices $\sqrt{W}\Psi_X\Psi_X^*\sqrt{W}$ and $\sqrt{W}\Psi_Y\Psi_X^*\sqrt{W}$ using (6.1) and the kernel trick.
- 2: Compute U and Σ from the eigendecomposition $\sqrt{W}\Psi_X\Psi_X^*\sqrt{W} = U\Sigma^2 U^*$.
- 3: Compute the dominant N''_K eigenvectors of $\widetilde{K}_{\text{EDMD}} = (\Sigma^{\dagger} U^*) \sqrt{W} \Psi_Y \Psi_X^* \sqrt{W} (U\Sigma^{\dagger})$ and stack them column-by-column into $Z \in \mathbb{C}^{M' \times N''_K}$.
- 4: Apply a QR decomposition to orthogonalize Z to $Q = \begin{bmatrix} Q_1 & \cdots & Q_{N''_k} \end{bmatrix} \in \mathbb{C}^{M' \times N''_k}$.
- 5: Apply Algorithms 2 to 4 with $\{\hat{x}^{(j)}, \hat{y}^{(j)}\}_{j=1}^{M''}$ and the dictionary $\{\psi_j\}_{j=1}^{N''_{\mathcal{K}}}$, where

$$\psi_j(\boldsymbol{x}) = \begin{bmatrix} \mathcal{S}(\boldsymbol{x}, \boldsymbol{x}^{(1)}) & \mathcal{S}(\boldsymbol{x}, \boldsymbol{x}^{(2)}) & \cdots & \mathcal{S}(\boldsymbol{x}, \boldsymbol{x}^{(M')}) \end{bmatrix} (U\Sigma^{\dagger}) Q_j, \qquad 1 \le j \le N_K''.$$

Output: Spectral properties of Koopman operator according to Algorithms 2 to 4.

1.2 | Paper structure

In Section 2, we provide background material to introduce spectral measures and pseudospectra of operators. In Section 3, we describe data-driven algorithms for computing the spectral measures of Koopman operators associated with measure-preserving dynamics when given long-time trajectory data. In Section 4, we present ResDMD for computing spectral properties of general Koopman operators. In Section 5, we compute spectral measures of \mathcal{K} associated with measure-preserving dynamics using snapshot data. Finally, in Section 6, we develop a kernelized approach and apply our algorithms to dynamical systems with high-dimensional state-spaces. Throughout, we use $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$ to denote the inner product and norm corresponding to $L^2(\Omega, \omega)$, respectively. We use $\sigma(T)$ to denote the spectrum of a linear operator T and $\|T\|$ to denote its operator norm if it is bounded.

2 | SPECTRAL MEASURES, SPECTRA AND PSEUDOSPECTRA

We are interested in the spectral measures, spectra, and pseudospectra of Koopman operators. This section introduces the relevant background for these spectral properties.

2.1 | Spectral measures for measure-preserving dynamical systems

Suppose that the associated dynamics is measure-preserving so that $\omega(E) = \omega(\{\mathbf{x} : F(\mathbf{x}) \in E\})$ for any Borel measurable subset $E \subset \Omega$. Equivalently, this means that the Koopman operator \mathcal{K} associated with the dynamical system in (1.1) is an isometry, that is, $||\mathcal{K}g|| = ||g||$ for all observables $g \in \mathcal{D}(\mathcal{K}) = L^2(\Omega, \omega)$. Dynamical systems such as Hamiltonian flows [4], geodesic flows on Riemannian manifolds [44, Chapter 5], Bernoulli schemes in probability theory [119], and ergodic systems [135] are all measuring-preserving. Moreover, many dynamical systems become measure-preserving in the long run [89].

Spectral measures provide a way of diagonalizing normal operators, including self-adjoint and unitary operators, even in the presence of continuous spectra. Unfortunately, a Koopman operator

that is an isometry does not necessarily commute with its adjoint (see Section 3.5.1). Therefore, we must consider a unitary extension (see Section 2.1.1) before defining a spectral measure (see Section 2.1.3) and Koopman mode decomposition (see Section 2.1.4).

2.1.1 | Unitary extensions of isometries

Given a Koopman operator \mathcal{K} of a measure-preserving dynamical system, we use the concept of unitary extension to formally construct a related normal operator \mathcal{K}' . That is, suppose that $\mathcal{K} : L^2(\Omega, \omega) \to L^2(\Omega, \omega)$ is an isometry, then there exists a unitary extension \mathcal{K}' defined on an extended Hilbert space \mathcal{H}' with $L^2(\Omega, \omega) \subset \mathcal{H}'$ [96, Proposition I.2.3].⁷ Even though such an extension is not unique, it allows us to understand the spectral information of \mathcal{K} by considering \mathcal{K}' , which is a normal operator. If F is invertible and measure-preserving, \mathcal{K} is unitary and we can simply take $\mathcal{K}' = \mathcal{K}$ and $\mathcal{H}' = L^2(\Omega, \omega)$.

2.1.2 | Spectral measures of normal operators

To explain the notion of spectral measures for a normal operator, we begin with the more familiar finite-dimensional case. Any linear operator acting on a finite-dimensional Hilbert space has a purely discrete spectrum consisting of eigenvalues. In particular, the spectral theorem for a normal matrix $A \in \mathbb{C}^{n \times n}$, that is, $A^*A = AA^*$, states that there exists an orthonormal basis of eigenvectors v_1, \dots, v_n for \mathbb{C}^n such that

$$v = \left(\sum_{k=1}^{n} v_k v_k^*\right) v, \quad v \in \mathbb{C}^n \quad \text{and} \quad Av = \left(\sum_{k=1}^{n} \lambda_k v_k v_k^*\right) v, \quad v \in \mathbb{C}^n, \tag{2.1}$$

where $\lambda_1, ..., \lambda_n$ are eigenvalues of A, that is, $Av_k = \lambda_k v_k$ for $1 \le k \le n$. In other words, the projections $v_k v_k^*$ simultaneously decompose the space \mathbb{C}^n and diagonalize the operator A. This intuition carries over to the infinite-dimensional setting by replacing $v \in \mathbb{C}^n$ by $f \in \mathcal{H}'$, and A by a normal operator \mathcal{K}' . However, the eigenvectors of \mathcal{K}' need not form a basis for \mathcal{H}' or diagonalize \mathcal{K}' . Instead, the spectral theorem for normal operators states that the projections $v_k v_k^*$ in (2.1) can be replaced by a projection-valued measure \mathcal{E} supported on the spectrum of \mathcal{K}' [104, Thm. VIII.6]. In our setting, \mathcal{K}' is unitary; hence, its spectrum is contained inside the unit circle \mathbb{T} . The measure \mathcal{E} assigns an orthogonal projector to each Borel measurable subset of \mathbb{T} such that

$$f = \left(\int_{\mathbb{T}} d\mathcal{E}(y)\right) f$$
 and $\mathcal{K}' f = \left(\int_{\mathbb{T}} y \, d\mathcal{E}(y)\right) f$, $f \in \mathcal{H}'$.

Analogous to (2.1), \mathcal{E} decomposes \mathcal{H}' and diagonalizes the operator \mathcal{K}' . For example, if $U \subset \mathbb{T}$ consists only of eigenvalues of \mathcal{K}' and no other types of spectra, then $\mathcal{E}(U)$ is simply the spectral projector onto the invariant subspace spanned by the corresponding eigenfunctions. More generally, \mathcal{E} decomposes elements of \mathcal{H}' along the discrete and continuous spectrum of \mathcal{K}' (see Section 2.1.4).

⁷ To see how to extend \mathcal{K} to a unitary operator \mathcal{K}' , consider the Wold–von Neumann decomposition [96, Theorem I.1.1]. This decomposition states that \mathcal{K} can be written as $\mathcal{K} = (\bigoplus_{\alpha \in I} S_{\alpha}) \oplus U$ for some index set I, where S_{α} is the unilateral shift on a Hilbert space \mathcal{H}_{α} and U is a unitary operator. Since one can extend any unilateral shift to a unitary bilateral shift, one can extend \mathcal{K} to a unitary operator \mathcal{K}' .

2.1.3 | Spectral measures of Koopman operators

Given an observable $g \in L^2(\Omega, \omega) \subset \mathcal{H}'$ of interest such that ||g|| = 1, the spectral measure of \mathcal{K}' with respect to g is a scalar measure defined as $\mu_g(U) := \langle \mathcal{E}(U)g, g \rangle$, where $U \subset \mathbb{T}$ is a Borel measurable set [104]. For plotting and visualization, it is more convenient to equivalently consider the corresponding probability measures ν_g defined on the periodic interval $[-\pi, \pi]_{per}$ after a change of variables $\lambda = \exp(i\theta)$ so that $d\mu_g(\lambda) = d\nu_g(\theta)$. Therefore, throughout this paper, we compute and visualize ν_g . We use the notation $\int_{[-\pi,\pi]_{per}} to denote integration along the periodic interval <math>[-\pi,\pi]_{per}$ as $\int_{-\pi}^{\pi}$ is ambiguous since spectral measures can have atoms at $\pm \pi$. The particular choice of g is up to the practitioner: smooth g makes ν_g easier to compute for nonsmooth g but can give better resolution of \mathcal{K} , whereas ν_g is more challenging to compute for nonsmooth g but can give better resolution of the underlying dynamics. In other situations, the application dictates that a particular g is interesting (see Section 6.2).

To compute ν_g , we start by noting that the Fourier coefficients of ν_g are given by

$$\hat{\nu_g}(n) := \frac{1}{2\pi} \int_{[-\pi,\pi]_{\text{per}}} e^{-in\theta} \, d\nu_g(\theta) = \frac{1}{2\pi} \int_{\mathbb{T}} \lambda^{-n} \, d\mu_g(\lambda) = \frac{1}{2\pi} \langle \mathcal{K}'^{-n}g,g\rangle, \qquad n \in \mathbb{Z}.$$
(2.2)

Since \mathcal{K}' is unitary, its inverse is its adjoint and thus, the Fourier coefficients of ν_g can be expressed in terms of correlations $\langle \mathcal{K}^n g, g \rangle$ and $\langle g, \mathcal{K}^n g \rangle$. That is, for $g \in L^2(\Omega, \omega)$,

$$\hat{\nu_g}(n) = \frac{1}{2\pi} \langle \mathcal{K}^{-n} g, g \rangle, \quad n < 0, \qquad \hat{\nu_g}(n) = \frac{1}{2\pi} \langle g, \mathcal{K}^n g \rangle, \quad n \ge 0.$$
(2.3)

Since $g \in L^2(\Omega, \omega)$ and (2.3) only depend on correlations with \mathcal{K} , and ν_g is determined by its Fourier coefficients, we find that ν_g is independent of the choice of unitary extension \mathcal{K}' . Henceforth, we can safely dispense with the extension \mathcal{K}' , and call ν_g the spectral measure of \mathcal{K} with respect to g.

From (2.3), we find that $\hat{\nu}_g(-n) = \overline{\hat{\nu}_g(n)}$ for $n \in \mathbb{Z}$, which tells us that ν_g is completely determined by the forward-time dynamical *autocorrelations* $\langle g, \mathcal{K}^n g \rangle$ with $n \ge 0$. Equivalently, the spectral measure of \mathcal{K} with respect to $g \in L^2(\Omega, \omega)$ is a signature for the forward-time dynamics of (1.1) if the closure of span $\{g, \mathcal{K}g, \mathcal{K}^2g, ...\}$ is $L^2(\Omega, \omega)$. If the closure of span $\{g, \mathcal{K}g, \mathcal{K}^2g, ...\}$ is $L^2(\Omega, \omega)$, then g is called cyclic. If g is not cyclic, then ν_g only determines the action of \mathcal{K} on the closure of span $\{g, \mathcal{K}g, \mathcal{K}^2g, ...\}$, which is often sufficient for many applications.

2.1.4 | Continuous and discrete parts of spectra, and Koopman mode decompositions

Of particular importance to dynamical systems is Lebesgue's decomposition of ν_g [121]:

$$d\nu_{g}(y) = \sum_{\substack{\lambda = \exp(i\theta) \in \sigma_{p}(\mathcal{K}) \\ \text{discrete part}}} \langle \mathcal{P}_{\lambda}g, g \rangle \, \delta(y - \theta) dy + \underbrace{\rho_{g}(y) \, dy + d\nu_{g}^{(sc)}(y)}_{\text{continuous part}}.$$
(2.4)

The discrete (or atomic) part of ν_g is a sum of Dirac delta distributions, supported on $\sigma_p(\mathcal{K})$, the set of eigenvalues of \mathcal{K} .⁸ The coefficient of each δ in the sum is $\langle \mathcal{P}_{\lambda}g,g \rangle = \|\mathcal{P}_{\lambda}g\|^2$, where

⁸ After mapping to the periodic interval, the discrete part of ν_g is supported on the closure of $\sigma_p(\mathcal{K}')$. However, we can always choose the extension \mathcal{K}' so that $\sigma_p(\mathcal{K}')=\sigma_p(\mathcal{K})$ with the same eigenspaces.

 \mathcal{P}_{λ} is the orthogonal spectral projector associated with the eigenvalue λ . The continuous part of ν_g consists of a part that is absolutely continuous with respect to the Lebesgue measure, with Radon–Nikodym derivative $\rho_g \in L^1([-\pi, \pi]_{\text{per}})$, and a singular continuous component $\nu_g^{(\text{sc})}$. The decomposition in (2.4) provides important information on the evolution of dynamical systems. For example, suppose that there is no singular continuous spectrum. Then any $g \in L^2(\Omega, \omega)$ can be written as

$$g = \sum_{\lambda \in \sigma_{\rm p}(\mathcal{K})} c_{\lambda} \varphi_{\lambda} + \int_{[-\pi,\pi]_{\rm per}} \phi_{\theta,g} \, d\theta,$$

where the φ_{λ} are the eigenfunctions of \mathcal{K} , c_{λ} are expansion coefficients and $\varphi_{\theta,g}$ is a "continuously parametrized" collection of eigenfunctions.⁹ Then, one obtains the Koopman mode decomposition [89]

$$g(\boldsymbol{x}_n) = [\mathcal{K}^n g](\boldsymbol{x}_0) = \sum_{\lambda \in \sigma_p(\mathcal{K})} c_\lambda \lambda^n \varphi_\lambda(\boldsymbol{x}_0) + \int_{[-\pi,\pi]_{\text{per}}} e^{in\theta} \phi_{\theta,g}(\boldsymbol{x}_0) \, d\theta.$$
(2.5)

One can often characterize a dynamical system in terms of these decompositions. For example, suppose *F* is measure-preserving and bijective, and ω is a probability measure. In that case, the dynamical system is: (1) ergodic if and only if $\lambda = 1$ is a simple eigenvalue of \mathcal{K} , (2) weakly mixing if and only if $\lambda = 1$ is a simple eigenvalue of \mathcal{K} and there are no other eigenvalues, and (3) mixing if $\lambda = 1$ is a simple eigenvalue of \mathcal{K} and \mathcal{K} has absolutely continuous spectrum on span $\{1\}^{\perp}$ [59]. Different spectral types also have interpretations in the context of fluid mechanics [90], and weakly anomalous transport where the Koopman operator has singular continuous spectra [142].

2.2 | Spectra and pseudospectra

Since the Koopman operators of interest in this paper are not always normal, the spectrum of \mathcal{K} can be sensitive to small perturbations [129]. We care about pseudospectra as they are a way to determine which regions of the computed spectra are accurate and trustworthy. Moreover, if the Koopman operator is nonnormal, the system's transient behavior can differ greatly from the behavior at large times. Pseudospectra can be used to detect and quantify transients that are not captured by the spectrum [129, Section IV] [131]. For a finite matrix A and $\epsilon > 0$, the ϵ -pseudospectrum of A is defined as $\sigma_{\epsilon}(A) = \{\lambda \in \mathbb{C} : ||(A - \lambda I)^{-1}|| \ge 1/\epsilon\} = \bigcup_{\|B\| \le \epsilon} \sigma(A + B)$, where $\sigma(A + B)$ is the set of eigenvalues of A + B.¹⁰ Therefore, the ϵ -pseudospectra of A are regions in the complex plane enclosing eigenvalues of A that tell us how far an ϵ sized perturbation can perturb an eigenvalue. The ϵ -pseudospectra of a Koopman operator must be defined with some care because \mathcal{K} may be an unbounded operator [117]. We define the ϵ -pseudospectra of \mathcal{K} to be the following [105, Prop. 4.15]:¹¹

$$\sigma_{\epsilon}(\mathcal{K}) := \operatorname{cl}(\{\lambda \in \mathbb{C} : \|(\mathcal{K} - \lambda)^{-1}\| > 1/\epsilon\}) = \operatorname{cl}(\bigcup_{\|\mathcal{B}\| < \epsilon} \sigma(\mathcal{K} + \mathcal{B})),$$

where cl is the closure of the set.

11

⁹ To be precise, $\phi_{\theta,g} d\theta$ is the absolutely continuous component of $d\mathcal{E}(\theta)g$ and $\rho_g(\theta) = \langle \phi_{\theta,g}, g \rangle$.

¹⁰ Some authors use a strict inequality in the definition of ϵ -pseudospectra.

¹¹ While [105, Prop. 4.15] considers bounded operators, it can be adjusted to cover unbounded operators [129, Thm. 4.3].



FIGURE 1 The ϵ -pseudospectra for the nonlinear pendulum and $\epsilon = 0.25$ (shaded region) computed using Algorithm 3 with discretization sizes N_K . Discretization sizes correspond to a hyperbolic cross approximation (see Section 4.3.1). The computed ϵ -pseudospectra converge as $N_K \to \infty$. The unit circle (red line) is shown with the EDMD eigenvalues (magenta dots), many of which are spurious. Our algorithm called ResDMD removes spurious eigenvalues by computing pseudospectra (see Section 4).

Pseudospectra allow us to detect so-called *spectral pollution*, which are spurious eigenvalues that are caused by discretization and have no relation with the underlying Koopman operator. Spectral pollution can cluster at points not in the spectrum of \mathcal{K} , even when \mathcal{K} is a normal operator, as well as persist as the discretization size increases [83]. For example, consider the dynamical system of the nonlinear pendulum. Let $\mathbf{x} = (x_1, x_2) = (\theta, \dot{\theta})$ be the state variable governed by the following equations of motion:

$$\dot{x}_1 = x_2, \quad \dot{x}_2 = -\sin(x_1), \quad \text{with} \quad \Omega = [-\pi, \pi]_{\text{per}} \times \mathbb{R},$$

$$(2.6)$$

where ω is the standard Lebesgue measure. We consider the corresponding discrete-time dynamical system by sampling with a time-step $\Delta_t = 0.5$. The system is Hamiltonian, and hence the Koopman operator is unitary. It follows that $\sigma_{\epsilon}(\mathcal{K}) = \{z \in \mathbb{C} : \text{dist}(z, \mathbb{T}) \leq \epsilon\}$. Figure 1 shows the pseudospectrum computed using Algorithm 3 for $\epsilon = 0.25$ (see Section 4.3.1). The algorithm uses a discretization size of N_K to compute a set guaranteed to be inside the ϵ -pseudospectrum (i.e., no spectral pollution) that also converges to the ϵ -pseudospectrum as $N_K \to \infty$. We also show the corresponding EDMD eigenvalues. Some of these EDMD eigenvalues are reliable, but most are not, demonstrating severe spectral pollution. Note that this spectral pollution has nothing to do with any stability issues but instead is due to the discretization of the infinite-dimensional operator \mathcal{K} by a finite matrix. Using the ϵ -pseudospectrum for different ϵ , we can detect which eigenvalues are reliable (see Algorithm 2).

3 | COMPUTING SPECTRAL MEASURES FROM TRAJECTORY DATA

This section assumes that the Koopman operator \mathcal{K} is associated with a measure-preserving dynamical system and hence is an isometry. We aim to compute smoothed approximations of

 ν_g from trajectory data of the dynamical system, even in the presence of continuous spectra. More precisely, we want to select a smoothing parameter $\epsilon > 0$ and then compute a smooth periodic function ν_g^{ϵ} such that for any continuous periodic function ϕ we have [14, Ch. 1]

$$\int_{[-\pi,\pi]_{\text{per}}} \phi(\theta) \nu_g^{\epsilon}(\theta) \, d\theta \approx \int_{[-\pi,\pi]_{\text{per}}} \phi(\theta) \, d\nu_g(\theta).$$
(3.1)

Moreover, we want our smoothed approximation ν_g^{ϵ} to converge weakly to $\nu_g \epsilon \rightarrow 0$, meaning that the error in (3.1) goes to zero.¹² We describe an algorithm that achieves this from trajectory data (see Algorithm 1). The algorithm deals with general spectral measures, including those with a singular continuous component. We also evaluate approximations of the density ρ_g pointwise and can evaluate the atomic part of ν_g , corresponding to eigenvalues and eigenfunctions of \mathcal{K} .

3.1 Computing autocorrelations from trajectory data

We want to collect trajectory data from (1.1) and use this data to recover as much information as possible about the spectral properties of \mathcal{K} . We assume that (1.1) is observed for M_2 timesteps, starting at M_1 initial conditions. It is helpful to view the trajectory data as an $M_1 \times M_2$ matrix

$$B_{\text{data}} = \begin{bmatrix} \mathbf{x}_{0}^{(1)} & \cdots & \mathbf{x}_{M_{2}-1}^{(1)} \\ \vdots & \ddots & \vdots \\ \mathbf{x}_{0}^{(M_{1})} & \cdots & \mathbf{x}_{M_{2}-1}^{(M_{1})} \end{bmatrix}.$$
 (3.2)

Each row of B_{data} corresponds to an observation of M_2 time steps of a single trajectory of the dynamical system. In particular, $\mathbf{x}_{i+1}^{(j)} = F(\mathbf{x}_i^{(j)})$ for $0 \le i \le M_2 - 2$ and $1 \le j \le M_1$.

There are two main ways that trajectory data might be collected: (1) Experimentally-determined initial states of the dynamical system, where one must do the best one can with predetermined initial states, and (2) Algorithmically-determined initial states, where the algorithm can select the initial states and then record the dynamics. When recovering properties of \mathcal{K} , often it is best to have lots of initial states that explore the whole state-space Ω with a preference of having M_1 large and M_2 small. If M_2 is too large, then each trajectory could quickly get trapped in attracting states, making it difficult to recover the global properties of the dynamical system.

The type of trajectory data determines how we calculate autocorrelations. There are three main ways to compute autocorrelations corresponding to different types of initial conditions $\{\mathbf{x}_{0}^{(j)}\}_{i=1}^{M_{1}}$.

1. Initial conditions at quadrature nodes: Suppose that we are free to choose $\{\boldsymbol{x}_{0}^{(j)}\}_{j=1}^{M_{1}}$ so that they are an M_{1} -point quadrature rule with weights $\{w_{j}\}_{j=1}^{M_{1}}$. Integrals and inner products are then approximated with numerical integration by evaluating functions at the data points. This

¹² Since $\int_{[-\pi,\pi]_{\text{per}}} \phi(\theta) d\nu_g(\theta) = \langle \phi(-i \log(\mathcal{K}))g, g \rangle$, this convergence also approximates the functional calculus of \mathcal{K} . For example, if the dynamical system (1.1) corresponds to sampling a continuous-time dynamical system at discrete time steps, one could use (3.1) to recover spectral properties of Koopman operators that generate continuous-time dynamics.

means that autocorrelations that determine ν_g (see Section 2.1.3) can be approximated as

$$\langle g, \mathcal{K}^n g \rangle = \int_{\Omega} g(\mathbf{x}) \overline{[\mathcal{K}^n g](\mathbf{x})} \, d\omega(\mathbf{x}) \approx \sum_{j=1}^{M_1} w_j g(\mathbf{x}_0^{(j)}) \overline{[\mathcal{K}^n g](\mathbf{x}_0^{(j)})} = \sum_{j=1}^{M_1} w_j g(\mathbf{x}_0^{(j)}) \overline{g(\mathbf{x}_n^{(j)})}, \quad n \ge 0.$$

High-order quadrature rules can lead to fast rates of convergence. If $\mathcal{K}^n g$ is analytic in a neighborhood of Ω , then we can often select a quadrature rule that even converges exponentially as $M_1 \to \infty$ [130].

2. **Random initial conditions**: Suppose that $\{\mathbf{x}_{0}^{(j)}\}_{j=1}^{M_{1}}$ are drawn independently at random from a fixed probability distribution over Ω that is absolutely continuous with respect to ω with a sufficiently regular Radon–Nikodym derivative κ . Then, the autocorrelations of \mathcal{K} can be approximated using Monte Carlo integration, that is,

$$\langle g, \mathcal{K}^n g \rangle \approx \frac{1}{M_1} \sum_{j=1}^{M_1} \kappa(\mathbf{x}_0^{(j)}) g(\mathbf{x}_0^{(j)}) \overline{[\mathcal{K}^n g](\mathbf{x}_0^{(j)})} = \frac{1}{M_1} \sum_{j=1}^{M_1} \kappa(\mathbf{x}_0^{(j)}) g(\mathbf{x}_0^{(j)}) \overline{g(\mathbf{x}_n^{(j)})}.$$

This typically converges at a rate of $\mathcal{O}(M_1^{-1/2})$ [26], but is a practical approach if the state-space dimension is large. One could also consider quasi-Monte Carlo integration, which can achieve a faster rate of $\mathcal{O}(M_1^{-1})$ (up to logarithmic factors) under suitable conditions [26].

3. A single fixed initial condition: Intuitively, a system is ergodic if any trajectory eventually visits all parts of the state space. Formally, this means that F in (1.1) is measure-preserving, ω is a probability measure, and if A is a Borel measurable subset of Ω with $\{x : F(x) \in A\} \subset A$, then $\omega(A) = 0$ or $\omega(A) = 1$. If a dynamical system is ergodic, then Birkhoff's ergodic theorem [15] implies that

$$\langle g, \mathcal{K}^n g \rangle = \lim_{M_2 \to \infty} \frac{1}{M_2 - n} \sum_{j=0}^{M_2 - n-1} g(\boldsymbol{x}_j) \overline{[\mathcal{K}^n g](\boldsymbol{x}_j)} = \lim_{M_2 \to \infty} \frac{1}{M_2 - n} \sum_{j=0}^{M_2 - n-1} g(\boldsymbol{x}_j) \overline{g(\boldsymbol{x}_{j+n})},$$
(3.3)

for almost any initial condition \boldsymbol{x}_0 .¹³ However, the sampling scheme in (3.3) is restricted to ergodic dynamical systems and often requires very long trajectories (see [67] for convergence rates).

If one is entirely free to select the initial conditions of the trajectory data, and d is not too large, then we recommend picking them based on a high-order quadrature rule (see Sections 4.1.3 and 4.3.2). We can use sparse grids when the state-space dimension d is moderately large. When d is large, we can use a kernelized approach (see Section 6). One may also want to study the dynamics near attractors, where setting up explicit quadrature nodes for initial conditions can be tricky. In this case, small M_1 and large M_2 can be ideal with ergodic sampling [80, Section 7.2]. We have no control over the initial conditions for experimental data, and we may need to combine the above methods of computing autocorrelations. Often it is assumed that the initial conditions correspond to random initial conditions or that we have access to long trajectories. We deal with two examples of real-world data sets in Section 6.

¹³ Here we mean ω -almost any initial condition. In practice, it often holds in the Lebesgue sense for almost all initial conditions, and then one can randomly pick \mathbf{x}_0 from a suitable distribution over Ω .

3.2 | Recovering the spectral measure from autocorrelations

We now suppose that one has already computed the autocorrelations $\langle g, \mathcal{K}^n g \rangle$ for $0 \le n \le N$ and would like to recover a smoothed approximation of ν_g . Since the Fourier coefficients of ν_g are given by autocorrelations (see (2.2)), the task is similar to Fourier recovery [1, 57]. We are particularly interested in approaches with good convergence properties as $N \to \infty$, as this reduces the number of computed autocorrelations and the sample size $M = M_1 M_2$ required for good recovery of the spectral measure.

Motivated by the classical task of recovering a continuous function by its partial Fourier series [48], we start by considering the "windowing trick" from sampling theory. That is, we define a smoothed approximation to v_g as a measure with density

$$\nu_{g,N}(\theta) = \sum_{n=-N}^{N} \varphi\left(\frac{n}{N}\right) \hat{\nu_g}(n) e^{in\theta} = \frac{1}{2\pi} \sum_{n=-N}^{-1} \varphi\left(\frac{n}{N}\right) \overline{\langle g, \mathcal{K}^{-n}g \rangle} e^{in\theta} + \frac{1}{2\pi} \sum_{n=0}^{N} \varphi\left(\frac{n}{N}\right) \langle g, \mathcal{K}^{n}g \rangle e^{in\theta}.$$
(3.4)

The function $\varphi : [-1, 1] \to \mathbb{R}$ is often called a filter function [63, 127]. The idea of φ is that $\varphi(x)$ is close to 1 when x is close to 0, and φ tapers to 0 near $x = \pm 1$. By carefully tapering φ , the partial sum in (3.4) converges to ν_g as $N \to \infty$ (see Section 3.4). For fast pointwise or weak convergence of $\nu_{g,N}$ to ν_g , it is desirable for φ to be an even function that smoothly tapers from 1 to 0.

Using the definition of $\hat{\nu}_{g}(n)$, we have

$$\nu_{g,N}(\theta_0) = \sum_{n=-N}^{N} \varphi\left(\frac{n}{N}\right) \frac{1}{2\pi} \int_{[-\pi,\pi]_{\text{per}}} e^{-in\theta_0} d\nu_g(\theta) e^{in\theta_0} = \int_{[-\pi,\pi]_{\text{per}}} \frac{1}{2\pi} \sum_{n=-N}^{N} \varphi\left(\frac{n}{N}\right) e^{in(\theta_0-\theta)} d\nu_g(\theta),$$
(3.5)

which represents our smoothed approximation as a convolution with a smoothing kernel. One of the simplest filters is the hat function $\varphi_{hat}(x) = 1 - |x|$, for which (3.4) corresponds to the classical Cesàro summation of Fourier series and (3.5) is the convolution of ν_g with the Fejèr kernel, that is, $F_N(\theta) = \sum_{n=-N}^N (1 - |n|/N)e^{in\theta}$. From this viewpoint, the fact that $\nu_{g,N}$ weakly converges to ν_g is immediate from Parseval's formula [71, Section 1.7]. In particular, for any Lipschitz continuous test function ϕ with Lipschitz constant *L*,

$$\left|\phi(\theta) - \frac{1}{2\pi} [\phi * F_N](\theta)\right| \le \frac{1}{2\pi} \int_{[-\pi,\pi]_{\text{per}}} \underbrace{|\phi(\theta - s) - \phi(\theta)|}_{\le L|s|} F_N(s) \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(s)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(Ns)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(Ns)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(Ns)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(Ns)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(Ns)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s \frac{1 - \cos(Ns)}{1 - \cos(Ns)} \, ds \le \frac{L}{N\pi} \int_0^{\pi} s$$

where we use the fact that F_N is nonnegative and integrates to one. Using Fubini's theorem, we can write

$$\int_{[-\pi,\pi]_{\text{per}}} \phi(\theta) \nu_{g,N}(\theta) \, d\theta = \frac{1}{2\pi} \int_{[-\pi,\pi]_{\text{per}}} [\phi * F_N](\theta) \, d\nu_g(\theta),$$

where * denotes convolution. Since we can bound $s^2/(1 - \cos(s))$ from above by $\pi^2/2$ for $s \in [0, \pi]$, we have

$$\left| \int_{[-\pi,\pi]_{\text{per}}} \phi(\theta) \, d\nu_g(\theta) - \int_{[-\pi,\pi]_{\text{per}}} \phi(\theta) \nu_{g,N}(\theta) \, d\theta \right| \le \frac{\pi L}{2N} \int_0^\pi \frac{1 - \cos(Ns)}{s} \, ds \tag{3.6}$$

$$= \frac{\pi L}{2N} \int_0^{\pi N} \frac{1 - \cos(x)}{x} \, dx = \mathcal{O}(N^{-1} \log(N)), \tag{3.7}$$

where the last equality follows from a change of variables. Therefore, there is slow weak convergence of $\nu_{g,N}$ to ν_g as $N \to \infty$, providing us with an algorithm for ensuring (3.1) with $\epsilon = 1/N$. Algorithm 1 summarizes our computational framework for recovering a smoothed version of ν_g from autocorrelations of the trajectory data.

Other filter functions can provide a faster rate of convergence than $\varphi_{hat}(x) = 1 - |x|$, including the cosine and fourth-order filters [57, 134]:

$$\varphi_{\cos}(x) = \frac{1}{2}(1 + \cos(\pi x)), \qquad \varphi_{four}(x) = 1 - x^4(-20|x|^3 + 70x^2 - 84|x| + 35).$$

For the recovery of measures, we find that a particularly good choice is

$$\varphi_{\text{bump}}(x) = \exp\left(-\frac{2}{1-|x|}\exp\left(-\frac{c}{x^4}\right)\right), \quad c \approx 0.109550455106347,$$
 (3.8)

where the value of *c* is selected so that $\varphi_{\text{bump}}(1/2) = 1/2$. This filter can lead to arbitrary high orders of convergence with errors between $\nu_{g,N}$ and ν_g that go to zero faster than any polynomial in N^{-1} (see Section 3.4). A further useful property is that $\nu_{g,N}$ localizes any singular behavior of ν_g (see Section 5.1.3).¹⁴

To demonstrate the application of various filters for the recovery of spectral measures from autocorrelations, we consider the pedagogical dynamical system given by the shift operator. The observed orders of convergence of Algorithm 1 are predicted by our theory (see Section 3.3).

Example 3.1 (Shift operator). Consider the shift operator with state-space $\Omega = \mathbb{Z}$ (and counting measure ω) given by

$$x_{n+1} = F(x_n), \qquad F(x) = x + 1.$$

We seek to compute the spectral measure ν_g with respect to $g \in L^2(\mathbb{Z}, \omega) = \ell^2(\mathbb{Z})$, where $\ell^2(\mathbb{Z})$ is the space of square summable doubly infinite vectors. This example is a building block of many dynamical systems, such as Bernoulli shifts, with so-called Lebesgue spectrum [5, Chapter 2]. We consider the observable $g(k) = C \sin(k)/k$, where $C \approx 0.564189583547756$ is a normalization constant so that ||g|| = 1. For this example, ν_g is absolutely continuous, but ρ_g has discontinuities at $\theta = \pm 1$. Figure 2 shows the weak convergence (left) and pointwise convergence (right) for various filters.

3.3 | High-order kernels for spectral measure recovery

To develop convergence theory for Algorithm 1, we view filtering in the broader context of convolution with kernels. This provides a unified framework with Section 5 and simplifies the arguments when analyzing the recovery of measures (that may be highly singular), as opposed to the recovery of functions.

¹⁴ This because the kernel associated with φ_{bump} (see Proposition 3.1) is highly localized due to the smoothness of φ_{bump} .



FIGURE 2 Relative errors between $\nu_{g,N}$ and ν_g for the shift operator computed with filters φ_{hat} (blue), φ_{cos} (red), φ_{four} (yellow), and φ_{bump} (purple). Left: Relative error between $\nu_{g,N}$ to ν_g in the sense of (3.7), illustrating weak convergence in (3.1) for the test function $\phi(\theta) = \cos(5\theta)/(2 + \cos(\theta))$. Right: Relative error between $\nu_{g,N}$ to ρ_g at $\theta = 0$, illustrating pointwise convergence. The errors are calculated using an exact analytical expression for ν_g .

Instead of viewing $\nu_{g,N}$ as constructed by "windowing" the Fourier series of ν_g in (3.4) by a filter, we form an approximation to ν_g by convolution. That is, we define

$$u_g^{\epsilon}(\theta_0) = \int_{[-\pi,\pi]_{\mathrm{per}}} K_{\epsilon}(\theta_0 - \theta) d\nu_g(\theta),$$

where K_{ϵ} are a family of integrable functions { $K_{\epsilon} : 0 < \epsilon \leq 1$ } satisfying certain properties (see Definition 3.1) so that ν_g^{ϵ} converges to ν_g in some sense. We saw in (3.5) that $\nu_g^{\epsilon} = \nu_{g,N}$ when $K_{\epsilon}(\theta) = \frac{1}{2\pi} \sum_{n=-N}^{N} \varphi(\frac{n}{N}) e^{in\theta}$ and $N = \lfloor \epsilon^{-1} \rfloor$, say. The most famous example of K_{ϵ} is the Poisson kernel for the unit disc given by [71, p.16]

$$K_{\varepsilon}(\theta) = \frac{1}{2\pi} \frac{(1+\varepsilon)^2 - 1}{1 + (1+\varepsilon)^2 - 2(1+\varepsilon)\cos(\theta)},$$
(3.9)

in polar coordinates with $r = (1 + \epsilon)^{-1}$. The Poisson kernel is a first-order kernel because, up to a logarithmic factor, it leads to a first-order algebraic convergence rate of ν_g^{ϵ} to ν_g . We now give the following general definition of an *m*th order kernel and justify their name by showing that they lead to an *m*th order rate of convergence of ν_g^{ϵ} to ν_g in a weak and pointwise sense (see Section 3.4).

Definition 3.1 (*m*th order periodic kernel). Let $\{K_{\epsilon} : 0 < \epsilon \leq 1\}$ be a family of integrable functions on the periodic interval $[-\pi, \pi]_{per}$. We say that $\{K_{\epsilon}\}$ is an *m*th order kernel for $[-\pi, \pi]_{per}$ if

- (i) (Normalized) $\int_{[-\pi,\pi]_{\text{per}}} K_{\epsilon}(\theta) d\theta = 1.$
- (ii) (Approximately vanishing moments) There exists a constant C_K such that¹⁵

$$\left| \int_{[-\pi,\pi]_{\text{per}}} \theta^n K_{\epsilon}(\theta) \, d\theta \right| \le C_K \epsilon^m \log(\epsilon^{-1}), \quad \text{for any integer } 1 \le n \le m-1.$$
(3.10)

¹⁵ One may wonder if insisting on exactly vanishing moments or removing the logarithmic term improves convergence rates as $\varepsilon \downarrow 0$. This is not the case (see the discussion in Section 3.4).

(iii) (Decay away from 0) For any $\theta \in [-\pi, \pi]$ and $0 < \epsilon \leq 1$,

$$|K_{\varepsilon}(\theta)| \le \frac{C_K \varepsilon^m}{(\varepsilon + |\theta|)^{m+1}}.$$
(3.11)

The conditions in Definition 3.1 are primarily technical assumptions that allow one to prove appropriate convergence rates of v_g^e to v_g . For pointwise convergence, property (iii) is required to apply a local cut-off argument away from singular parts of the measure. Properties (i) and (ii) are used to show that terms vanish in a local Taylor series expansion of the Radon-Nikodym derivative, and the remainder is bounded by (iii). For weak convergence, we apply similar arguments to the test function by Fubini's theorem.

The properties of an *m*th order kernel can be translated to properties of a filter.

Proposition 3.1. Let $m \in \mathbb{N}$ and suppose that φ is an even continuous function that is compactly supported on [-1, 1] such that (a) $\varphi \in C^{m-1}([-1, 1])$, (b) $\varphi(0) = 1$ and $\varphi^{(n)}(0) = 0$ for any integer $1 \le n \le m-1$, (c) $\varphi^{(n)}(1) = 0$ for any integer $0 \le n \le m-1$, and (d) $\varphi|_{[0,1]} \in C^{m+1}([0,1])$. Then,

$$K_{\epsilon}(\theta) = \frac{1}{2\pi} \sum_{n=-N}^{N} \varphi\left(\frac{n}{N}\right) e^{in\theta}, \qquad N = \lfloor \epsilon^{-1} \rfloor$$
(3.12)

is an mth order kernel for $[-\pi, \pi]_{per}$.

Proof. We need to verify that K_{ϵ} in (3.12) satisfies the three properties in Definition 3.1. To see that

 K_{ε} is normalized, we integrate (3.12) term-by-term and note that $\varphi(0) = 1$. For property (iii), we define $\Phi(k) = \int_{-1}^{1} \varphi(x) e^{ikx} dx$ and note that by the Poisson summation formula, $K_{\epsilon}(\theta) = \frac{N}{2\pi} \sum_{j=-\infty}^{\infty} \Phi(N(\theta + 2\pi j))$. Since φ is an even function, $\Phi(k) =$ $2\int_0^1 \varphi(x) \cos(kx) dx$, and since φ satisfies (b), (c) and (d), we find that for $k \neq 0$

$$|\Phi(k)| = \frac{2}{k^m} \begin{cases} \left| \int_0^1 \varphi^m(x) \cos(kx) \, dx \right|, & \text{if } m \text{ is even,} \\ \left| \int_0^1 \varphi^m(x) \sin(kx) \, dx \right|, & \text{if } m \text{ is odd.} \end{cases}$$

Finally, one can use property (d) to perform integration-by-parts (now with possibly nonvanishing endpoint conditions) to deduce that $|\Phi(k)| \leq (1 + |k|)^{-(m+1)}$. Therefore, we find that

$$|K_{\varepsilon}(\theta)| \lesssim \sum_{j=-\infty}^{\infty} \frac{N}{(1+|N(\theta+2\pi j)|)^{m+1}} \lesssim \frac{N}{(1+N|\theta|)^{m+1}}.$$

This bound implies the decay condition in (3.11).

For property (ii), note that for any $N \ge |n|$, the following analytical expression holds:

$$\int_{[-\pi,\pi]_{\rm per}} K_{\varepsilon}(-\theta) e^{in\theta} \, d\theta - 1 = \varphi\left(\frac{n}{N}\right) - \varphi(0), \qquad N = \lfloor \epsilon^{-1} \rfloor.$$

19

 \square

As a consequence of Taylor's theorem and $\varphi^{(j)}(0) = 0$,

$$\left| \int_{[-\pi,\pi]_{\text{per}}} K_{\epsilon}(-\theta) e^{in\theta} \, d\theta - 1 \right| \leq \frac{n^m \|\varphi\|_{[0,1]}^{(m)}\|_{L^{\infty}}}{m! N^m} = \mathcal{O}(\epsilon^m) \quad \text{as } \epsilon \to 0.$$

The result now follows from Lemma A.1.

Therefore, it can be verified that: φ_{hat} , φ_{cos} , and φ_{four} induce first-order, second-order and fourth-order kernels in (3.12), respectively. Similarly, φ_{bump} induces a kernel that is *m*th order for any $m \in \mathbb{N}$. For example, up to a logarithmic factor, the rate of convergence between ν_g^{ϵ} (resp. $\nu_{g,N}$) and ν_g for φ_{four} is $\mathcal{O}(\epsilon^4)$ as $\epsilon \to 0$ (resp. $\mathcal{O}(N^{-4})$ as $N \to \infty$) in a weak and pointwise sense (see Section 3.4).

Readers who are experts on filters may be surprised by condition (d) in Proposition 3.1. This additional condition, which holds for all common choices of filters, is theoretically needed to obtain the optimal convergence rates between ν_g^{ϵ} and ν_g . It means that our convergence theory is one algebraic order better than the convergence theorems derived in the literature [57, Theorem 3.3].

3.4 | Convergence results

We now provide convergence theorems for recovering the spectral measure of \mathcal{K} . A reader concerned with only the practical aspects of our algorithms can safely skip over this section while appreciating that the convergence guarantees are rigorous.

3.4.1 | Pointwise convergence

For a point $\theta_0 \in [-\pi, \pi]$, the value of the approximate spectral measure $\nu_g^{\epsilon}(\theta_0)$ converges to the Radon–Nikodym derivative $\rho_g(\theta_0)$, provided that ν_g is absolutely continuous in an interval containing θ_0 (without this separation condition it still converges for almost every θ_0). The convergence rate depends on the smoothness of ρ_g in a small interval *I* containing θ_0 . In particular, we write $\rho_g \in C^{n,\alpha}(I)$ if ρ_g is *n*-times continuously differentiable on *I* and the *n*th derivative is Hölder continuous with parameter $0 \le \alpha < 1$. For $h_1 \in C^{0,\alpha}(I)$ and $h_2 \in C^{k,\alpha}(I)$ we define the seminorm and norm, respectively, as

$$|h_1|_{\mathcal{C}^{0,\alpha}(I)} = \sup_{x \neq y \in I} \frac{|h_1(x) - h_1(y)|}{|x - y|^{\alpha}}, \quad ||h_2||_{\mathcal{C}^{k,\alpha}(I)} = |h_2^{(k)}|_{\mathcal{C}^{0,\alpha}(I)} + \max_{0 \leq j \leq k} ||h_2^{(j)}||_{\infty,I}.$$

We state the following pointwise convergence theorem for general complex-valued measures ν as we apply it to measures corresponding to test functions to prove Theorem 3.2. The choice $\nu = \nu_g$ with $\|\nu_g\| = 1$ in Theorem 3.1 gives pointwise convergence of spectral measures.

Theorem 3.1 (Pointwise convergence). Let $\{K_{\epsilon}\}$ be an *m*th order kernel for $[-\pi, \pi]_{per}$ and let ν be a complex-valued measure on $[-\pi, \pi]_{per}$ with finite total variation $\|\nu\|$. Suppose that for some $\theta_0 \in [-\pi, \pi]$ and $\eta \in (0, \pi)$, ν is absolutely continuous on $I = (\theta_0 - \eta, \theta_0 + \eta)$ with Radon–Nikodym derivative $\rho \in C^{n,\alpha}(I)$ ($\alpha \in [0, 1)$). Then the following hold for any $0 \le \epsilon < 1$:

(i) If $n + \alpha < m$, then

$$\left|\rho(\theta_{0}) - \int_{[-\pi,\pi]_{\text{per}}} K_{\epsilon}(\theta_{0} - \theta) \, d\nu(\theta) \right| \lesssim C_{K} \left(\|\nu\| + \|\rho\|_{C^{n,\alpha}(I)} \right) \left(\epsilon^{n+\alpha} + \frac{\epsilon^{m}}{(\epsilon+\eta)^{m+1}} \right) (1 + \eta^{-n-\alpha}).$$

$$(3.13)$$

(ii) If $n + \alpha \ge m$, then

$$\left|\rho(\theta_{0}) - \int_{[-\pi,\pi]_{\text{per}}} K_{\varepsilon}(\theta_{0} - \theta) \, d\nu(\theta) \right| \lesssim C_{K} \left(\|\nu\| + \|\rho\|_{C^{m}(I)} \right) \left(\epsilon^{m} \log(\epsilon^{-1}) + \frac{\epsilon^{m}}{(\epsilon + \eta)^{m+1}} \right) (1 + \eta^{-m}). \tag{3.14}$$

Here, " \leq " *means that the inequality holds up to a constant that only depends on n* + α *and m.*

Proof. By periodicity, we can assume without loss of generality that $\theta_0 = 0$. First, we decompose ρ into two parts $\rho = \rho_1 + \rho_2$, where $\rho_1 \in C^{n,\alpha}(I)$ is compactly supported on *I* and ρ_2 vanishes on $(-\eta/2, +\eta/2)$. Using (3.11), we have

$$\left|\rho(0) - \int_{[-\pi,\pi]_{\text{per}}} K_{\epsilon}(-\theta) d\nu(\theta)\right| \leq \left|\rho_{1}(0) - \int_{[-\pi,\pi]_{\text{per}}} K_{\epsilon}(-\theta)\rho_{1}(\theta) d\theta\right| + \int_{|\theta| > \eta/2} \frac{C_{K} \epsilon^{m} d|\nu^{r}|(\theta)}{(\epsilon + \eta/2)^{m+1}},$$
(3.15)

where $d\nu^{r}(\theta) = d\nu(\theta) - \rho_{1}(\theta) d\theta$. The second term on the right-hand side of (3.15) is bounded by $\hat{C}_{1}C_{K}(\|\nu\| + \|\rho_{1}\|_{L^{\infty}(I)})\epsilon^{m}(\epsilon + \eta)^{-(m+1)}$ for some constant \hat{C}_{1} independent of all parameters. To bound the first term, we expand ρ_{1} using Taylor's theorem:

$$\rho_1(\theta) = \sum_{j=0}^{k-1} \frac{\rho_1^{(j)}(0)}{j!} \theta^j + \frac{\rho_1^{(k)}(\xi_\theta)}{k!} \theta^k, \qquad k = \min(n, m), \tag{3.16}$$

where $|\xi_{\theta}| \leq |\theta|$. We now consider the two cases of the theorem separately.

Case i (**n** + α < **m**). In this case, k = n and we can select ρ_1 so that,

$$\|\rho_1\|_{C^{n,\alpha}(I)} \le C(n,\alpha) \|\rho\|_{C^{n,\alpha}(I)} (1+\eta^{-n-\alpha}), \quad \|\rho_1\|_{L^{\infty}(I)} \le C(n,\alpha) \|\rho\|_{C^{n,\alpha}(I)}$$
(3.17)

for some universal constant $C(n, \alpha)$ that only depends on *n* and α . The existence of such a decomposition follows from standard arguments with cut-off functions. Using (3.16), part (ii) of Definition 3.1 and the first bound of (3.17), we obtain

$$\begin{aligned} \left| \rho_{1}(0) - \int_{[-\pi,\pi]_{\text{per}}} K_{\varepsilon}(-\theta) \rho_{1}(\theta) \, d\theta \right| &\leq \hat{C}_{2} C_{K} \|\rho\|_{C^{n,\alpha}(I)} \epsilon^{m} \log(\epsilon^{-1}) (1+\eta^{-n-\alpha}) \\ &+ \left| \int_{[-\pi,\pi]_{\text{per}}} K_{\varepsilon}(-\theta) \frac{\rho_{1}^{(n)}(\xi_{\theta}) - \rho_{1}^{(n)}(0)}{n!} \theta^{n} \, d\theta \right|, \end{aligned}$$

$$(3.18)$$

for some constant \hat{C}_2 independent of ϵ , η and ν (or ρ , ρ_1 , ρ_2). Note that we have added the factor of $\rho_1^{(n)}(0)$ into the integrand by a second application of part (ii) of Definition 3.1 and the fact

П

that n < m. The Hölder continuity of $\rho_1^{(n)}$ implies that $|\rho_1^{(n)}(\xi_\theta) - \rho_1^{(n)}(0)| \le C(n, \alpha) \|\rho\|_{C^{n,\alpha}(I)}(1 + \eta^{-n-\alpha})\theta^{\alpha}$. Using this bound in the integrand on the right-hand side of (3.18) and (3.11), we obtain

$$\left|\rho_1(0) - \int_{[-\pi,\pi]_{\text{per}}} K_{\varepsilon}(-\theta)\rho_1(\theta) \, d\theta \right| \leq \hat{C}_3 C_K \|\rho\|_{C^{n,\alpha}(I)} \left(\epsilon^m \log(\epsilon^{-1}) + \epsilon^{n+\alpha} \int_0^{\pi/\epsilon} \frac{\tau^{n+\alpha} d\tau}{(1+\tau)^{m+1}}\right) (1+\eta^{-n-\alpha}),$$

for some constant \hat{C}_3 independent of ϵ , η and ν (or ρ , ρ_1 , ρ_2). Since $m > n + \alpha$, the integral in brackets converges as $\epsilon \downarrow 0$, and the bound in (3.13) now follows.

Case ii ($\mathbf{n} + \boldsymbol{\alpha} \ge \mathbf{m}$). In this case, k = m and we can select ρ_1 such that

$$\|\rho_1\|_{\mathcal{C}^m(I)} \le C(m) \|\rho\|_{\mathcal{C}^m(I)} (1+\eta^{-m})$$

for some universal constant C(m) that only depends on m. Again, the existence of such a decomposition follows from standard arguments with cut-off functions. Using (3.16) and applying (3.10) to the powers θ_j for j < m and (3.11) to the θ^m term, we obtain

$$\left|\rho_1(0) - \int_{[-\pi,\pi]_{\text{per}}} K_{\epsilon}(-\theta)\rho_1(\theta) \, d\theta\right| \leq \hat{C}_2 C_K \|\rho\|_{C^m(I)} \left(\epsilon^m \log(\epsilon^{-1}) + \epsilon^m \int_0^{\pi/\epsilon} \frac{\tau^m d\tau}{(1+\tau)^{m+1}}\right) (1+\eta^{-m}),$$

for some constant \hat{C}_2 independent of ϵ , η and ν (or ρ , ρ_1 , ρ_2). The bound in (3.14) now follows.

The logarithmic term in (3.14) is due to the divergence of the integral $\int_0^{\pi/\epsilon} \frac{\tau^m d\tau}{(1+\tau)^{m+1}}$ as $\epsilon \downarrow 0$ and cannot, in general, be removed even with exactly vanishing moments replacing (3.10).

3.4.2 | Weak convergence

We now prove weak convergence in the sense of (3.1). Using Theorem 3.1 and a duality argument, we can also provide convergence rates for (3.1).

Theorem 3.2 (Weak convergence). Let $\{K_{\epsilon}\}$ be an *m*th order kernel for $[-\pi, \pi]_{per}$, $\phi \in C^{n,\alpha}([-\pi, \pi]_{per})$, and let ν_g be a spectral measure on the periodic interval $[-\pi, \pi]_{per}$. Then

$$\left| \int_{[-\pi,\pi]_{\text{per}}} \phi(\theta) \nu_g^{\epsilon}(\theta) \, d\theta - \int_{[-\pi,\pi]_{\text{per}}} \phi(\theta) \, d\nu_g(\theta) \right| \lesssim C_K \|\phi\|_{\mathcal{C}^{n,\alpha}([-\pi,\pi]_{\text{per}})} \left(\epsilon^{n+\alpha} + \epsilon^m \log(\epsilon^{-1})\right), \tag{3.19}$$

where ' \leq ' means that the inequality holds up to a constant that only depends on $n + \alpha$ and m.

Proof. Let $\tilde{K}_{\varepsilon}(\theta) = K_{\varepsilon}(-\theta)$, then it is easily seen that $\{\tilde{K}_{\varepsilon}\}$ is an *m*th order kernel for $[-\pi, \pi]_{\text{per}}$. Fubini's theorem allows us to exchange the order of integration to see that

$$\int_{[-\pi,\pi]_{\text{per}}} \phi(\theta) \nu_g^{\epsilon}(\theta) \, d\theta = \int_{[-\pi,\pi]_{\text{per}}} \phi(\theta) [K_{\epsilon} * \nu_g](\theta) \, d\theta = \int_{[-\pi,\pi]_{\text{per}}} [\tilde{K}_{\epsilon} * \phi](\theta) \, d\nu_g(\theta).$$

We can now apply Theorem 3.1 to the absolutely continuous measure with Radon–Nikodym derivative ϕ and the kernel \tilde{K}_{ε} (e.g., with $\eta = \pi/2$) to see that

$$\left| \left[\tilde{K}_{\epsilon} * \phi \right](\theta) - \phi(\theta) \right| \le C_1 C_K \|\phi\|_{\mathcal{C}^{n,\alpha}([-\pi,\pi]_{\operatorname{ner}})} \left(\epsilon^{n+\alpha} + \epsilon^m \log(\epsilon^{-1}) \right),$$

for some constant C_1 depending on n, α and m. Since ν_g is a probability measure, (3.19) follows.

The high-order convergence in Theorem 3.2 does not require regularity assumptions on ν_g . Moreover, though not covered by the theorem, weak convergence still holds for any *m*th order kernel and continuous periodic function ϕ .

3.4.3 | Recovery of the atomic parts of the spectral measure

Finally, we consider the recovery of the atomic parts of spectral measures or, equivalently, $\sigma_p(\mathcal{K})$ - the set of eigenvalues of \mathcal{K} (see (2.4)). This convergence is achieved by *rescaling* the smoothed approximation $K_{\varepsilon} * \nu_g$. The following theorem means that Algorithm 1 converges to both the eigenvalues of \mathcal{K} and the continuous part of the spectrum of \mathcal{K} .

Theorem 3.3 (Recovery of atoms). Let $\{K_{\varepsilon}\}$ be an *m*th order kernel for $[-\pi, \pi]_{\text{per}}$ that satisfies $\limsup_{\varepsilon \downarrow 0} \frac{\varepsilon^{-1}}{|K_{\varepsilon}(0)|} < \infty$, and let ν_g be a spectral measure on $[-\pi, \pi]_{\text{per}}$. Then, for any $\theta_0 \in [-\pi, \pi]_{\text{per}}$,

$$\nu_g(\{\theta_0\}) = \lim_{\epsilon \downarrow 0} \frac{1}{K_{\epsilon}(0)} [K_{\epsilon} * \nu_g](\theta_0).$$
(3.20)

Proof. By periodicity, we may assume without loss of generality that $\theta_0 = 0$. Let $\nu'_g = \nu_g - \nu_g(\{0\})\delta_0$, then

$$\frac{1}{K_{\varepsilon}(0)}[K_{\varepsilon} * \nu_{g}](0) = \nu_{g}(\{0\}) + \frac{1}{K_{\varepsilon}(0)}[K_{\varepsilon} * \nu_{g}'](0).$$
(3.21)

Consider the function $K_{\epsilon}(-\theta)/K_{\epsilon}(0)$, which is uniformly bounded for sufficiently small ϵ using (3.11) and the assumption $\limsup_{\epsilon \downarrow 0} \frac{\epsilon^{-1}}{|K_{\epsilon}(0)|} < \infty$. Since $\lim_{\epsilon \downarrow 0} K_{\epsilon}(-\theta)/K_{\epsilon}(0) = 0$ for any $\theta \neq 0$ and $\nu'_{g}(\{0\}) = 0$,

$$\lim_{\varepsilon \downarrow 0} \frac{1}{K_{\varepsilon}(0)} [K_{\varepsilon} * \nu'_{g}](0) = \lim_{\varepsilon \downarrow 0} \int_{[-\pi,\pi]_{\text{per}}} \frac{K_{\varepsilon}(-\theta)}{K_{\varepsilon}(0)} d\nu'_{g} = 0,$$

where we used the dominated convergence theorem. The theorem now follows from (3.21).

The condition that $\limsup_{\epsilon \downarrow 0} \frac{\epsilon^{-1}}{|K_{\epsilon}(0)|} < \infty$ is a technical condition that is satisfied by all the kernels constructed in this paper. A condition such as this is required to recover the atomic part of ν_{g} , as it says that K_{ϵ} must become localized around 0 sufficiently quickly as $\epsilon \to 0$.

However, we end this section with the following warning to the reader about recovering atomic parts of spectral measures. As soon as ν_g has atoms (i.e., \mathcal{K} has eigenvalues and g is not orthogonal to all the eigenspaces), the map $\theta \mapsto \nu_g(\{\theta\})$ is discontinuous. One can prove that, in



FIGURE 3 The approximation $[K_{\epsilon} * \nu_g](\theta)/K_{\epsilon}(0)$ of the atomic parts of the spectral measure using φ_{hat} and $g(x) = C \cos(11x)/\sqrt{1.001 + \sin(x)}$, where $C \approx 0.150803789609385$ is a normalization constant so that ||g|| = 1. As $\epsilon \downarrow 0$, the curves converge pointwise to the map $\theta \mapsto \nu_g(\{\theta\})$. Left column: The approximation for the rotation in (3.22) with c = 0.7 and N = 100 (top-left) and N = 10000 (bottom-left) ($\epsilon = 1/N$). Right column: The approximation for the rotation in (3.22) with $c = 1/\sqrt{2} \approx 0.707106781186548$ and and N = 1000 (top-left) and N = 10000 (bottom-left) ($\epsilon = 1/N$).

general, separating the point spectrum from the rest of the spectrum, either in terms of spectral measures or spectral sets, is impossible for any algorithm. This holds even for simple classes of operators [30, 35], unless we know apriori that the spectrum is discrete in a region of interest [28, Section 7.3]. We use smoothing with convolution kernels to regularize the map $\theta \mapsto \nu_g(\{\theta\})$. Analogues of Theorem 3.3 were proven in [80] using the Christoffel–Darboux kernel (which should not be confused with a convolution kernel as we have defined it) and in [39] using reproducing kernel Hilbert spaces to regularize harmonic averaging techniques. An apparent advantage of the framework of convolution kernels is that we need minimal assumptions on our dynamical system and trajectory data, and can provide explicit convergence rates in Theorems 3.1 and 3.2.

Example 3.2 (Irrational circle rotation). Consider the rotation operator with state-space $\Omega = [-\pi, \pi]_{per}$ (and standard Lebesgue measure ω) given by

$$x_{n+1} = F(x_n), \qquad F(x) = x + 2c\pi.$$
 (3.22)

If c = p/q is rational, with p and q coprime, then the Koopman operator has pure point spectrum at $\theta = 0, 2\pi/q, ..., 2\pi(q-1)/q$. Otherwise, the Koopman operator has dense point spectrum. We consider the observable $g(x) = C \cos(11x)/\sqrt{1.001 + \sin(x)}$, where $C \approx 0.150803789609385$ is a normalization constant so that ||g|| = 1. This particular choice is so that g has a relatively slowly decaying Fourier series. Figure 3 shows $[K_{\epsilon} * \nu_g](\theta)/K_{\epsilon}(0)$ for the filter $\varphi_{hat}(x) = 1 - |x|$ and N = 100, 10000. On the left, we plot the results for c = 0.7 and on the right for $c = 1/\sqrt{2} \approx 0.707106781186548$. We see the highly discontinuous nature of $\theta_0 \mapsto \nu_g(\{\theta_0\})$, both in terms of θ_0 and the map F itself.

3.5 | Numerical examples

We now consider two numerical examples of Algorithm 1 and the convergence theory in Section 3.4.



FIGURE 4 Computed approximate spectral measures with respect to the function *g* in (3.23) using Algorithm 1 with the filter in (3.8) for the tent map. These are computed with (blue) and without (red) the filter in (3.8) for discretization sizes N = 100 (left) and N = 1000 (right). The function $\nu_{g,N}$ is highly oscillatory if no filter is used (see zoomed-in subplot).

3.5.1 | Tent map

The tent map with parameter 2 is the function $F : [0,1] \rightarrow [0,1]$ given by $F(x) = 2 \min\{x, 1-x\}$. It generates a chaotic system with discrete and continuous spectra. We consider $\Omega = [0,1]$ with the usual Lebesgue measure. The corresponding Koopman operator \mathcal{K} is an isometry. However, it is not unitary since the function $\mathcal{K}g$ is symmetric about x = 1/2 for any function g, and hence \mathcal{K} is not onto. Furthermore, the decomposition in (2.4) reduces to an atomic part at $\theta = 0$ of size $(\int_0^1 g(x)dx)^2$ and an absolutely continuous part. The tent map thus demonstrates that our algorithm deals with mixed spectral types without a priori knowledge of the eigenvalues of \mathcal{K} .

To compute the inner products $\langle g, \mathcal{K}^n g \rangle$ for Algorithm 1, we sample the observable g on an equally spaced dyadic grid with equal weights.¹⁶ As an example, consider the arbitrary discontinuous function

$$g(\theta) = C|\theta - 1/3| + C\sin(20\theta) + \begin{cases} C, & \theta > 0.78, \\ 0, & \theta \le 0.78, \end{cases}$$
(3.23)

where $C \approx 1.035030525813683$ is a normalization constant. Applying Algorithm 1, we found convergence to the Radon–Nikodym derivative away from the singular part of the measure (the atom at zero) behaved as predicted by Theorem 3.1. To see the importance of the filter, we plot the reconstruction from the Fourier coefficients both with and without the filter (3.8) in Figure 4. As N increases, we see that the filter localizes the severe oscillations near the origin, and we gain convergence to the Radon–Nikodym derivative away from the eigenvalue at 0. This is not the case without the filter, where severe oscillations pollute the entire interval. Using the filter in (3.8), the atomic part approximated via Theorem 3.3 was correct to 4.4×10^{-4} for N = 1000 and converges in the limit $N \to \infty$, whereas convergence via (3.20) does not hold without the filter.

¹⁶ While this quadrature rule may seem suboptimal, it is selected because of the dyadic structure of the tent map and to avoid issues when computing the integrals $\langle g, \mathcal{K}^n g \rangle$ for large *n* due to the chaotic nature of *F*.



FIGURE 5 Maximum absolute error in computing $\hat{v}_g(n)$ for $-10 \le n \le 10$ with a sample size of M_1M_2 for quadrature (blue) and stabilized ergodic (red). Each g is normalized so that ||g|| = 1.

Finally, we consider the sample complexity M_1M_2 needed to recover the Fourier coefficients of the measures. We consider the maximum absolute error of the computed $\hat{\nu}_g(n)$ for $|n| \leq 10$ for our quadrature-based method and the ergodic sampling method from (3.3). For the ergodic method, $M_1 = 1$ and hence the sample complexity is the length of the trajectory used. However, for this example, the naive application of (3.3) for the ergodic method is severely unstable. Because of the binary nature of the tent map, each application of *F* loses a single bit of precision. Using floating-point arithmetic means that trajectories eventually stagnate at the fixed point 0. This error is not a random numerical error but a structured one. A good solution is to perturb each evaluation of *F* by a small random amount. The instability is absent for the tent map with parameters that avoid exact binary representations. Figure 5 shows the results for a range of different *g*. The ergodic method converges like $\mathcal{O}(M_2^{-1/2})$ as $M_2 \to \infty$, and the quadrature-based method converges with a faster rate in M_1 .

3.5.2 | Nonlinear pendulum

The nonlinear pendulum is a nonchaotic Hamiltonian system with continuous spectra and challenging Koopman operator theory. Here, we consider a corresponding discrete-time system by sampling (2.6) with a time-step of $\Delta_t = 1$. We collect M_1 data points on an equispaced tensor product grid corresponding to the periodic trapezoidal quadrature rule in the x_1 direction and a truncated¹⁷ trapezoidal quadrature rule in the x_2 direction. To simulate the collection of trajectory data, we compute trajectories starting at each initial condition using the ode45 command in MATLAB. We stress that we only use ode45 as a black-box integrator – all of our algorithms in this paper are purely data-driven.

We look at the following observable that involves nontrivial dynamics in each coordinate:

$$g(x_1, x_2) = C(1 + i\sin(x_1))(1 - \sqrt{2x_2})e^{-x_2^2/2},$$

where $C \approx 0.24466788518668$ is a normalization constant. Figure 6 shows high resolution approximations of the spectral measure ν_g for N = 100 using $M_1 = 50000$ and N = 1000 using $M_1 = 10^6$. The spectral measure is purely continuous (no atoms) away from $\theta = 0$, consistent with the general theory of integrable Hamiltonian systems with one degree of freedom [92]. Note that the constant function 1 is not in $L^2([0, 2\pi]_{\text{per}} \times \mathbb{R})$ and hence cannot be an eigenfunction.

¹⁷ We select a truncation to $x_2 \in [-L, L]$ so that the g and $\mathcal{K}^n g$ are negligible for $|x_2| > L$.



FIGURE 6 Computed spectral measure using Algorithm 1 with (3.8) for the nonlinear pendulum.

We confirmed this by using Theorem 3.3 for larger *N* and observing that the peak at $\theta = 0$ seen in Figure 6 does not grow as fast as $\propto N$. However, the spectral measure behaves singularly at $\theta = 0$.

4 | RESIDUAL DMD (ResDMD)

In this section, we consider dynamical systems of the form (1.1) that are not necessarily measurepreserving. Therefore, we cannot assume that \mathcal{K} is an isometry. Instead, all we assume about \mathcal{K} is that it is a closed and densely defined operator. We assume that we have access to a sequence of snapshots, that is, a trajectory data matrix (see (3.2)) with two columns:

$$B_{\text{data}} = \begin{bmatrix} \mathbf{x}^{(1)} & \mathbf{y}^{(1)} \\ \vdots & \vdots \\ \mathbf{x}^{(M)} & \mathbf{y}^{(M)} \end{bmatrix},$$
(4.1)

where $\mathbf{y}^{(j)} = F(\mathbf{x}^{(j)})$ for $1 \le j \le M = M_1$. An $M_1 \times M_2$ trajectory data matrix can be converted to the form of (4.1) by forgetting that the data comes from longer runs and reshaping the matrix.

Using (4.1), we develop a new algorithm, Residual DMD (ResDMD), that approximates the associated Koopman operator of the dynamics. Our approach allows for Koopman operators \mathcal{K} with no non-trivial finite-dimensional invariant subspace. The critical difference between ResDMD and other DMD algorithms (such as EDMD) is that we construct Galerkin approximations for not only \mathcal{K} , but also $\mathcal{K}^*\mathcal{K}$. This difference allows us to have rigorous convergence guarantees for ResDMD when recovering the spectral information of \mathcal{K} and computing spectra and pseudospectra. In particular, we avoid spectral pollution (see Figure 1).

4.1 | Extended DMD (EDMD) and a new matrix for computing residuals

Before discussing our ResDMD approach, we describe EDMD. EDMD constructs a matrix $K_{\text{EDMD}} \in \mathbb{C}^{N_K \times N_K}$ that approximates the action of \mathcal{K} from the snapshot data in (4.1). The original description of EDMD assumes that $\{\boldsymbol{x}^{(j)}\}_{j=1}^M \subset \Omega$ are drawn independently according to ω [138]. Here, we describe EDMD for arbitrary initial states and use $\{\boldsymbol{x}^{(j)}\}_{j=1}^M$ as quadrature nodes.

10970121, 0, Downloaded from https://onlinebitary.wiley.com/doi/10.1092/cpa.22125 by Cambridge University Library, Wiley Online Library on (27/07/2023). See the Terms and Conditions (https://onlinebitary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA arisels are governed by the applicable Creative Commons License

4.1.1 | EDMD viewed as a Galerkin method

Given a dictionary $\{\psi_1, \dots, \psi_{N_K}\} \subset \mathcal{D}(\mathcal{K})$ of observables, EDMD selects a matrix K_{EDMD} that approximates \mathcal{K} on the subspace $V_{N_K} = \text{span}\{\psi_1, \dots, \psi_{N_K}\}$, that is, $[\mathcal{K}\psi_j](\mathbf{x}) = \psi_j(F(\mathbf{x})) \approx \sum_{i=1}^{N_K} (K_{\text{EDMD}})_{ij}\psi_i(\mathbf{x})$ for $1 \leq j \leq N_K$. Define the vector-valued feature map $\Psi(\mathbf{x}) = [\psi_1(\mathbf{x}) \cdots \psi_{N_K}(\mathbf{x})] \in \mathbb{C}^{1 \times N_K}$. Then any $g \in V_{N_K}$ can be written as $g(\mathbf{x}) = \sum_{j=1}^{N_K} \psi_j(\mathbf{x})g_j = \Psi(\mathbf{x})g$ for some vector $\mathbf{g} \in \mathbb{C}^{N_K}$. It follows that

$$[\mathcal{K}g](\mathbf{x}) = \Psi(F(\mathbf{x}))\mathbf{g} = \Psi(\mathbf{x})(K_{\text{EDMD}}\mathbf{g}) + \underbrace{\left(\sum_{j=1}^{N_K} \psi_j(F(\mathbf{x}))g_j - \Psi(\mathbf{x})(K_{\text{EDMD}}\mathbf{g})\right)}_{r(\mathbf{g},\mathbf{x})}.$$

Typically, V_{N_K} is not an invariant subspace of \mathcal{K} so there is no choice of K_{EDMD} that makes $r(\mathbf{g}, \mathbf{x})$ zero for all $\mathbf{g} \in V_N$ and $\mathbf{x} \in \Omega$. Instead, it is natural to select K_{EDMD} as a solution of

$$\operatorname{argmin}_{B \in \mathbb{C}^{N_{K} \times N_{K}}} \left\{ \int_{\Omega} \max_{\boldsymbol{g} \in \mathbb{C}^{N_{K}}, \|\boldsymbol{g}\|=1} |r(\boldsymbol{g}, \boldsymbol{x})|^{2} d\omega(\boldsymbol{x}) = \int_{\Omega} \|\Psi(F(\boldsymbol{x})) - \Psi(\boldsymbol{x})B\|_{\ell^{2}}^{2} d\omega(\boldsymbol{x}) \right\}.$$
(4.2)

Here, $\|\cdot\|_{\ell^2}$ denotes the standard Euclidean norm of a vector.

In practice, one cannot directly evaluate the integral in (4.2). Instead, we approximate it via a quadrature rule with nodes $\{\mathbf{x}^{(j)}\}_{j=1}^{M}$ and weights $\{w_j\}_{j=1}^{M}$. The discretized version of (4.2) is the following weighted least-squares problem:

$$\operatorname{argmin}_{B \in \mathbb{C}^{N_K \times N_K}} \sum_{j=1}^M w_j \left\| \Psi(\boldsymbol{y}^{(j)}) - \Psi(\boldsymbol{x}^{(j)}) B \right\|_{\ell^2}^2.$$
(4.3)

A solution to (4.3) can be written down explicitly as $K_{EDMD} = (\Psi_X^* W \Psi_X)^{\dagger} (\Psi_X^* W \Psi_Y)$, where '†' denotes the pseudoinverse and $W = \text{diag}(w_1, \dots, w_M)$. Here, Ψ_X and Ψ_Y are the $M \times N_K$ matrices given by

$$\Psi_{X} = \begin{pmatrix} \Psi(\boldsymbol{x}^{(1)}) \\ \vdots \\ \Psi(\boldsymbol{x}^{(M)}) \end{pmatrix} \in \mathbb{C}^{M \times N_{K}}, \quad \Psi_{Y} = \begin{pmatrix} \Psi(\boldsymbol{y}^{(1)}) \\ \vdots \\ \Psi(\boldsymbol{y}^{(M)}) \end{pmatrix} \in \mathbb{C}^{M \times N_{K}}.$$
(4.4)

By reducing the size of the dictionary if necessary, we may assume without loss of generality that $\Psi_X^* W \Psi_X$ is invertible. Regularization through truncated singular value decompositions or a change of basis representation may also be considered. Since $\Psi_X^* W \Psi_X = \sum_{j=1}^M w_j \Psi(\mathbf{x}^{(j)})^* \Psi(\mathbf{x}^{(j)})$ and $\Psi_X^* W \Psi_Y = \sum_{j=1}^M w_j \Psi(\mathbf{x}^{(j)})^* \Psi(\mathbf{y}^{(j)})$, if the quadrature converges (see Section 4.1.3) then

$$\lim_{M \to \infty} [\Psi_X^* W \Psi_X]_{jk} = \langle \psi_k, \psi_j \rangle \quad \text{and} \quad \lim_{M \to \infty} [\Psi_X^* W \Psi_Y]_{jk} = \langle \mathcal{K} \psi_k, \psi_j \rangle,$$

where $\langle \cdot, \cdot \rangle$ is the inner product associated with $L^2(\Omega, \omega)$. Thus, EDMD can be viewed as a Galerkin method in the large data limit $M \to \infty$. Let $\mathcal{P}_{V_{N_K}}$ denote the orthogonal projection onto V_{N_K} . In the large data limit, K_{EDMD} approaches a matrix representation of $\mathcal{P}_{V_{N_K}} \mathcal{KP}^*_{V_{N_K}}$ and the EDMD eigenvalues approach the spectrum of $\mathcal{P}_{V_{N_K}} \mathcal{KP}^*_{V_{N_K}}$. Thus, approximating $\sigma(\mathcal{K})$ by the eigenvalues of K_{EDMD} is closely related to the so-called finite section method [16]. Since the finite

section method can suffer from spectral pollution (see Section 2.2), spectral pollution is also a concern for EDMD¹⁸ and it is important to have an independent way to measure the accuracy of the candidate eigenvalue-eigenvector pairs.

4.1.2 | Measuring the accuracy of candidate eigenvalue-eigenvector pairs

Suppose that we have a candidate eigenvalue-eigenvector pair (λ, g) of \mathcal{K} , where $\lambda \in \mathbb{C}$ and $g = \Psi \mathbf{g} \in V_{N_{\mathcal{K}}}$. One way to measure the accuracy of (λ, g) is by estimating the squared relative residual

$$\frac{\int_{\Omega} \left| \left[\mathcal{K}g \right](\boldsymbol{x}) - \lambda g(\boldsymbol{x}) \right|^{2} d\omega(\boldsymbol{x})}{\int_{\Omega} \left| g(\boldsymbol{x}) \right|^{2} d\omega(\boldsymbol{x})} = \frac{\langle (\mathcal{K} - \lambda)g, (\mathcal{K} - \lambda)g \rangle}{\langle g, g \rangle}$$

$$= \frac{\sum_{j,k=1}^{N_{K}} \overline{g_{j}}g_{k} \Big[\langle \mathcal{K}\psi_{k}, \mathcal{K}\psi_{j} \rangle - \lambda \langle \psi_{k}, \mathcal{K}\psi_{j} \rangle - \overline{\lambda} \langle \mathcal{K}\psi_{k}, \psi_{j} \rangle + |\lambda|^{2} \langle \psi_{k}, \psi_{j} \rangle \Big]}{\sum_{i,k=1}^{N_{K}} \overline{g_{j}}g_{k} \langle \psi_{k}, \psi_{j} \rangle}.$$
(4.5)

If \mathcal{K} is a normal operator, then the minimum of (4.5) over all normalized $g \in \mathcal{D}(\mathcal{K})$ is exactly the square distance of λ to the spectrum of \mathcal{K} ; otherwise, for nonnormal \mathcal{K} the residual can still provide a measure of accuracy (see Appendix B). One can also use the residual to bound the distance between g and the eigenspace associated with λ , assuming λ is a point in the discrete spectrum of \mathcal{K} [122, Chapter V].

We approximate the residual in (4.5) by

$$\operatorname{res}(\lambda,g)^{2} = \frac{\sum_{j,k=1}^{N_{K}} \overline{g_{j}} g_{k} \Big[(\Psi_{Y}^{*} W \Psi_{Y})_{jk} - \lambda (\Psi_{Y}^{*} W \Psi_{X})_{jk} - \overline{\lambda} (\Psi_{X}^{*} W \Psi_{Y})_{jk} + |\lambda|^{2} (\Psi_{X}^{*} W \Psi_{X})_{jk} \Big]}{\sum_{j,k=1}^{N_{K}} \overline{g_{j}} g_{k} (\Psi_{X}^{*} W \Psi_{X})_{jk}}$$
(4.6)

All the terms in this residual can be computed using the snapshot data. Note that, as well as the matrices found in EDMD, (4.6) has the *additional matrix* $\Psi_Y^*W\Psi_Y$. Moreover, under certain conditions, $\lim_{M\to\infty} \operatorname{res}(\lambda, g)^2 = \int_{\Omega} |[\mathcal{K}g](\mathbf{x}) - \lambda g(\mathbf{x})|^2 d\omega(\mathbf{x}) / \int_{\Omega} |g(\mathbf{x})|^2 d\omega(\mathbf{x})$ for any $g \in V_{N_K}$ (see Section 4.1.3). In particular, $\lim_{M\to\infty} [\Psi_Y^*W\Psi_Y]_{jk} = \langle \mathcal{K}\psi_k, \mathcal{K}\psi_j \rangle$ and $\Psi_Y^*W\Psi_Y$ formally corresponds to a Galerkin approximation of $\mathcal{K}^*\mathcal{K}$ as $M \to \infty$. In Appendix B, we show that the quantity $\operatorname{res}(\lambda, g)$ can be used to rigorously compute spectra and pseudospectra of \mathcal{K} .

4.1.3 | Convergence of matrices in the large data limit

Echoing Section 3.1, we focus on three situations where

$$\lim_{M \to \infty} [\Psi_X^* W \Psi_X]_{jk} = \langle \psi_k, \psi_j \rangle, \quad \lim_{M \to \infty} [\Psi_X^* W \Psi_Y]_{jk} = \langle \mathcal{K} \psi_k, \psi_j \rangle, \quad \lim_{M \to \infty} [\Psi_Y^* W \Psi_Y]_{jk} = \langle \mathcal{K} \psi_k, \mathcal{K} \psi_j \rangle. \tag{4.7}$$

For initial conditions at quadrature nodes, if the dictionary functions and *F* are sufficiently regular, then it is beneficial to select $\{\mathbf{x}^{(j)}\}_{j=1}^{M}$ as an *M*-point quadrature rule with weights $\{w_j\}_{j=1}^{M}$.

¹⁸ It is pointed out in [79] that if \mathcal{K} is bounded and the corresponding eigenvectors of a sequence of eigenvalues do not weakly converge to zero as the discretization size increases, then the EDMD eigenvalues have a convergent subsequence to an element of the spectrum. Unfortunately, one can prove that no algorithm can determine whether a sequence of eigenvectors converges weakly to zero.

This can lead to much faster convergence rates in (4.7) (see Section 4.3.2). For example, if Ω is unbounded, then we can use quadrature rules such as the trapezoidal rule (see Section 4.3.1), and if Ω is a simple bounded domain, then one can use Gaussian quadrature (see Section 4.3.2). When the state-space dimension *d* is moderately large, we can use sparse grids and a kernelized approach for large *d* (see Section 6).

If ω is a probability measure and the initial points $\{\mathbf{x}^{(j)}\}_{j=1}^{M}$ are drawn independently and at random according to ω , the strong law of large numbers shows that $\lim_{M\to\infty} [\Psi_X^* W \Psi_X]_{jk} = \langle \psi_k, \psi_j \rangle$ and $\lim_{M\to\infty} [\Psi_X^* W \Psi_Y]_{jk} = \langle \mathcal{K}\psi_k, \psi_j \rangle$ holds with probability one [72, Section 3.4] provided that ω is not supported on a zero level set that is a linear combination of the dictionary [79, Section 4]. This is with the quadrature weights $w_j = 1/M$, and the convergence is typically at a Monte Carlo rate of $\mathcal{O}(M^{-1/2})$. This argument is straightforward to adapt to show the convergence $\lim_{M\to\infty} [\Psi_Y^* W \Psi_Y]_{jk} = \langle \mathcal{K}\psi_k, \mathcal{K}\psi_j \rangle$.

For a single fixed initial condition, if the dynamical system is ergodic, then one can use Birkhoff's Ergodic theorem to show that $\lim_{M\to\infty} [\Psi_X^* W \Psi_X]_{jk} = \langle \psi_k, \psi_j \rangle$ and $\lim_{M\to\infty} [\Psi_X^* W \Psi_Y]_{jk} = \langle \mathcal{K}\psi_k, \psi_j \rangle$ [79]. One chooses $w_j = 1/M$, but the convergence rate is problem dependent [67]. This argument is straightforward to adapt to show the convergence $\lim_{M\to\infty} [\Psi_Y^* W \Psi_Y]_{jk} = \langle \mathcal{K}\psi_k, \mathcal{K}\psi_j \rangle$.

4.2 | ResDMD: Avoiding spectral pollution and computing pseudospectra

We now present two ResDMD algorithms. The first, shown in Algorithm 2, is a simple modification of EDMD to remove spectral pollution. However, this algorithm can miss parts of the spectrum because the finite section method is not guaranteed to approximate the whole of the spectrum. Our second algorithm (see Algorithm 3) overcomes the limitation of missing spectra and computes pseudospectra with error control. Further methods and results are given in Appendix B.

Throughout, our only assumption is that the dictionary $\{\psi_j\}_{j=1}^{N_K}$ lies in the domain of \mathcal{K} . For convenience, we assume that $\{\psi_j\}_{j=1}^{N_K}$ are linearly independent; otherwise, one can compute a truncated SVD and use that dictionary instead. To obtain convergence as $N_K \to \infty$, we assume that the span of all dictionary elements in the limit forms a core of \mathcal{K} (see Appendix B). Since our algorithms come with error bounds associated with residuals, we can also perform aposteri verification of a dictionary by simply looking at the errors, which is particularly useful in Section 6 (see also Appendix C).

Our first ResDMD algorithm computes the residual using the snapshot data to avoid spectral pollution. As is usually done, the algorithm assumes that K_{EDMD} is diagonalizable. First, we compute the three matrices $\Psi_X^* W \Psi_X$, $\Psi_X^* W \Psi_Y$, and $\Psi_Y^* W \Psi_Y$, where Ψ_X and Ψ_Y are given in (4.4). Then, we find the eigenvalues and eigenvectors of K_{EDMD} , that is, we solve $(\Psi_X^* W \Psi_X)^{\dagger} (\Psi_X^* W \Psi_Y) \mathbf{g} = \lambda \mathbf{g}$. One can solve this eigenproblem directly, but it is often numerically more stable to solve the generalized eigenproblem $(\Psi_X^* W \Psi_Y) \mathbf{g} = \lambda (\Psi_X^* W \Psi_X) \mathbf{g}$. Afterward, we discard computed eigenpairs with a larger relative residual than an accuracy goal of $\epsilon > 0$ to avoid spectral pollution.

Algorithm 2 summarizes the procedure and is a simple modification of EDMD, as the only difference is a clean-up where spurious eigenpairs are discarded based on their residual. This clean-up avoids spectral pollution and also removes eigenpairs that are inaccurate because

of numerical errors associated with nonnormal operators, up to the relative tolerance ϵ . The following result makes this precise.

Theorem 4.1. Let \mathcal{K} be the associated Koopman operator of (1.1) from which snapshot data is collected. Let $\Lambda_M(\epsilon)$ denote the eigenvalues in the output of Algorithm 2. Then, assuming (4.7),

$$\limsup_{M\to\infty} \max_{\lambda\in\Lambda_M(\epsilon)} \|(\mathcal{K}-\lambda)^{-1}\|^{-1} \leq \epsilon.$$

Proof. Seeking a contradiction, assume that $\limsup_{M\to\infty} \max_{\lambda\in\Lambda_M(\varepsilon)} \|(\mathcal{K}-\lambda)^{-1}\|^{-1} > \varepsilon$. Then, there is a subsequence of eigenpairs $(\lambda_k, \mathbf{g}_k)$ in the output of Algorithm 2 such that $\lambda_k \in \Lambda_{n_k}(\varepsilon)$, $n_k \to \infty$, and $\|(\mathcal{K}-\lambda_k)^{-1}\|^{-1} > \varepsilon + \delta$ for some $\delta > 0$ and all k. Since (4.7) is satisfied and $K_{\text{EDMD}} = (\Psi_X^* W \Psi_X)^{\dagger} (\Psi_X^* W \Psi_Y)$ is a sequence of bounded finite matrices for fixed N_K (recall that without loss of generality the Gram matrix $\lim_{M\to\infty} \Psi_X^* W \Psi_X$ is invertible), the sequence $\lambda_1, \lambda_2, \dots$ stays bounded. By taking a subsequence if necessary, we may assume that $\lim_{k\to\infty} \lambda_k = \lambda$. It follows that

$$\epsilon + \delta \leq \|(\mathcal{K} - \lambda)^{-1}\|^{-1} = \lim_{k \to \infty} \|(\mathcal{K} - \lambda_k)^{-1}\|^{-1} \leq \limsup_{k \to \infty} \operatorname{res}(\lambda_k, \mathbf{g}_k) \leq \epsilon,$$

which is the desired contradiction.

In the large data limit, Theorem 4.1 tells us that ResDMD computes eigenvalues inside the ϵ -pseudospectrum of \mathcal{K} and hence, avoids spectral pollution and returns reasonable eigenvalues. Despite this, Algorithm 2 may not approximate the whole ϵ -pseudospectrum of \mathcal{K} , even as $M \to \infty$ and $N_K \to \infty$. This is because the eigenvalues of K_{EDMD} may not approximate the whole spectrum of \mathcal{K} . For example, consider the shift operator of Example 3.1, which is unitary. Suppose our dictionary consists of the functions $\psi_j(k) = \delta_{k,q(j)}$, where $q : \mathbb{N} \to \mathbb{Z}$ is an enumeration of \mathbb{Z} . Then, in the large data limit, K_{EDMD} corresponds to a finite section of the shift operator and has spectrum $\{0\}$, whereas $\sigma(\mathcal{K}) = \mathbb{T}$. Hence, for $\epsilon < 1$, the output of Algorithm 2 is the empty set. This issue is known as *spectral inclusion*.

To overcome this issue, we discuss how to compute spectra and pseudospectra in Appendix B. For example, Algorithm 3 computes practical approximations of ϵ -pseudospectra with rigorous convergence guarantees. Assuming (4.7), the output of Algorithm 3 is guaranteed to be inside the ϵ -pseudospectrum of \mathcal{K} . Algorithm 3 also computes observables g with res $(\lambda, g) < \epsilon$, which are known as ϵ -approximate eigenfunctions.¹⁹

4.3 | Numerical examples

We now apply ResDMD to several dynamical systems: (1) Nonlinear pendulum, (2) Gauss iterated map, and (3) Lorenz system.

¹⁹ Previous methods to compute ϵ -approximate eigenfunctions for Koopman operators include [91], which requires the absence of continuous spectra, [53], which uses delay coordinate maps [38, 52] to deal with certain isolated eigenvalues, and [40], which uses a compactification to jointly approximate ϵ -approximate eigenfunctions of the Koopman operator associated with ergodic systems for multiple time steps.



FIGURE 7 The approximate eigenfunctions of the nonlinear pendulum visualized as phase portraits, where the color illustrates the complex argument of the eigenfunction [136]. We also plot lines of constant modulus as shadowed steps. All of these approximate eigenfunctions have residuals at most $\epsilon = 0.05$ as judged by (4.6), which can be made smaller by increasing N_K .

4.3.1 | Nonlinear pendulum revisited

We first return to the nonlinear pendulum from Section 2.2. For the dictionary of observables $\psi_1, ..., \psi_{N_K}$, we use a hyperbolic cross approximation with the standard Fourier basis (in $x_1 \in [-\pi, \pi]_{per}$) and Hermite functions (in $x_2 \in \mathbb{R}$). The hyperbolic cross approximation is an efficient way to represent functions that have bounded mixed derivatives [45, Chapter 3]. We use the trapezoidal quadrature rule discussed in Section 3.5.2 to compute $\Psi_X^* W \Psi_X$, $\Psi_X^* W \Psi_Y$, and $\Psi_Y^* W \Psi_Y$, where Ψ_X and Ψ_Y are given in (4.4). The pseudospectrum in Figure 1 (left) is computed using Algorithm 3 with $N_K = 152$ basis functions corresponding to a hyperbolic cross approximation of order 20 and $M = 10^4$ data points. The pseudospectrum in Figure 1 (right) is computed using Algorithm 3 with $N_K = 964$ basis functions corresponding to a hyperbolic cross approximation of order 100 and $M = 9 \times 10^4$ data points. Note that we can now see which EDMD eigenvalues are reliable. The output of Algorithm 3 lies inside $\sigma_{\epsilon}(\mathcal{K})$ and converges to the set $\sigma_{\epsilon}(\mathcal{K})$ as N_K increases (see Theorem B.1). Using Algorithm 3 and $N_K = 964$, we also compute some approximate eigenfunctions corresponding to $\lambda = \exp(0.4932i)$, $\lambda = \exp(0.9765i)$, $\lambda = \exp(1.4452i)$, and $\lambda = \exp(1.8951i)$ (see Figure 7). As λ moves further from 1, we typically see increased oscillations in the approximate eigenfunctions.

4.3.2 | Gauss iterated map

The Gauss iterated map is a real-valued function on the real line given by $F(x) = \exp(-\alpha x^2) + \beta$, where α and β are real parameters. It has a bifurcation diagram that resembles a mouse [65, Figure 5.17]. We consider the choice $\alpha = 2$ and $\beta = -1 - \exp(-\alpha)$, and restrict ourselves to the state-space $\Omega = [-1, 0]$ with the usual Lebesgue measure. We select a dictionary of Legendre polynomials transplanted to the interval [-1, 0], that is, $c_0P_0(2x + 1), \dots, c_{N_K-1}P_{N_K-1}(2x + 1)$, where P_j is the degree *j* Legendre polynomial on [-1, 1] and c_j are such that the dictionary is orthonormal. Figure 8 (left) shows the pseudospectra computed with $N_K = 40$ using Algorithm 3. We also show the EDMD eigenvalues, several of which are reliable (blue crosses) and correspond to the output of Algorithm 2 with $\epsilon = 0.01$.

Rather than showing the convergence of the eigenvalues and pseudospectra as $N_K \rightarrow \infty$, we demonstrate the importance of the initial conditions of the snapshot data as it determines the



Pseudospectral contours (black lines) of the Koopman operator associated with the Gauss FIGURE 8 iterated map computed using $N_{K} = 40$ for $\epsilon = 0.001, 0.01, 0.1$ and 0.3. The EDMD eigenvalues (magenta dots and blue crosses) contain a mix of accurate and spurious eigenvalues. Reliable eigenvalues corresponding to the output of Algorithm 2 with $\epsilon = 0.01$ are shown as blue crosses. Right: Convergence of the Galerkin matrix $\Psi_{Y}^{*}W\Psi_{Y}$ for different quadrature rules as the number of data points M increases.

quadrature rule in Section 4.1. If M is not sufficiently large, or the quadrature rule is poor, the Galerkin method in Algorithms 2 and 3 is not accurate. To compare initial conditions, we consider the computed matrix $\Psi_X^* W \Psi_Y$ for $N_K = 40$ with four choices of quadrature:

- Gauss-Legendre: Select $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)}$ as the Gauss-Legendre quadrature nodes transplanted to [-1, 0] and w_1, \dots, w_M as the corresponding Gauss–Legendre weights.
- Trapezoidal: Select $\mathbf{x}^{(j)} = -1 + \frac{j-1}{M-1}$ with $w_j = \frac{1}{M-1}$ for 1 < j < M and $w_1 = w_M = \frac{1}{2(M-1)}$. Riemann sum: Select $\mathbf{x}^{(j)} = -1 + \frac{j-1}{M-1}$ and $w_j = \frac{1}{M-1}$. Monte Carlo: Select $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)}$ independently at random from a uniform distribution over
- [-1,0] and $w_i = \frac{1}{N}$.

The final two rules are the most commonly used in the EDMD/Koopman literature. The trapezoidal rule uses the same data points as the Riemann sum but changes the weights at the endpoints.

Figure 8 (right) shows the accuracy, measured as $\max_{1 \le i,k \le N_{\mathcal{V}}} |\langle \Psi_x^* W \Psi_Y \rangle_{ik} - \langle \mathcal{K} \psi_k, \psi_i \rangle|$, of the four quadrature methods. Gauss-Legendre quadrature converges exponentially, trapezoidal at a rate $\mathcal{O}(M^{-2})$, Riemann sums at a rate $\mathcal{O}(M^{-1})$ and Monte Carlo at a rate $\mathcal{O}(M^{-1/2})$. The benefit of choosing the initial states at quadrature nodes is clear. Wisely choosing the weights can drastically affect convergence even when restricted to equispaced data. We never recommend using the Riemann sum, as one immediately obtains a better convergence rate by altering just two weights to employ a trapezoidal rule.

Lorenz system 4.3.3

The Lorenz system [85] is the following system of three coupled ordinary differential equations [109]:

$$\dot{X} = \sigma(Y - X), \quad \dot{Y} = X(\rho - Z) - Y, \quad \dot{Z} = XY - \beta Z.$$

33



FIGURE 9 The approximate eigenfunctions of the Lorenz system (real part shown), computed using Algorithm 3 with delay embedding. The spectral parameter is given by $\lambda = \exp(i\theta)$, and we show the residuals computed using (4.6).

The parameters σ , ρ , and β are proportional to the Prandtl number, Rayleigh number, and the physical proportions of the fluid, respectively. The system describes a truncated model of Rayleigh–Bénard convection. We take $\sigma = 10$, $\beta = 8/3$, and $\rho = 28$, corresponding to the original system studied by Lorenz.

We consider the Koopman operator corresponding to the time-step $\Delta_t = 0.05$, with dynamics on the attractor. For this Ω , one cannot explicitly write down a suitable dictionary with knowledge of the attractor. We use delay-embedding of the observables $g_j(\mathbf{x}) = [\mathbf{x}]_j$, where the subscript corresponds to the *j*th coordinate. As a dictionary, we consider a union of Krylov subspaces

$$\{g_1, \mathcal{K}g_1, \dots, \mathcal{K}^{99}g_1, g_2, \mathcal{K}g_2, \dots, \mathcal{K}^{99}g_2, g_3, \mathcal{K}g_3, \dots, \mathcal{K}^{99}g_3\}.$$

Time-delay embedding is a popular method for DMD-type algorithms [2, 70]. For this example, we build the matrices Ψ_X and Ψ_Y using a single trajectory of length $M + 1 = 10^5 + 1$ with a random initial condition on the attractor.²⁰ Figure 9 shows various approximate eigenfunctions computed using Algorithm 3. We have visualized each function as a cloud of points on the attractor. Some of

²⁰ Another suitable choice of dictionary for this example is a set of radial basis functions with centers on the attractor sampled according to the SRB measure. One could use relatively short trajectories to compute the centers and build Ψ_X and Ψ_Y using multiple trajectories of length one from the centers.

these modes resemble unstable periodic orbits, which in a sense, form a backbone of the attractor [46, 133].

5 | COMPUTING SPECTRAL MEASURES USING A ResDMD FRAMEWORK

In Section 3, we developed algorithms that computed the spectral properties of isometries from autocorrelations. Computing autocorrelations typically requires data collected as long trajectories, that is, large M_2 . In this section, we develop methods that build on ResDMD to compute spectral measures of Koopman operators for measure-preserving dynamical systems from arbitrary snapshot data. Throughout this section, we assume that \mathcal{K} is an isometry and that we have access to snapshot data, B_{data} , of the form given in (4.1). That is, we have pairs $\{(\mathbf{x}^{(j)}, \mathbf{y}^{(j)} = F(\mathbf{x}^{(j)})\}_{j=1}^{M}\}$. We develop rational kernels that allow us to compute smoothed approximations of spectral measures from the ResDMD matrices $\Psi_X^* W \Psi_X$, $\Psi_X^* W \Psi_Y$, and $\Psi_Y^* W \Psi_Y$ (see Section 4.1). Moreover, these matrices can be reused to compute spectral measures with respect to different observable functions g.

5.1 | Generalized Cauchy transform and the Poisson kernel for the unit disc

As in Section 2.1.1, we begin by considering a unitary extension \mathcal{K}' of \mathcal{K} , which is defined on a Hilbert space \mathcal{H}' that is an extension of $L^2(\Omega, \omega)$. Let $z \in \mathbb{C}$ with |z| > 1 and $g \in L^2(\Omega, \omega)$. Since $\|\mathcal{K}\| = 1 < |z|$ for $z \notin \sigma(\mathcal{K}), (\mathcal{K}' - z)^{-1}g = (\mathcal{K} - z)^{-1}g$ and

$$\langle (\mathcal{K}-z)^{-1}g, \mathcal{K}^*g \rangle = \langle \mathcal{K}'(\mathcal{K}'-z)^{-1}g, g \rangle_{\mathcal{H}'} = \int_{\mathbb{T}} \frac{\lambda}{\lambda-z} \, d\mu_g(\lambda) = \int_{[-\pi,\pi]_{\text{per}}} \frac{e^{i\theta}}{e^{i\theta}-z} \, d\nu_g(\theta), \quad (5.1)$$

where the last equality follows from a change-of-variables.²¹ If $z \neq 0$ with |z| < 1, then z may be in $\sigma(\mathcal{K})$ since \mathcal{K} is not necessarily unitary. However, since $|\overline{z}^{-1}| > 1$, $\overline{z}^{-1} \notin \sigma(\mathcal{K})$ and hence $(\mathcal{K}' - \overline{z}^{-1})^{-1}g = (\mathcal{K} - \overline{z}^{-1})^{-1}g$. Since ν_g is a real-valued measure, we find that

$$\langle g, (\mathcal{K} - \overline{z}^{-1})^{-1}g \rangle = \langle g, (\mathcal{K}' - \overline{z}^{-1})^{-1}g \rangle_{\mathcal{H}'} = \overline{\int_{[-\pi,\pi]_{\text{per}}} \frac{1}{e^{i\theta} - \overline{z}^{-1}} d\nu_g(\theta)} = -z \int_{[-\pi,\pi]_{\text{per}}} \frac{e^{i\theta}}{e^{i\theta} - z} d\nu_g(\theta).$$
(5.2)

The leftmost and rightmost sides of (5.1) and (5.2) are independent of \mathcal{K}' , so we can safely dispense with the extension and have an expression for a generalized Cauchy transform of ν_g , that is,

$$C_{\nu_g}(z) = \frac{1}{2\pi} \int_{[-\pi,\pi]_{\text{per}}} \frac{e^{i\theta}}{e^{i\theta} - z} \, d\nu_g(\theta) = \frac{1}{2\pi} \begin{cases} \langle (\mathcal{K} - z)^{-1}g, \mathcal{K}^*g \rangle, & \text{if } |z| > 1, \\ -z^{-1}\langle g, (\mathcal{K} - \overline{z}^{-1})^{-1}g \rangle, & \text{if } z \neq 0 \text{ with } |z| < 1. \end{cases}$$
(5.3)

The importance of (5.3) is that it relates C_{ν_g} to the resolvent operator $(\mathcal{K} - z)^{-1}$ for |z| > 1. In Section 5.3, we show how to compute the resolvent operator from snapshot data for |z| > 1.

²¹ Here we use the fact that the resolvent can be written as $(\mathcal{K}' - z)^{-1} = \int_{\mathbb{T}} (\lambda - z)^{-1} d\mathcal{E}(\lambda)$, where \mathcal{E} is the projection-valued spectral measure of \mathcal{K}' . For expansions of the resolvent in the context of different dynamical systems, see [124].

Since |z| > 1, we can provide convergence and stability results even when we replace \mathcal{K} with a discretization.

To recover ν_g from C_{ν_a} , a derivation motivated by the Sokhotski–Plemelj formula shows that

$$C_{\nu_g}(e^{i\theta_0}(1+\epsilon)^{-1}) - C_{\nu_g}(e^{i\theta_0}(1+\epsilon)) = \int_{[-\pi,\pi]_{\text{per}}} K_{\epsilon}(\theta_0 - \theta) \, d\nu_g(\theta), \tag{5.4}$$

where K_{ϵ} is the Poisson kernel for the unit disc given in (3.9). The Poisson kernel for the unit disc is a first-order kernel (see Proposition 5.1). In practice, it is important to develop high-order kernels for at least three reasons: (1) A numerical tradeoff, (2) A stability tradeoff, and (3) For localization of the kernel. We now explain each of these reasons in turn.

5.1.1 | The need for high-order kernels I: A numerical tradeoff

As a model problem, consider the following generalization of the shift operator in Example 3.1:

$$U = \begin{pmatrix} \alpha_{0} & \alpha_{1}\rho_{0} & \rho_{1}\rho_{0} & & \\ \rho_{0} & -\overline{\alpha_{1}}\alpha_{0} & -\rho_{1}\alpha_{0} & 0 & \\ 0 & \overline{\alpha_{2}}\rho_{1} & -\overline{\alpha_{2}}\alpha_{1} & \overline{\alpha_{3}}\rho_{2} & \rho_{3}\rho_{2} & \\ & \rho_{2}\rho_{1} & -\rho_{2}\alpha_{1} & -\overline{\alpha_{3}}\alpha_{2} & -\rho_{3}\alpha_{2} & \ddots \\ & 0 & \overline{\alpha_{4}}\rho_{3} & -\overline{\alpha_{4}}\alpha_{3} & \ddots \\ & & \ddots & \ddots & \ddots \end{pmatrix}, \qquad \alpha_{j} = (-1)^{j}0.95^{\frac{j+1}{2}}, \quad \rho_{j} = \sqrt{1 - |\alpha_{j}|^{2}}.$$
(5.5)

We seek the spectral measure with respect to the first canonical unit vector $g = e_1$, which is a cyclic vector. *U* is the Cantero–Moral–Velázquez (CMV) matrix for the Rogers–Szegő orthogonal polynomials on the unit circle [106]. CMV matrices for measures on the unit circle are analogous to Jacobi matrices for measures on the real line [120].²² The spectral measure ν_g is absolutely continuous with an analytic Radon–Nikodym derivative given by

$$\rho_g(\theta) = \sqrt{\frac{2\pi}{\log(0.95^{-1})}} \sum_{m=-\infty}^{\infty} \exp\left(-\frac{(\theta - 2\pi m)^2}{2\log(0.95^{-1})}\right).$$

To compute the resolvent and hence C_{ν_g} , we use finite square truncations of the matrix U (see Section 5.3). Since U is banded, once it is truncated to an $N_K \times N_K$ matrix it costs $\mathcal{O}(N_K)$ operations to compute C_{ν_g} . The dominating computational cost in evaluating $[K_{\epsilon} * \nu_g]$ is solving the shifted linear systems to compute the resolvent $(\mathcal{K} - e^{i\theta_0}(1 + \epsilon))^{-1}$. This cost increases as $\epsilon \downarrow 0$ and $e^{i\theta_0}(1 + \epsilon)$ approaches the continuous spectrum of \mathcal{K} . This increase in cost is due to a larger truncation size of U needed to capture the function $(\mathcal{K} - e^{i\theta_0}(1 + \epsilon))^{-1}g$ because it becomes more oscillatory.

Figure 10 (left) shows the approximations $[K_{\epsilon} * \nu_g]$ for the Poisson kernel and $\epsilon = 0.1, 0.01$ using $N_K = 40$ basis functions for the operator U in (5.5). The exact Radon–Nikodym derivative is also shown. Whilst the computation for $\epsilon = 0.1$ is an accurate approximation of $[K_{\epsilon} * \nu_g]$, we must decrease ϵ to approximate ρ_g . However, decreasing ϵ to 0.01 introduces oscillatory artifacts.

²² One applies Gram–Schmidt to $\{1, z, z^{-1}, z^2, z^{-2}, ...\}$, which leads to a pentadiagonal matrix.



FIGURE 10 Left: Approximation of the spectral measure of U in (5.5) with respect to $g = e_1$ for a fixed truncation size of $N_K = 40$. The oscillations for $\epsilon = 0.01$ are an artifact caused by the discrete spectrum of the truncation. Right: The relative error in the approximation v_{g,N_K}^{ϵ} with varying N_K of the smoothed measure $v_g^{\epsilon} = [K_{\epsilon} * v_g]$ for the Poisson kernel evaluated at $\theta = 0.2$.

This is because we have fixed the truncation size, and the truncated operator's spectral measure is a sum of Dirac delta distributions located in the closed unit disc. To avoid this, we must select N_K adaptively for a given ϵ (see Section 5.3). Figure 10 (right) shows the relative error of approximating $[K_{\epsilon} * \nu_g](0.2)$ for a truncation size N_K and varying ϵ . A smaller ϵ leads to slower convergence as the truncation size increases.

Therefore, there is a numerical balancing act. On the one hand, we wish to stay as far away from the spectrum as possible so that the evaluation of $[K_{\epsilon} * \nu_g]$ is computationally cheap. On the other hand, we desire a good approximation of ν_g , which requires small $\epsilon > 0$. The Poisson kernel leads to $\mathcal{O}(\epsilon \log(\epsilon^{-1}))$ convergence as $\epsilon \downarrow 0$, making it computationally infeasible to compute the spectral measure to more than a few digits of accuracy. High-order kernels allow us to compute spectral measures with better accuracy. This tradeoff is significant for data-driven computations since we want to keep N_K as small as possible.

5.1.2 | The need for high-order kernels II: A stability tradeoff

Continuing with (5.5), the shifted linear system involved in computing the resolvent operator becomes increasingly ill-conditioned as $\epsilon \downarrow 0$ when $e^{i\theta_0} \in \sigma(\mathcal{K})$. This is because $\|(\mathcal{K} - e^{i\theta_0}(1 + \epsilon))^{-1}\| = \mathcal{O}(\epsilon^{-1})$ (see Section 5.2.2). The ResDMD framework developed in Section 4 uses approximations of Galerkin truncations of \mathcal{K} , with convergence to the exact truncation in the large data limit $M \to \infty$ as the number of data points increases (see Section 4.1.3). Therefore, the stability of the resolvent operator is crucial for data-driven approximations as we have inexact matrix entries.

We now repeat the experiment from Section 5.1.1, but with an adaptive N_K for fixed ϵ , and we perturb each entry of an $N_K \times N_K$ truncation of U with independent Gaussian random variables of mean 0 and standard deviation $\delta = 0.001$. Figure 11 (left) shows the approximation of the spectral measure ν_g with $g = e_1$. While $\epsilon = 0.1$ appears stable to the eye, decreasing to $\epsilon = 0.01$ yields an inaccurate approximation since $\|(\mathcal{K} - e^{i\theta_0}(1 + \epsilon))^{-1}\|$ increases and the shifted linear systems become more sensitive to noise. Figure 11 (right) shows the mean error of the approximation of $[K_{\epsilon} * \nu_g]$ over 50 independent realizations. In contrast to the convergence to machine precision in



FIGURE 11 Left: Approximation of the spectral measure of *U* in (5.5) with respect to $g = e_1$ for an adaptive choice of N_K , but with the truncated matrix perturbed by normal random variables of mean 0 and standard deviation $\delta = 0.001$. The oscillations for $\epsilon = 0.01$ are an artifact caused by the ill-conditioning of the resolvent as $\epsilon \downarrow 0$. Right: The relative error in the numerical approximation ν_{g,N_K}^{ϵ} , corresponding to truncation size N_K , of the smoothed measure $\nu_g^{\epsilon} = [K_{\epsilon} * \nu_g]$ for the Poisson kernel in (3.9) with the same model of matrix perturbation. The error is averaged over 50 realizations.

Figure 10 (right), the errors in Figure 11 (right) plateau because of the perturbation of the matrix, with a larger final error for smaller ϵ .

The error in approximating ν_g with a perturbed truncation of the operator is unavoidable, but the Poisson kernel is suboptimal in this regard. For a perturbation of order δ , the accuracy achieved by the Poisson kernel scales as $\mathcal{O}(\epsilon \log(\epsilon^{-1}) + \delta/\epsilon)$. Balancing these terms while ignoring logarithmic factors leads to the noise-dependent choice $\epsilon \sim \delta^{1/2}$. This choice yields an overall error scaling of $\mathcal{O}(\delta^{1/2} \log(\delta^{-1}))$. High-order kernels can get arbitrarily close to the optimal scaling of $\mathcal{O}(\delta)$ as $\delta \to 0$.

5.1.3 | The need for high-order kernels III: Localization of the kernel

We revisit the tent map of Section 3.5.1 with the choice of g given in (3.23). Figure 12 shows the computed measure using the Poisson kernel, where we have chosen N_K adaptively to avoid the problem in Section 5.1.1 and $\epsilon = 0.1$ so that with $M = 2^5 = 32$ the error in Section 5.1.2 is negligible. For comparison, we do the same computation, but with a 6th-order kernel (see Section 5.2.1). As a function, the Poisson kernel decays slowly away from 0, and hence the singular part of the spectral measure at $\theta = 0$ is not localized. In contrast, *m*th order kernels obey a decay condition (see (3.11)) so are more localized for a particular ϵ . For this example, this localization leads to a better approximation of the Radon–Nikodym derivative away from the singularity. We have also considered examples where the spectral measure is singular and observed that increased localization has the added benefit of avoiding over-smoothing.

5.2 | Building high-order rational kernels

We now aim to generalize the Sokhotski–Plemelj-like formula in (5.4) to develop high-order rational kernels. Let $\{z_j\}_{i=1}^m$ be distinct points with positive real part and consider the rational



FIGURE 12 Convolutions of spectral measure with Poisson kernel (m = 1) and a 6th order rational kernel (see Section 5.2). High-order kernels have better localization properties.

function:

$$K_{\varepsilon}(\theta) = \frac{e^{-i\theta}}{2\pi} \sum_{j=1}^{m} \left[\frac{c_j}{e^{-i\theta} - (1 + \varepsilon \overline{z_j})^{-1}} - \frac{d_j}{e^{-i\theta} - (1 + \varepsilon z_j)} \right].$$
(5.6)

A short derivation using (5.3) shows that

$$\int_{[-\pi,\pi]_{\text{per}}} K_{\varepsilon}(\theta_{0}-\theta) \, d\nu_{g}(\theta) = \sum_{j=1}^{m} \left[c_{j} C_{\nu_{g}} \left(e^{i\theta_{0}} (1+\epsilon\overline{z_{j}})^{-1} \right) - d_{j} C_{\nu_{g}} \left(e^{i\theta_{0}} (1+\epsilon\overline{z_{j}}) \right) \right]$$
$$= \frac{-1}{2\pi} \sum_{j=1}^{m} \left[c_{j} e^{-i\theta_{0}} (1+\epsilon\overline{z_{j}}) \langle g, (\mathcal{K}-e^{i\theta_{0}} (1+\epsilon\overline{z_{j}}))^{-1}g \rangle + d_{j} \langle (\mathcal{K}-e^{i\theta_{0}} (1+\epsilon\overline{z_{j}}))^{-1}g, \mathcal{K}^{*}g \rangle \right].$$
(5.7)

It follows that we can compute the convolution $[K_{\epsilon} * \nu_g](\theta_0)$ by evaluating the resolvent at the *m* points $\{e^{i\theta_0}(1 + \epsilon z_j)\}_{j=1}^m$. We use rational kernels because they allow us to compute smoothed approximations of the spectral measure by applying the resolvent operator to functions.

However, for (5.7) to be a good approximation of ν_g , we must carefully select the points z_j and the coefficients $\{c_j, d_j\}$ in (5.6). In particular, we would like $\{K_{\varepsilon}\}$ to be an *m*th order kernel. To design such a kernel, we carefully consider Lemma A.1. First, we define $\zeta_j(\varepsilon)$ by the relationship $1 + \varepsilon \zeta_j(\varepsilon) = (1 + \varepsilon \overline{z_j})^{-1}$ and use Cauchy's Residue theorem to show that for any integer $n \ge 1$,

$$\begin{split} \int_{[-\pi,\pi]_{\text{per}}} K_{\epsilon}(-\theta) e^{in\theta} \, d\theta &= \frac{1}{2\pi i} \int_{\mathbb{T}} \left[\sum_{j=1}^{m} \frac{c_j}{\lambda - (1 + \epsilon \overline{z_j})^{-1}} - \sum_{j=1}^{m} \frac{d_j}{\lambda - (1 + \epsilon z_j)} \right] \lambda^n \, d\lambda \\ &= \sum_{j=1}^{m} c_j (1 + \epsilon \overline{z_j})^{-n} = \left(\sum_{j=1}^{m} c_j \right) + \sum_{k=1}^{n} \epsilon^k \binom{n}{k} \sum_{j=1}^{m} c_j \zeta_j(\epsilon)^k. \end{split}$$

It follows that condition (A.1) in Lemma A.1 holds if

$$\begin{pmatrix} 1 & \dots & 1 \\ \zeta_1(\varepsilon) & \dots & \zeta_m(\varepsilon) \\ \vdots & \ddots & \vdots \\ \zeta_1(\varepsilon)^{m-1} & \dots & \zeta_m(\varepsilon)^{m-1} \end{pmatrix} \begin{pmatrix} c_1(\varepsilon) \\ c_2(\varepsilon) \\ \vdots \\ c_m(\varepsilon) \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}.$$
(5.8)

Note also that, if this holds, the coefficients $c_j = c_j(\epsilon)$ remain bounded as $\epsilon \downarrow 0$. To ensure that the decay condition in (A.2) is satisfied, let $\omega = (e^{-i\theta} - 1)/\epsilon$. The kernel in (5.6) can then be re-written as

$$K_{\varepsilon}(\theta) = \frac{\varepsilon^{-1} e^{-i\theta}}{2\pi} \sum_{j=1}^{m} \left[\frac{c_j}{\omega - \zeta_j(\varepsilon)} - \frac{d_j}{\omega - z_j} \right].$$
(5.9)

Therefore, we have

$$\begin{split} \omega K_{\varepsilon}(\theta) &= \frac{\varepsilon^{-1} e^{-i\theta}}{2\pi} \sum_{j=1}^{m} \left[c_j + \frac{c_j \zeta_j(\varepsilon)}{\omega - \zeta_j(\varepsilon)} - d_j - \frac{d_j z_j}{\omega - z_j} \right] \\ &= \frac{\varepsilon^{-1} e^{-i\theta}}{2\pi} \sum_{j=1}^{m} (c_j - d_j) + \frac{\varepsilon^{-1} e^{-i\theta}}{2\pi} \sum_{j=1}^{m} \left[\frac{c_j \zeta_j(\varepsilon)}{\omega - \zeta_j(\varepsilon)} - \frac{d_j z_j}{\omega - z_j} \right]. \end{split}$$

By repeating the same argument *m* times, we arrive at

$$\omega^m K_{\epsilon}(\theta) = \frac{\epsilon^{-1} e^{-i\theta}}{2\pi} \left[\sum_{k=0}^{m-1} \omega^{m-1-k} \sum_{j=1}^m (c_j \zeta_j(\epsilon)^k - d_j z_j^k) + \sum_{j=1}^m \left(\frac{c_j \zeta_j(\epsilon)^m}{\omega - \zeta_j(\epsilon)} - \frac{d_j z_j^m}{\omega - z_j} \right) \right].$$
(5.10)

This means that we should select the d_k 's so that

$$\begin{pmatrix} 1 & \dots & 1 \\ z_1 & \dots & z_m \\ \vdots & \ddots & \vdots \\ z_1^{m-1} & \dots & z_m^{m-1} \end{pmatrix} \begin{pmatrix} d_1 \\ d_2 \\ \vdots \\ d_m \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}.$$
(5.11)

We conclude that if the coefficients $\{c_j\}_{j=1}^m$ and $\{d_j\}_{j=1}^m$ satisfy (5.8) and (5.11), respectively, then

$$\sum_{k=0}^{m-1} \omega^{m-1-k} \sum_{j=1}^m (c_j \zeta_j(\epsilon)^k - d_j z_j^k) = 0, \qquad \left| \sum_{j=1}^m \left(\frac{c_j \zeta_j(\epsilon)^m}{\omega - \zeta_j(\epsilon)} - \frac{d_j z_j^m}{\omega - z_j} \right) \right| \lesssim |\omega|^{-1}.$$

By (5.10), this means that $|\omega|^m |K_{\epsilon}(\theta)| \leq \epsilon^{-1} |\omega|^{-1}$. Moreover, since $|K_{\epsilon}(\theta)| \leq \epsilon^{-1}$ we see that $|K_{\epsilon}(\theta)| \leq \min\{\epsilon^{-1}, \epsilon^m |\theta|^{-(m+1)}\} \leq \epsilon^m (\epsilon + |\theta|)^{-(m+1)}$. Using Lemma A.1, we have proved the following proposition.

Proposition 5.1. Let $\{z_j\}_{j=1}^m$ be distinct points with positive real part and let K_{ε} be given by (5.6). Then, $\{K_{\varepsilon}\}$ is an mth order kernel for $[-\pi, \pi]_{\text{per}}$ if the coefficients $\{c_j\}_{j=1}^m$ and $\{d_j\}_{j=1}^m$ satisfy (5.8) and (5.11), respectively.



FIGURE 13 The *m*th-order kernels (5.6) constructed with the choice (5.12) for $\epsilon = 1$ and $1 \le m \le 6$.

TABLE 1 Coefficients in the rational kernels in (5.6) for $1 \le m \le 6$, the choice (5.12), and $\epsilon = 0.1$. We give the first $\lfloor m/2 \rfloor$ coefficients as $c_{m+1-j} = \overline{c_j}$ and $d_{m+1-j} = \overline{d_j}$.

m	$\{d_1,\ldots,d_{\lceil m/2\rceil}\}$	$\{c_1(\epsilon),, c_{\lceil m/2 \rceil}(\epsilon)\}, \epsilon = 0.1$
2	$\left\{\frac{1-3i}{2}\right\}$	$\left\{ \frac{3+10i}{6} \right\}$
3	$\{-2-i,5\}$	$\left\{\frac{-202+79i}{80},\frac{121}{20}\right\}$
4	$\left\{\frac{-39+65i}{24},\frac{17-85i}{8}\right\}$	$\left\{\frac{-1165710-2944643i}{750000},\frac{513570+3570527i}{250000}\right\}$
5	$\left\{\frac{15+10i}{4}, \frac{-39-13i}{2}, \frac{65}{2}\right\}$	$\left\{\frac{4052283 - 1460282i}{648000}, \frac{-2393157 + 486551i}{81000}, \frac{190333}{4000}\right\}$
6	$\left\{\frac{725-1015i}{192}, \frac{-2775+6475i}{192}, \frac{1073-7511i}{96}\right\}$	$\left\{\frac{24883929805+81589072062i}{8067360000}, \frac{-19967590755-93596942182i}{1613472000}, \frac{7898770397+102424504746i}{806736000}\right\}$

5.2.1 | The choice of rational kernel

We are free to choose the points $\{z_j\}_{j=1}^m$ in Proposition 5.1 subject to $\operatorname{Re}(z_j) > 0$, after which the linear systems (5.8) and (5.11) provide suitable $\{c_j\}_{j=1}^m$ (dependent on ϵ) and $\{d_j\}_{j=1}^m$ (independent of ϵ). As a natural extension of the Poisson kernel in (3.9), we select the points $\{z_j\}_{j=1}^m$ as

$$z_j = 1 + \left(\frac{2j}{m+1} - 1\right)i, \qquad 1 \le j \le m.$$
 (5.12)

The kernels that we have developed are typically not real-valued. Since ν_g is a probability measure and hence real-valued, we often gain better accuracy for a particular ϵ by considering the kernel Re(K_{ϵ}), and this is what we do throughout this paper. The convolution with Re(K_{ϵ}) can be computed by taking the real part of the right-hand side of (5.7). The first six kernels with the choice $\epsilon = 1$ are shown in Figure 13. The exact coefficients $\{c_j, d_j\}_{j=1}^m$ for $\epsilon = 0.1$ are shown in Table 1 for the first six kernels.



FIGURE 14 Left: The pointwise relative error in smoothed measures of the operator in (5.5) computed using the high-order kernels in Section 5.2.1 for $1 \le m \le 6$. Right: The pointwise relative error in the computed smoothed measure using the perturbed K_{N_k} and $\epsilon \sim \delta^{1/(m+1)}$.

5.2.2 | Revisiting the examples in Section 5.1.1 and Section 5.1.2

To demonstrate the advantages of high-order rational kernels, we revisit the operator U given in (5.5) and look again at two of the numerical needs for high-order kernels.

- A numerical tradeoff. Figure 14 (left) shows the convergence of $[K_{\varepsilon} * \nu_g]$ for the rational kernels in Section 5.2.1. We see the convergence rates predicted from Theorem 3.1. In practice, this means that we rarely need to take ε below 0.01 for very accurate results, and often we take $\varepsilon = 0.1$.
- A stability tradeoff. Let K_{N_K} be an $N_K \times N_K$ truncation of \mathcal{K} (see Section 5.3 for how we do this). An important property of our truncations is that $||K_{N_K}|| \le ||\mathcal{K}|| = 1$. Using the standard Neumann series for the inverse, it follows that any λ such that $|\lambda| > 1$ has $||(K_{N_K} \lambda)^{-1}|| \le 1/(|\lambda| 1)$. In general, this is the best one can hope for. Suppose that instead of K_{N_K} , we have access to $\tilde{K}_{N_K} = K_{N_K} + \Delta$ with $||\Delta|| \le \delta$. If $|\lambda| = 1 + \mathcal{O}(\epsilon) > 1 + \delta$, then since $||\tilde{K}_{N_K}|| \le 1 + \delta$ we have $||(\tilde{K}_{N_K} \lambda)^{-1}|| \le 1/(|\lambda| 1 \delta)$. Using the second resolvent identity for the first equality, the relative error is bounded via

$$\frac{\|[(\tilde{K}_{N_{K}}-\lambda)^{-1}-(K_{N_{K}}-\lambda)^{-1}]g\|}{\|(K_{N_{K}}-\lambda)^{-1}g\|} = \frac{\|[-(\tilde{K}_{N_{K}}-\lambda)^{-1}\Delta](K_{N_{K}}-\lambda)^{-1}g\|}{\|(K_{N_{K}}-\lambda)^{-1}g\|} \le \|(\tilde{K}_{N_{K}}-\lambda)^{-1}\Delta\| \le \frac{\delta}{|\lambda|-1-\delta} = \mathcal{O}(\epsilon^{-1}\delta).$$

Theorems 3.1 and 3.2 show that an *m*th order kernel provides error rates of $\mathcal{O}(\epsilon^m \log(\epsilon^{-1}))$ as $\epsilon \to 0$. Therefore, we have an error that is asymptotically of size $\mathcal{O}(\epsilon^m \log(\epsilon^{-1}))$ from using a kernel and $\mathcal{O}(\epsilon^{-1}\delta)$ from using a perturbed version of K_{N_K} . We assume that N_K is sufficiently large to render the truncation error negligible. By selecting $\epsilon = \delta^{1/(m+1)}$, we balance these two errors and have an overall error of $\mathcal{O}(\delta^{\frac{m}{m+1}} \log(\delta^{-1}))$. For large *m*, the convergence rate is close to the optimal rate of $\mathcal{O}(\delta)$. In Figure 14 (right), we show the relative error $|\nu_g^{\epsilon}(0.2) - \rho_g(0.2)|/|\rho_g(0.2)|$ when using truncations where each entry is perturbed by independent normal

random variables of mean 0 and standard deviation $\delta^{.23}$ We have chosen N_K sufficiently large and $\epsilon \sim \delta^{1/(m+1)}$. The error is averaged over 50 trials so that we show the expected convergence rates.

5.3 | An algorithm for evaluating a smoothed spectral measure

We now seek to evaluate $[K_{\varepsilon} * \nu_g]$ at $\theta_1, ..., \theta_P \in [-\pi, \pi]_{\text{per}}$, where $\{K_{\varepsilon}\}$ is a rational kernel of the form (5.6). This requires the evaluation of $(\mathcal{K} - \lambda)^{-1}$ for different values of λ (see (5.7)). We first describe how to evaluate $[K_{\varepsilon} * \nu_g]$ at a single point θ_0 before precomputing a factorization to derive an efficient scheme for evaluating at $\theta_1, ..., \theta_P$.

5.3.1 | Evaluating a smoothed spectral measure at a single point

To evaluate $[K_{\varepsilon} * \nu_g]$ at a single point $\theta_0 \in [-\pi, \pi]_{per}$, we use the setup of ResDMD (see Section 4.1) so that we can obtain rigorous a posteriori error bounds on the computed resolvents, allowing us to adaptively select the dictionary size N_K based on the smoothing parameter ε . Since \mathcal{K} is an isometry, we only need to compute $(\mathcal{K} - z)^{-1}$ for $|z| > 1 = ||\mathcal{K}||$ and so we can achieve our goal with standard Galerkin truncations of \mathcal{K} .

Theorem 5.1. Suppose that \mathcal{K} is an isometry and $\lambda \in \mathbb{C}$ with $|\lambda| > 1$. Let $\psi_1, \psi_2, ...$ be a dictionary of observables and $V_{N_K} = \operatorname{span}\{\psi_1, ..., \psi_{N_K}\}$, so that $\bigcup_{N_K \in \mathbb{N}} V_{N_K}$ is dense in $L^2(\Omega, \omega)$. Then, for any sequence of observables $g_{N_K} \in V_{N_K}$ such that $\lim_{N_K \to \infty} g_{N_K} = g \in L^2(\Omega, \omega)$,

$$\lim_{N_K\to\infty} \left(\mathcal{P}_{V_{N_K}} \mathcal{K} \mathcal{P}^*_{V_{N_K}} - \lambda I_{N_K} \right)^{-1} g_{N_K} = (\mathcal{K} - \lambda)^{-1} g,$$

where $\mathcal{P}_{V_{N_{K}}}$ is the orthogonal projection operator onto $V_{N_{K}}$ and $I_{N_{K}}$ is the $N_{K} \times N_{K}$ identity matrix.

Proof. Since $|\lambda| > ||\mathcal{K}||$ and $||\mathcal{P}_{V_{N_{K}}}\mathcal{KP}^{*}_{V_{N_{K}}}|| \le ||\mathcal{K}||$, we have the following absolutely convergent Neumann series:

$$\left(\mathcal{P}_{V_{N_{K}}}\mathcal{K}\mathcal{P}_{V_{N_{K}}}^{*}-\lambda I_{N_{K}}\right)^{-1}=-\frac{1}{\lambda}\sum_{k=0}^{\infty}\frac{1}{\lambda^{k}}(\mathcal{P}_{V_{N_{K}}}\mathcal{K}\mathcal{P}_{V_{N_{K}}}^{*})^{k},\quad (\mathcal{K}-\lambda)^{-1}=-\frac{1}{\lambda}\sum_{k=0}^{\infty}\frac{1}{\lambda^{k}}\mathcal{K}^{k}.$$

Using $|\lambda| > ||\mathcal{K}||$, we can bound the tails of these series uniformly in N_K so we just need to show that $(\mathcal{P}_{V_{N_K}} \mathcal{KP}^*_{V_{N_K}})^k$ converges strongly to \mathcal{K}^k for $k \ge 0$. This clearly holds for k = 0 and k = 1. For k > 1, we apply induction using the fact that if a sequence of matrices T_{N_K} and S_{N_K} converge strongly to \mathcal{T} and \mathcal{S} , respectively, then the product $T_{N_K}S_{N_K}$ converges strongly to the composition $\mathcal{T} \circ \mathcal{S}$.

We now apply Theorem 5.1 to evaluate $[K_{\varepsilon} * \nu_g]$ at θ_0 . Recall from (5.7) that there are two types of inner products to compute: (i) $\langle g, (\mathcal{K} - \lambda)^{-1}g \rangle$ and (ii) $\langle (\mathcal{K} - \lambda)^{-1}g, \mathcal{K}^*g \rangle$ for some observable g. We form a sequence of observables $g_{N_K} \in V_{N_K}$ by setting $g_{N_K} = \mathcal{P}_{V_{N_K}}g$, which can be

²³ Strictly speaking, the operator norm of this perturbation need not be bounded by δ , but the scaling argument for the choice $\epsilon = \delta^{1/(m+1)}$ remains the same.

43

approximated from the snapshot data as

$$\tilde{g}_{N_{K}} = \sum_{j=1}^{N_{K}} \boldsymbol{a}_{j} \psi_{j}, \qquad \boldsymbol{a} = (\Psi_{X}^{*} W \Psi_{X})^{-1} \Psi_{X}^{*} W \begin{bmatrix} g(\boldsymbol{x}^{(1)}) \\ \vdots \\ g(\boldsymbol{x}^{(M)}) \end{bmatrix} \in \mathbb{C}^{N_{K}},$$
(5.13)

where Ψ_X and W are given in Section 4.1.1. Under suitable conditions, such as those discussed in Section 4.1.3, $\lim_{M\to\infty} \tilde{g}_{N_K} = g_{N_K}$. Since

$$\left(\mathcal{P}_{V_{N_{K}}}\mathcal{K}\mathcal{P}_{V_{N_{K}}}^{*}-\lambda I_{N_{K}}\right)^{-1}g_{N_{K}}=\lim_{M\to\infty}\sum_{j=1}^{N_{K}}\left[\left(\Psi_{X}^{*}W\Psi_{Y}-\lambda\Psi_{X}^{*}W\Psi_{X}\right)^{-1}\Psi_{X}^{*}W\Psi_{X}\boldsymbol{a}\right]_{j}\psi_{j},$$

it follows that our two types of inner products satisfy

$$\left\langle g_{N_{K}}, \left(\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}} \right)^{-1} g_{N_{K}} \right\rangle = \lim_{M \to \infty} \overline{\boldsymbol{a}^{*} \Psi_{X}^{*} W \Psi_{X} (\Psi_{X}^{*} W \Psi_{Y} - \lambda \Psi_{X}^{*} W \Psi_{X})^{-1} \Psi_{X}^{*} W \Psi_{X} \boldsymbol{a}},$$
(5.14)

$$\left\langle \left(\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}} \right)^{-1} g_{N_{K}}, \mathcal{K}^{*} g_{N_{K}} \right\rangle = \lim_{M \to \infty} \boldsymbol{a}^{*} \Psi_{X}^{*} W \Psi_{Y} (\Psi_{X}^{*} W \Psi_{Y} - \lambda \Psi_{X}^{*} W \Psi_{X})^{-1} \Psi_{X}^{*} W \Psi_{X} \boldsymbol{a}.$$

$$(5.15)$$

For a given value of M, the right-hand side of (5.14) and (5.15) can then be substituted into (5.7) to evaluate $[K_{\epsilon} * \nu_g](\theta_0)$. Often we can estimate the error between these computed inner products and the limiting value as $M \to \infty$ by comparing the computations for different M or by using a priori knowledge of the convergence rates. In Appendix C, we show how N_K can be adaptively chosen (by approximating the error in the large data limit and adding observables to the dictionary if required) so that we approximate the inner products $\langle g, (\mathcal{K} - \lambda)^{-1}g \rangle$ and $\langle (\mathcal{K} - \lambda)^{-1}g, \mathcal{K}^*g \rangle$ to a desired accuracy. Thus, for a given smoothing parameter ϵ , we have a principled way of selecting (a) the sample size M and (b) the truncation size N_K to ensure that our approximations of the inner products in (5.7) are accurate. In general, the cost of point evaluation of $[K_{\epsilon} * \nu_g]$ using these formulas is $\mathcal{O}(N_K^3)$ operations as it requires m solutions of $N_K \times N_K$ dense linear systems.

5.3.2 | Evaluating a smoothed spectral measure at multiple points

To evaluate $[K_{\varepsilon} * \nu_g]$ at $\theta_1, ..., \theta_P \in [-\pi, \pi]_{per}$, one can be more computationally efficient than independently computing each of the inner products in (5.14) and (5.15) for each θ_k for $1 \le k \le P$. Instead, one can compute a generalized Schur decomposition and to speed up the evaluation. Let $\Psi_X^* W \Psi_Y = QSZ^*$ and $\Psi_X^* W \Psi_X = QTZ^*$ be a generalized Schur decomposition, where Q and Z are unitary matrices, and S and T are upper-triangular matrices. With this decomposition in hand,

$$\boldsymbol{a}^* \Psi_X^* W \Psi_X (\Psi_X^* W \Psi_Y - \lambda \Psi_X^* W \Psi_X)^{-1} \Psi_X^* W \Psi_X \boldsymbol{a} = \boldsymbol{a}^* Q T (S - \lambda T)^{-1} T Z^* \boldsymbol{a},$$

$$\boldsymbol{a}^* \Psi_X^* W \Psi_Y (\Psi_X^* W \Psi_Y - \lambda \Psi_X^* W \Psi_X)^{-1} \Psi_X^* W \Psi_X \boldsymbol{a} = \boldsymbol{a}^* Q S (S - \lambda T)^{-1} T Z^* \boldsymbol{a}.$$



FIGURE 15 Left: Physical setup of the double pendulum. Right: Computed spectral measures for four observables g. The constants C_i are normalization factors so that ||g|| = 1.

Now, after computing the generalized Schur decomposition costing $\mathcal{O}(N_K^2)$ operations, each evaluation requires solving $N_K \times N_K$ upper-triangular linear systems in $\mathcal{O}(N_K^2)$ operations. Additional computational savings can be realized if one is willing to do each evaluation at $\theta_1, \dots, \theta_P$ in parallel. We summarize the evaluation scheme in Algorithm 4.

5.4 | Numerical example: The double pendulum

The double pendulum is shown in Figure 15 (left) and is a physical system that exhibits rich dynamic behavior. It is chaotic and, after a suitable scaling of physical constants, governed by the following system of ordinary differential equations [82]:

$$\begin{split} \dot{\theta}_1 &= \frac{2p_1 - 3p_2\cos(\theta_1 - \theta_2)}{16 - 9\cos^2(\theta_1 - \theta_2)}, & \dot{\theta}_2 &= \frac{8p_2 - 3p_1\cos(\theta_1 - \theta_2)}{16 - 9\cos^2(\theta_1 - \theta_2)}, \\ \dot{p}_1 &= -3\Big(\dot{\theta}_1\dot{\theta}_2\sin(\theta_1 - \theta_2) + \sin(\theta_1)\Big), \quad \dot{p}_2 &= -3\Big(-\dot{\theta}_1\dot{\theta}_2\sin(\theta_1 - \theta_2) + \frac{1}{3}\sin(\theta_2)\Big), \end{split}$$

where $p_1 = 8\dot{\theta}_1 + 3\dot{\theta}_2 \cos(\theta_1 - \theta_2)$ and $p_2 = 2\dot{\theta}_2 + 3\dot{\theta}_1 \cos(\theta_1 - \theta_2)$. The state is $\mathbf{x} = (\theta_1, \theta_2, p_1, p_2)$ in $\Omega = [-\pi, \pi]_{\text{per}}^2 \times \mathbb{R}^2$. We consider the Koopman operator corresponding to discrete time-steps with $\Delta_t = 1$. The system is Hamiltonian; hence, \mathcal{K} is unitary on $L^2(\Omega, \omega)$ with the usual Lebesgue measure.

To compute the resolvent, we use Algorithm 4 with a dictionary consisting of tensor-products of the Fourier basis in x_1 and x_2 , and Hermite functions in x_3 and x_4 , together with an ordering corresponding to a hyperbolic cross approximation. We simulate the dynamics using the ode45 command in MATLAB for a single time-step on an equispaced grid and use the trapezoidal rule.

The chaotic nature of the dynamical system is not a problem as we only collect snapshot data from trajectories with one time step. Following the procedure outlined in Section 5.3, the maximum value of M used in this example is 1,562,500, and the maximum value of N_K used is 2528.

Figure 15 (right) shows spectral measures for four observables computed using the 6th order kernel in Table 1 with $\epsilon = 0.1$. The results are verified against the 6th order kernel with smaller ϵ and a larger N_K . We also show the results for the Poisson kernel (m = 1), which produces an oversmoothed estimate of ν_g . The Fourier observables (top row) display a pronounced peak at $\theta = 0$ and relatively flat spectral measures away from zero. In contrast, the Hermite observables (bottom row) display no peak at $\theta = 0$ and a broader spectral measure away from zero. These observables have a richer interplay with the dynamics, and their spectral measures are more spread out over the spectrum of the Koopman operator.

6 | HIGH-DIMENSIONAL DYNAMICAL SYSTEMS

In the numerical examples of Sections 4 and 5, the dimension of the state-space of (1.1) is modest. However, it can be very difficult to explicitly store a dictionary for larger dimensions that arise in applications such as fluid dynamics, molecular dynamics, and climate forecasting. DMD is a popular approach to studying Koopman operators associated with high-dimensional dynamics. It can yield accurate results for periodic or quasiperiodic systems but can often not adequately capture relevant nonlinear dynamics [6, 21, 138] as it implicitly seeks to fit linear dynamics. Moreover, a rigorous connection with Galerkin approximations of Koopman operators only sometimes holds [132].

When *d* is large, we are naturally in the setting of $N_K \gg M$ as a sufficiently rich dictionary must be selected. In other words, the dimension of the dictionary subspace is usually much larger than the snapshot data when *d* is large. How can we approximate spectral properties of \mathcal{K} when *d* is large and $N_K \gg M$? This section extends our algorithms to high dimensions and provides applications to molecular dynamics and the study of noise leakage of turbulent flow past a cascade of airfoils.

6.1 | Adapting our algorithms to use data-driven dictionaries

Our main idea to tackle large *d* is to construct a data-driven dictionary using the kernel trick. We use the error control aspect of ResDMD to verify the learned dictionary aposteri.

6.1.1 | Kernelized EDMD

A naive construction of the matrix $\Psi_X^* W \Psi_Y$ in Section 4.1.1 requires $\mathcal{O}(N_K^2 M)$ operations, which becomes impractical when N_K is large. Kernelized EDMD [139] aims to make EDMD practical for large N_K . The idea is to compute a much smaller matrix \tilde{K}_{EDMD} that has a subset of the same eigenvalues as K_{EDMD} .

Proposition 6.1 Proposition 1 of [139] with additional weight matrix. Let $\sqrt{W}\Psi_X = U\Sigma V^*$ be a SVD, where $U \in \mathbb{C}^{M \times M}$ is a unitary matrix, $\Sigma \in \mathbb{C}^{M \times M}$ is a diagonal matrix with nonincreasing

nonnegative entries, and $V \in \mathbb{C}^{N_K \times M}$ is an isometry. Define the $M \times M$ matrix

$$\widetilde{K}_{\mathrm{EDMD}} = (\Sigma^{\dagger} U^{*})(\sqrt{W} \Psi_{Y} \Psi_{X}^{*} \sqrt{W})(U\Sigma^{\dagger}),$$

where '†' is the pseudoinverse. Then, for some $\lambda \neq 0$ and $\tilde{v} \in \mathbb{C}^M$, (λ, \tilde{v}) is an eigenvalue-eigenvector pair of \tilde{K}_{EDMD} if and only if $(\lambda, V\tilde{v})$ is an eigenvalue-eigenvector pair of K_{EDMD} .

Suitable matrices U and Σ in Proposition 6.1 can be recovered from the eigenvalue decomposition $\sqrt{W}\Psi_X\Psi_X^*\sqrt{W} = U\Sigma^2 U^*$. Moreover, both matrices $\sqrt{W}\Psi_X\Psi_X^*\sqrt{W} \in \mathbb{C}^{M \times M}$ and $\sqrt{W}\Psi_Y\Psi_X^*\sqrt{W} \in \mathbb{C}^{M \times M}$ can be computed using inner products:

$$[\sqrt{W}\Psi_X\Psi_X^*\sqrt{W}]_{jk} = \sqrt{w_j}\Psi(\mathbf{x}^{(j)})\Psi(\mathbf{x}^{(k)})^*\sqrt{w_k}, \quad [\sqrt{W}\Psi_Y\Psi_X^*\sqrt{W}]_{jk} = \sqrt{w_j}\Psi(\mathbf{y}^{(j)})\Psi(\mathbf{x}^{(k)})^*\sqrt{w_k},$$
(6.1)

where we recall that $\Psi(\mathbf{x})$ is a row vector of the dictionary evaluated at $\mathbf{x}^{.24}$ Kernelized EDMD applies the kernel trick to compute the inner products in (6.1) in an implicitly defined reproducing Hilbert space \mathcal{H} with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ [112]. A positive-definite kernel function S: $\Omega \times \Omega \to \mathbb{R}$ induces a feature map $\varphi : \mathbb{R}^d \to \mathcal{H}$ so that $\langle \varphi(\mathbf{x}), \varphi(\mathbf{x}') \rangle_{\mathcal{H}} = S(\mathbf{x}, \mathbf{x}')$. This leads to a choice of dictionary (or reweighted feature map) $\Psi(\mathbf{x})$ so that $\Psi(\mathbf{x})\Psi(\mathbf{x}')^* = \langle \varphi(\mathbf{x}), \varphi(\mathbf{x}') \rangle_{\mathcal{H}} =$ $S(\mathbf{x}, \mathbf{x}')$. Often S can be evaluated in $\mathcal{O}(d)$ operations so that the matrices $\sqrt{W}\Psi_X\Psi_X^*\sqrt{W}$ and $\sqrt{W}\Psi_Y\Psi_X^*\sqrt{W}$ can be computed in $\mathcal{O}(dM^2)$ operations. $\widetilde{K}_{\text{EDMD}}$ can thus be constructed in $\mathcal{O}(dM^2)$ operations, a considerable saving, with a significant reduction in memory consumption.

As a naive first attempt at extending our algorithms this way, one could consider Algorithm 2 with kernelized EDMD. The residual for a vector $\tilde{v} \in \mathbb{C}^M$ becomes

$$\frac{\tilde{\upsilon}^*[(\Psi_Y V)^* W(\Psi_Y V) - \lambda(\Psi_Y V)^* W(\Psi_X V) - \lambda(\Psi_X V)^* W(\Psi_Y V) + |\lambda|^2 (\Psi_X V)^* W(\Psi_X V)]\tilde{\upsilon}}{\tilde{\upsilon}^*(\Psi_X V)^* W(\Psi_X V)\tilde{\upsilon}^*}.$$
 (6.2)

Unfortunately, the following simple proposition shows that the residual in (6.2) always vanishes. This should be interpreted as *over-fitting of the snapshot data* when $N_K \ge M$.

Proposition 6.2. Suppose that $N_K \ge M$ and $\sqrt{W}\Psi_X$ has rank M (independent rows). For any eigenvalue-eigenvector pair λ and $\tilde{v} \in \mathbb{C}^M$ of \tilde{K}_{EDMD} , the residual in (6.2) vanishes.

Proof. The numerator in (6.2) is the square of the ℓ^2 -vector norm of $\sqrt{W}(\Psi_Y - \lambda \Psi_X)V\tilde{v}$. From the decomposition $\sqrt{W}\Psi_X = U\Sigma V^*$, we see that

$$\sqrt{W}(\Psi_Y - \lambda \Psi_X) V \tilde{v} = \sqrt{W} \Psi_Y V \tilde{v} - \lambda U \Sigma \tilde{v} = \sqrt{W} \Psi_Y \Psi_X^* \sqrt{W} U \Sigma^{\dagger} \tilde{v} - \lambda U \Sigma \tilde{v}$$

Since Σ is invertible, $\sqrt{W}\Psi_{Y}\Psi_{X}^{*}\sqrt{W}U\Sigma^{\dagger} = U\Sigma\widetilde{K}_{\text{EDMD}}$. Hence, $\sqrt{W}(\Psi_{Y} - \lambda\Psi_{X})V\tilde{v} = U\Sigma(\widetilde{K}_{\text{EDMD}} - \lambda)\tilde{v} = 0$.

In other words, the restriction $N_K \ge M$ prevents the large data convergence $(M \to \infty)$ of Algorithm 2 for fixed N_K . Moreover, the residual in (6.2) may not fit into the $L^2(\Omega, \omega)$ Galerkin

²⁴ This is the transpose of the convention in [139].

framework because of the implicit reproducing Hilbert space. A way to overcome these issues is to consider two sets of snapshot data (see Section 6.1.2).

6.1.2 | Kernelized versions of the algorithms

Kernelized EDMD computes the eigenvalues, eigenvectors, and corresponding Koopman modes of $\widetilde{K}_{\text{EDMD}}$. The required operations to compute these quantities are inner products and hence can be computed by evaluating S. In particular, the *k*th eigenfunction ψ_k corresponding to eigenvector $\widetilde{v}_k \in \mathbb{C}^M$ of $\widetilde{K}_{\text{EDMD}}$ can be evaluated as [139, Remark 8]

$$\psi_k(\mathbf{x}) = \left[\mathcal{S}(\mathbf{x}, \mathbf{x}^{(1)}) \quad \mathcal{S}(\mathbf{x}, \mathbf{x}^{(2)}) \quad \cdots \quad \mathcal{S}(\mathbf{x}, \mathbf{x}^{(M)}) \right] (U\Sigma^{\dagger}) \tilde{v}_k, \qquad k = 1, \dots, M.$$
(6.3)

This leads to the simple kernelization of our algorithms summarized in Algorithm 5. The idea is a two-step process: (1) We use kernelized EDMD with a data set $\{\boldsymbol{x}^{(j)}, \boldsymbol{y}^{(j)}\}_{j=1}^{M'}$ to compute a dictionary corresponding to the dominant eigenfunctions of \tilde{K}_{EDMD} and (2) We apply our algorithms using this dictionary on another set $\{\hat{\boldsymbol{x}}^{(j)}, \hat{\boldsymbol{y}}^{(j)}\}_{j=1}^{M''}$ of snapshot data. Of course, there are many algorithmic details regarding the choice of kernel, the second set of snapshot data, and how to a posteriori check that the learned dictionary is suitable. We describe these details now:

• **Choice of kernel**: The choice of kernel *S* determines the dictionary, and the best choice depends on the application. In the following experiments, we use the Gaussian radial basis function kernel

$$S(\boldsymbol{x}, \boldsymbol{y}) = \exp\left(-\gamma \|\boldsymbol{x} - \boldsymbol{y}\|^2\right). \tag{6.4}$$

We select γ as the squared reciprocal of the average ℓ^2 -norm of the snapshot data after it is shifted to mean zero.

- **Regularization**: The matrix $\sqrt{W}\Psi_X\Psi_X^*\sqrt{W}$ can be ill-conditioned [10], in which case we can regularize and consider $\sqrt{W}\Psi_X\Psi_X^*\sqrt{W} + \eta \|\sqrt{W}\Psi_X\Psi_X^*\sqrt{W}\|$ for small η or consider a more stable representation of the range of the feature map. This is standard practice in DMD. However, we found that it was often not needed for Algorithm 5, due to steps 3 and 4.
- Acquiring a second set of snapshot data: Given snapshot data with random and independent initial conditions, one can split the snapshot data up into two parts $\{\mathbf{x}^{(j)}, \mathbf{y}^{(j)}\}_{j=1}^{M'}$ and $\{\hat{\mathbf{x}}^{(j)}, \hat{\mathbf{y}}^{(j)}\}_{j=1}^{M''}$, where M = M' + M''. For initial conditions distributed according to a quadrature rule, if one already has access to M' snapshots, one must typically go back to the original dynamical system and request M'' further snapshots. In both cases, we denote the total amount of snapshot data as M = M' + M''. The first set of M' snapshots is used in the kernelized approach to refine a large initial dictionary of size N'_K (implicitly defined by the chosen kernel) to a smaller data-driven one of size N'_K , while the second set of M'' snapshots are used to apply our ResDMD algorithms. We recommend that $M'' \ge M'$.
- Selecting the second dictionary: We start with a large dictionary of size N'_K and apply a kernelized approach to construct a second dictionary of size N''_K . To construct the second dictionary, we compute the dominant N''_K eigenvectors of \tilde{K}_{EDMD} that hopefully capture the essential dynamics. In this paper, we compute a complete eigendecomposition of \tilde{K}_{EDMD} ; however, for problems with extremely large M', iterative methods are recommended to compute these dominant eigenvectors.



FIGURE 16 Left: Structure of ADK, which has three domains: CORE (green), LID (yellow) and NMP (red). Middle and right: Spectral measures with respect to the dihedral angles of the selected residues.

- **Rigorous results**: It is well-known that the eigenvalues computed by kernel EDMD may suffer from spectral pollution. In our setting, we do not directly use kernel EDMD to compute eigenvalues of \mathcal{K} . Instead, we only use kernel EDMD to select a good dictionary of size $N''_{\mathcal{K}}$, after which our rigorous algorithms can be used (see Algorithms 2 to 4, with convergence as $M'' \to \infty$).
- A posteriori check of the second dictionary: We use $\{\hat{x}^{(j)}, \hat{y}^{(j)}\}_{j=1}^{M''}$ to check the quality of the constructed second dictionary. By studying the residuals and using the error control in Theorem B.1, we can tell a posteriori whether the second dictionary is satisfactory and N_K'' is sufficiently large.

6.2 | Example: MD simulation of the Adenylate Kinase enzyme (d = 20046)

Molecular dynamics (MD) analyzes the movement of atoms and molecules by numerically solving Newton's equations of motion for a system of interacting particles. Energies and forces between particles are typically computed using potentials. MD is arguably one of the most robust approaches for simulating macromolecular dynamics, primarily due to the availability of full atomistic detail [43]. Recently, DMD-type and Koopman techniques are impacting MD [74, 97, 113, 114]. For example, [73] applies kernel EDMD to the positions of the carbon atoms in *n*-butane (d = 12) and shows that the EDMD eigenfunctions parameterize a dihedral angle that controls crucial dynamics.

Here, we study trajectory data from the dynamics of Adenylate Kinase (ADK), which is an enzyme (see Figure 16) that catalyzes important phosphate reactions in cellular biology [56]. ADK is a common benchmark enzyme in MD [116] and consists of 3341 atoms split into 214 residues (specific monomers that can be thought of as parts). The trajectory data comes from an all-atom equilibrium simulation for 1.004×10^{-6} s, with a so-called CHARMM force field that is produced on PSC Anton [118] and publicly available [8]. The data consists of a single trajectory of the positions of all atoms as ADK moves. To make the system Hamiltonian, we append the data with approximations of the velocities computed using centered finite differences. This leads to $d = 6 \times 3341 = 20046$. We vertically stack the data (as discussed in Section 4) and sample the trajectory data every 240×10^{-12} s so that M = 4184.

To apply the kernelized version of Algorithm 4, we subselect M' = 2000 initial conditions from the trajectory data and used the Gaussian radial basis function kernel in (6.4). We select $N''_K =$ 1000 EDMD eigenfunctions and append the dictionary with the four observables of interest that are discussed below. Accuracy of the corresponding matrices in (4.7) is verified by comparing to smaller M'' and also by computing pseudospectra with Algorithm 3.

ADK has three parts of its molecule called CORE, LID, and NMP (see Figure 16 (left)). The LID and NMP domains move around the stable CORE. We select the most mobile residue from the LID and NMP domains by computing root-mean-square-fluctuations. These residues have canonical dihedral angles (ϕ, ψ) defined on the backbone atoms [103] that determine the overall shape of the residue. Figure 16 (middle, right) shows the spectral measures with respect to these dihedral angles (where we have subtracted the mean angle value) for both selected residues. These spectral measures are computed using the sixth-order rational kernel with $\epsilon = 0.1$ (see Table 1). The computed spectral measures are verified with higher order kernels and smaller ϵ , and through comparison with a polynomial kernel S instead of a Gaussian radial basis function kernel. The spectral measures for the angles in the LID residue are much broader than for the NMP residue. This hints at a more complicated dynamical interaction and may have biological consequences. We hope that this example can be a catalyst for the use of spectral measures in MD.

6.3 | Example: Turbulent flow past a cascade of aerofoils (d = 295122)

Some of the most successful applications of DMD are in fluid dynamics [24, 27, 90, 107, 128]. Here, we consider a large-scale wall-resolved turbulent flow past a periodic cascade of aerofoils with a stagger angle 56.9° and a one-sided tip gap. The setup is motivated by the need to reduce noise sources from flying objects [100]. We use a high-fidelity simulation that solves the fully nonlinear Navier–Stokes equations [76] with Reynolds number 3.88×10^5 and Mach number 0.07. The data consists of a 2D slice of the pressure field, measured at 295,122 points, for trajectories of length 798 that are sampled every 2×10^{-5} s. We stack the data in the form of (4.1) so that M = 797.

To apply the kernelized version of Algorithm 2, we subselect M' = 350 initial conditions from the trajectory data and used the Gaussian radial basis function kernel in (6.4). We select $N''_K =$ 250 EDMD eigenfunctions as our dictionary. Accuracy of the corresponding matrices in (4.7) is verified by comparing to smaller M'' and also by computing pseudospectra with Algorithm 3.

Figure 17 shows the computed Koopman modes for a range of representative frequencies. We also show the corresponding Koopman modes computed using DMD. In the case of DMD and ResDMD, each mode is only defined up to a global phase. Moreover, the color maps for each λ differ because of the pressure modulus variations with frequency. For the first row of Figure 17, ResDMD shows stronger acoustic waves between the cascades. Detecting these vibrations is important as they can damage turbines in engines [99]. For the second and third row of Figure 17, ResDMD shows larger scale turbulent fluctuations past the trailing edge corresponding to the figure's right. This can be crucial for understanding acoustic interactions with nearby structures such as subsequent blade rows [140]. DMD shows a clear acoustic source for the third row just above the wing. This source is less distinct in the case of ResDMD because of nonlinear interference. The residuals for ResDMD are also shown and are small, particularly given the enormous state-space dimension. This example demonstrates two benefits of the kernelized version of ResDMD (see Algorithms 2 and 3) compared with DMD: (1) ResDMD can capture the nonlinear dynamics, and (2) It computes residuals as well, thus providing an accuracy certificate.



FIGURE 17 Left column: Modes computed by DMD. Right column: Modes computed by ResDMD with computed residuals. Each row corresponds to the different values of λ that correspond to certain physical frequencies of noise pollution.

ACKNOWLEDGMENTS

The work of the first author was supported by a Research Fellowship at Trinity College, Cambridge; a Fondation Sciences Mathématiques de Paris Postdoctoral Fellowship at École Normale Supérieure; and an LMS Cecil King Travel Scholarship. The work of the second author was supported by the National Science Foundation grants DMS-1818757, DMS-1952757 and DMS-2045646. The authors are grateful to John Guckenheimer for stimulating discussions during the completion of this work, Andrew Horning and Jennifer Zvonek for reading a draft of this manuscript, Dimitris Giannakis and Igor Mezić for feedback, Regis Koch, Marléne Sanjosé, and Stephane Moreau for sharing their flow data for the example in Section 6.3, and to the anonymous referees whose careful comments helped improve the clarity of the manuscript.

REFERENCES

- 1. B. Adcock and A. C. Hansen, Stable reconstructions in Hilbert spaces and the resolution of the Gibbs phenomenon, Appl. Comput. Harm. Anal. **32** (2012), no. 3, 357–388.
- 2. H. Arbabi and I. Mezic, Ergodic theory, dynamic mode decomposition, and computation of spectral properties of the Koopman operator, SIAM J. Appl. Dyn. Syst. **16** (2017), no. 4, 2096–2126.
- 3. H. Arbabi and I. Mezić, *Study of dynamics in post-transient flows using Koopman mode decomposition*, Phys. Rev. Fluids **2** (2017), no. 12, 124402.
- 4. V. I. Arnold, Mathematical Methods of Classical Mechanics, Springer-Verlag, New York, 1989.
- 5. V. I. Arnold and A. Avez, Ergodic Problems of Classical Mechanics, vol. 9, Benjamin, New York, 1968.
- 6. P. J. Baddoo, B. Herrmann, B. J. McKeon, and S. L. Brunton, *Kernel learning for robust dynamic mode decomposition: Linear and nonlinear disambiguation optimization, Proc. R. Soc. A*, 478 (2021), no. 2260.
- 7. J. Basley, L. Pastur, F. Lusseyran, T. M. Faure, and N. Delprat, *Experimental investigation of global structures in an incompressible cavity flow using time-resolved PIV*, Exper. Fluids **50** (2011), no. 4, 905–918.
- O. Beckstein, S. Fan, R. J. Gowers, M. Matta, and L. Wang, *AdK equilibrium dataset*, https://www.mdanalysis. org/MDAnalysisData/adk_equilibrium.html.
- 9. G. Beer, *Topologies on Closed and Closed Convex Sets*, vol. 268, Springer Science & Business Media, Berlin/Heidelberg, 1993.
- M. Belkin, Approximation beats concentration? An approximation view on inference with smooth radial kernels, Conference On Learning Theory, PMLR, pp. 1348–1361.
- J. Ben-Artzi, M. J. Colbrook, A. C. Hansen, O. Nevanlinna, and M. Seidel, *Computing spectra On the solvability complexity index hierarchy and towers of algorithms*, arXiv preprint arXiv:1508.03280v4 (2020).
- P. Benner, S. Gugercin, and K. Willcox, A survey of projection-based model reduction methods for parametric dynamical systems, SIAM Rev. 57 (2015), no. 4, 483–531.
- E. Berger, M. Sastuba, D. Vogt, B. Jung, and H. Ben Amor, *Estimation of perturbations in robotic behavior using dynamic mode decomposition*, Adv. Robot. 29 (2015), no. 5, 331–343.
- 14. P. Billingsley, Convergence of Probability Measures, 2nd edn., John Wiley & Sons, Hoboken, NJ, 1999.
- 15. G. D. Birkhoff, Proof of the ergodic theorem, Proc. Natl. Acad. Sci. 17 (1931), no. 12, 656-660.
- A. Böttcher and B. Silbermann, The finite section method for Toeplitz operators on the quarter-plane with piecewise continuous symbols, Math. Nachrichten 110 (1983), no. 1, 279–291.
- 17. D. Bruder, B. Gillespie, C. D. Remy, and R. Vasudevan, *Modeling and control of soft robots using the Koopman operator and model predictive control*, arXiv preprint arXiv:1902.02827 (2019).
- B. W. Brunton, L. A. Johnson, J. G. Ojemann, and J. N. Kutz, Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition, J. Neurosci. Meth. 258 (2016), 1–15.
- 19. S. L. Brunton and J. N. Kutz, *Data-driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*, Cambridge University Press, Cambridge, 2019.
- S. L. Brunton, J. L. Proctor, and J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, Proc. Natl. Acad. Sci. 113 (2016), no. 15, 3932–3937.
- 21. S. L. Brunton, B. W. Brunton, J. L. Proctor, and J. N. Kutz, Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control, PloS one **11** (2016), no. 2, e0150171.
- S. L. Brunton, B. W. Brunton, J. L. Proctor, E. Kaiser, and J. N. Kutz, *Chaos as an intermittently forced linear system*, Nat. Commun. 8 (2017), no. 1, 1–9.

- S. L. Brunton, M. Budišić, E. Kaiser, and J. N. Kutz, Modern Koopman theory for dynamical systems, arXiv preprint arXiv:2102.12086 (2021).
- 24. M. Budišić, R. Mohr, and I. Mezić, Applied Koopmanism, Chaos 22 (2012), no. 4, 047510.
- 25. D. Burov, D. Giannakis, K. Manohar, and A. Stuart, *Kernel analog forecasting: multiscale test problems*, Multiscale Model. Simul. **19** (2021), no. 2, 1011–1040.
- 26. R. E. Caflisch, Monte Carlo and quasi-Monte Carlo methods, Acta Numer. 7 (1998), 1-49.
- 27. K. K. Chen, J. H. Tu, and C. W. Rowley, Variants of dynamic mode decomposition: boundary condition, Koopman, and Fourier analyses, J. Nonl. Sci. 22 (2012), no. 6, 887–915.
- M. Colbrook, A. Horning, and A. Townsend, *Computing spectral measures of self-adjoint operators*, SIAM Rev. 63 (2021), no. 3, 489–524.
- 29. M. J. Colbrook, *The Foundations of Infinite-Dimensional Spectral Computations*, Ph.D. thesis, University of Cambridge, 2020.
- 30. M. J. Colbrook, Computing spectral measures and spectral types, Commun. Math. Phys. 384 (2021), 433-501.
- 31. M. J. Colbrook, On the computation of geometric features of spectra of linear operators on Hilbert spaces, Found. Comut. Math. (2022), 1–82.
- 32. M. J. Colbrook, *The mpEDMD algorithm for data-driven computations of measure-preserving dynamical systems*, SIAM J. Numer. Anal. **61** (2023), no. 3, 1585–1608.
- 33. M. J. Colbrook, L. J. Ayton, and M. Szőke, *Residual dynamic mode decomposition: robust and verified Koopmanism*, J. Fluid Mech. **955** (2023), A21.
- 34. M. J. Colbrook and A. C. Hansen, *On the infinite-dimensional QR algorithm*, Numer. Math. **143** (2019), no. 1, 17–83.
- 35. M. J. Colbrook and A. C. Hansen, *The foundations of spectral computations via the solvability complexity index hierarchy*, J. Eur. Math. Soc. (2022).
- M. J. Colbrook, B. Roman, and A. C. Hansen, *How to compute spectra with error control*, Phys. Rev. Lett. 122 (2019), no. 25, 250201.
- 37. M. J. Colbrook and A. Townsend, Avoiding discretization issues for nonlinear eigenvalue problems, arXiv preprint arXiv:2305.01691 (2023).
- S. Das and D. Giannakis, Delay-coordinate maps and the spectra of Koopman operators, J. Stat. Phys. 175 (2019), no. 6, 1107–1145.
- S. Das and D. Giannakis, Koopman spectra in reproducing kernel Hilbert spaces, Appl. Comput. Harmon. Anal. 49 (2020), no. 2, 573–607.
- 40. S. Das, D. Giannakis, and J. Slawinska, *Reproducing kernel Hilbert space compactification of unitary evolution groups*, Appl. Comput. Harm. Anal. **54** (2021), 75–136.
- 41. M. Dellnitz and O. Junge, *On the approximation of complicated dynamical behavior*, SIAM J. Numer. Anal. **36** (1999), no. 2, 491–515.
- 42. Z. Drmac, I. Mezic, and R. Mohr, *Data driven modal decompositions: analysis and enhancements*, SIAM J. Sci. Comput. **40** (2018), no. 4, A2253–A2285.
- 43. R. O. Dror, R. M. Dirks, J. Grossman, H. Xu, and D. E. Shaw, *Biomolecular simulation: a computational microscope for molecular biology*, Ann. Rev. Biophys. **41** (2012), 429–452.
- B. A. Dubrovin, A. T. Fomenko, and S. P. Novikov, Modern Geometry Methods and Applications Part I. The Geometry of Surfaces, Transformation Groups, and Fields, vol. 104, Springer Science & Business Media, Berlin/Heidelberg, 21984.
- 45. D. D ung, V. Temlyakov, and T. Ullrich, Hyperbolic Cross Approximation, Springer, Cham, 2018.
- J.-P. Eckmann and D. Ruelle, *Ergodic theory of chaos and strange attractors*, Rev. Mod. Phys. 57 (1985), no. 3, 617.
- 47. E. S. Epstein, Stochastic dynamic prediction, Tellus 21 (1969), no. 6, 739–759.
- 48. L. Fejér, Sur les fonctions bornées et intégrables, C.R. Acad. Sci. Paris (1900), 984–987.
- 49. G. Froyland, *Computer-assisted bounds for the rate of decay of correlations*, Commun. Math. Phys. **189** (1997), 237–257.
- 50. G. Froyland, On Ulam approximation of the isolated spectrum and eigenfunctions of hyperbolic maps, Discrete Contin. Dyn. Syst. Ser. A **17** (2007), no. 3, 203–221.
- 51. M. Ghil, M. Allen, M. Dettinger, K. Ide, D. Kondrashov, M. Mann, A. W. Robertson, A. Saunders, Y. Tian, F. Varadi, et al., *Advanced spectral methods for climatic time series*, Rev. Geophys. **40** (2002), no. 1, 3-1–3-41.

- 52. D. Giannakis, *Data-driven spectral decomposition and forecasting of ergodic dynamical systems*, Appl. Comput. Harm. Anal. **47** (2019), no. 2, 338–396.
- D. Giannakis, *Delay-coordinate maps, coherence, and approximate spectra of evolution operators*, Res. Math. Sci. 8 (2021), no. 1, 1–33.
- D. Giannakis, A. Kolchinskaya, D. Krasnov, and J. Schumacher, Koopman analysis of the long-term evolution in a turbulent convection cell, J. Fluid Mech. 847 (2018), 735–767.
- 55. D. Giannakis, A. Henriksen, J. A. Tropp, and R. Ward, *Learning to forecast dynamical systems from streaming data*, arXiv preprint arXiv:2109.09703 (2021).
- R. N. Goldberg, Y. B. Tewari, and T. N. Bhat, *Thermodynamics of enzyme-catalyzed reactions a database for quantitative biochemistry*, Bioinfo. 20 (2004), no. 16, 2874–2877.
- 57. D. Gottlieb and C.-W. Shu, On the Gibbs phenomenon and its resolution, SIAM Rev. 39 (1997), no. 4, 644-668.
- M. Guo and J. S. Hesthaven, Data-driven reduced order modeling for time-dependent problems, Comput. Meth. Appl. Mech. Eng. 345 (2019), 75–99.
- 59. P. R. Halmos, Lectures on Ergodic Theory, Courier Dover Publications, Mineola, NY, 2017.
- A. Hansen, On the solvability complexity index, the n-pseudospectrum and approximations of spectra of operators, J. Am. Math. Soc. 24 (2011), no. 1, 81–124.
- 61. T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning, Springer-Verlag, Berlin, 2009.
- 62. J. Hendrickx and A. Olshevsky, *Matrix p-norms are NP-hard to approximate if* $p \neq 1, 2, \infty$, SIAM J. Matrix Anal. Appl. **31** (2010), no. 5, 2802.
- 63. J. S. Hesthaven, Numerical Methods for Conservation Laws: From Analysis to Algorithms, SIAM, Philadelphia, 2017.
- 64. A. J. Hey, S. Tansley, K. M. Tolle, et al., *The Fourth Paradigm: Data-intensive Scientific Discovery*, vol. 1, Microsoft Research, Redmond, WA, 2009.
- R. C. Hilborn, et al., Chaos and Nonlinear Dynamics: An Introduction for Scientists and Engineers, 2nd edn., Oxford University Press, Oxford, 2004.
- 66. O. Junge and P. Koltai, Discretization of the Frobenius–Perron operator using a sparse haar tensor basis: the sparse Ulam method, SIAM J. Numer. Anal. 47 (2009), no. 5, 3464–3485.
- 67. A. G. Kachurovskii, The rate of convergence in ergodic theorems, Russian Math. Sur. 51 (1996), no. 4, 653–703.
- E. Kaiser, J. N. Kutz, and S. L. Brunton, Data-driven discovery of Koopman eigenfunctions for control, arXiv preprint arXiv:1707.01146 (2017).
- 69. R. E. Kalman, *Mathematical description of linear dynamical systems*, J. Soc. Ind. Appl. Math. **1** (1963), no. 2, 152–192.
- 70. M. Kamb, E. Kaiser, S. L. Brunton, and J. N. Kutz, *Time-delay observables for Koopman: theory and applications*, SIAM J. Appl. Dyn. Syst. **19** (2020), no. 2, 886–917.
- 71. Y. Katznelson, An Introduction to Harmonic Analysis, Cambridge University Press, Cambridge, 2004.
- 72. S. Klus, P. Koltai, and C. Schütte, On the numerical approximation of the Perron-Frobenius and Koopman operator, J. Comput. Dyn. 3 (2016), no. 1, 51–79.
- 73. S. Klus, I. Schuster, and K. Muandet, *Eigendecompositions of transfer operators in reproducing kernel Hilbert spaces*, J. Nonlin. Sci. **30** (2020), no. 1, 283–315.
- 74. S. Klus, F. Nüske, P. Koltai, H. Wu, I. Kevrekidis, C. Schütte, and F. Noé, *Data-driven model reduction and transfer operator approximation*, J. Nonlin. Sci. **28** (2018), no. 3, 985–1010.
- S. Klus, F. Nüske, S. Peitz, J.-H. Niemann, C. Clementi, and C. Schütte, *Data-driven approximation of the Koopman generator: Model reduction, system identification, and control*, Phys. D: Nonl. Pheno. 406 (2020), 132416.
- 76. R. Koch, M. Sanjosé, and S. Moreau, *Large-eddy simulation of a linear compressor cascade with tip gap: aerodynamic and acoustic analysis*, AIAA Aviation 2021 Forum, pp. 2312.
- 77. B. O. Koopman, *Hamiltonian systems and transformation in Hilbert space*, Proc. Natl. Acad. Sci. USA **17** (1931), no. 5, 315.
- B. O. Koopman and J. von Neumann, *Dynamical systems of continuous spectra*, Proc. Natl. Acad. Sci. USA 18 (1932), no. 3, 255.
- M. Korda and I. Mezić, On convergence of extended dynamic mode decomposition to the Koopman operator, J. Nonlin. Sci. 28 (2018), no. 2, 687–710.

- M. Korda, M. Putinar, and I. Mezić, Data-driven spectral analysis of the Koopman operator, Appl. Comput. Harm. Anal. 48 (2020), no. 2, 599–629.
- J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor, Dynamic Mode Decomposition: Data-driven Modeling of Complex Systems, SIAM, Philadelphia, 2016.
- 82. R. Levien and S. Tan, Double pendulum: an experiment in chaos, Amer. J. Phys. 61 (1993), no. 11, 1038–1044.
- M. Lewin and É. Séré, Spectral pollution and how to avoid it, Proc. London Math. Soc. 100 (2010), no. 3, 864–900.
- Q. Li, F. Dietrich, E. M. Bollt, and I. G. Kevrekidis, *Extended dynamic mode decomposition with dictionary learning: a data-driven adaptive spectral decomposition of the Koopman operator*, Chaos 27 (2017), no. 10, 103111.
- 85. E. N. Lorenz, Deterministic nonperiodic flow, J. Atmos. Sci. 20 (1963), no. 2, 130-141.
- B. Lusch, J. N. Kutz, and S. L. Brunton, *Deep learning for universal linear embeddings of nonlinear dynamics*, Nat. Commun. 9 (2018), no. 1, 1–10.
- J. Mann and J. N. Kutz, Dynamic mode decomposition for financial trading strategies, Quant. Finance 16 (2016), no. 11, 1643–1655.
- A. Mardt, L. Pasquali, H. Wu, and F. Noé, VAMPnets for deep learning of molecular kinetics, Nat. Commun. 9 (2018), no. 1, 1–11.
- I. Mezić, Spectral properties of dynamical systems, model reduction and decompositions, Nonlin. Dyn. 41 (2005), no. 1, 309–325.
- I. Mezić, Analysis of fluid flows via spectral properties of the Koopman operator, Ann. Rev. Fluid Mech. 45 (2013), 357–378.
- 91. I. Mezic, On numerical approximations of the Koopman operator, arXiv preprint arXiv:2009.05883 (2020).
- I. Mezić, Spectrum of the Koopman operator, spectral expansions in functional spaces, and state-space geometry, J. Nonlin. Sci. 30 (2020), no. 5, 2091–2145.
- I. Mezić, Koopman operator, geometry, and learning of dynamical systems, Not. Amer. Math. Soc. 68 (2021), no. 7, 1087–1105.
- I. Mezić and A. Banaszuk, Comparison of systems with complex behavior, Phys. D: Nonlin. Phen. 197 (2004), no. 1-2, 101–133.
- M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of Machine Learning*, MIT Press, Cambridge, MA, 2018.
- B. S. Nagy, C. Foias, H. Bercovici, and L. Kérchy, *Harmonic Analysis of Operators on Hilbert Space*, Springer Science & Business Media, Berlin, 2010.
- F. Nuske, B. G. Keller, G. Pérez-Hernández, A. S. Mey, and F. Noé, *Variational approach to molecular kinetics*, J. Chem. Theory Comput. **10** (2014), no. 4, 1739–1752.
- S. E. Otto and C. W. Rowley, *Linearly recurrent autoencoder networks for learning dynamics*, SIAM J. Appl. Dyn. Syst. 18 (2019), no. 1, 558–593.
- R. Parker, Acoustic resonances and blade vibration in axial flow compressors, J. Sound Vibr. 92 (1984), no. 4, 529–539.
- N. Peake and A. B. Parry, Modern challenges facing turbomachinery aeroacoustics, Ann. Rev. Fluid Mech. 44 (2012), 227–248.
- 101. H. Poincaré, Les Méthodes Nouvelles de la Mécanique Céleste, vol. 3, Gauthier-Villars, 1899.
- J. L. Proctor and P. A. Eckhoff, Discovering dynamic patterns from infectious disease data using dynamic mode decomposition, Inter. Health 7 (2015), no. 2, 139–145.
- G. Ramachandran, C. Ramakrishnan, and V. Sasisekharan, *Stereochemistry of polypeptide chain configura*tions, J. Mol. Biol. 7 (1963), no. 1, 95–99.
- M. Reed and B. Simon, *Methods of Modern Mathematical Physics*. I, 2nd edn., Academic Press, New York, 1980, pp. xv+400.
- 105. S. Roch and B. Silbermann, C*-algebra techniques in numerical analysis, J. Oper. Theory (1996), 241–280.
- L. J. Rogers, Second memoir on the expansion of certain infinite products, Proc. Lond. Math. Soc. 25 (1893/94), 318–343.
- C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, *Spectral analysis of nonlinear flows*, J. Fluid Mech. 641 (2009), 115–127.

- S. H. Rudy, S. L. Brunton, J. L. Proctor, and J. N. Kutz, *Data-driven discovery of partial differential equations*, Sci. Adv. 3 (2017), no. 4, e1602614.
- B. Saltzman, Finite amplitude free convection as an initial value problem I, J. Atmos. Sci. 19 (1962), no. 4, 329–341.
- P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, J. Fluid Mech. 656 (2010), 5–28.
- M. Schmidt and H. Lipson, *Distilling free-form natural laws from experimental data*, Science 324 (2009), no. 5923, 81–85.
- 112. B. Scholkopf, The kernel trick for distances, Adv. Neur. Info. Proc. Syst. (2001), 301-307.
- 113. C. R. Schwantes and V. S. Pande, *Improvements in Markov state model construction reveal many non-native interactions in the folding of NTL9*, J. Chem. Theory Comput. **9** (2013), no. 4, 2000–2009.
- C. R. Schwantes and V. S. Pande, Modeling molecular kinetics with tICA and the kernel trick, J. Chem. Theory Comput. 11 (2015), no. 2, 600–608.
- 115. M. Seidel, $On(N, \epsilon)$ -pseudospectra of operators on Banach spaces, J. Funct. Anal. **262** (2012), no. 11, 4916–4927.
- S. L. Seyler and O. Beckstein, Sampling large conformational transitions: adenylate kinase as a testing ground, Mole. Simul. 40 (2014), no. 10-11, 855–877.
- 117. E. Shargorodsky, On the level sets of the resolvent norm of a linear operator, Bull. Lond. Math. Soc. **40** (2008), no. 3, 493–504.
- 118. D. E. Shaw, R. O. Dror, J. K. Salmon, J. Grossman, K. M. Mackenzie, J. A. Bank, C. Young, M. M. Deneroff, B. Batson, K. J. Bowers, et al., *Millisecond-scale molecular dynamics simulations on Anton*, Proc. Conf. High Perfor. Comput. Net., Stor. and Anal., pp. 1–11.
- 119. P. C. Shields, The Theory of Bernoulli Shifts, University of Chicago Press, Chicago, 1973.
- 120. B. Simon, CMV matrices: Five years after, J. Comput. Appl. Math. 208 (2007), no. 1, 120–154.
- 121. E. M. Stein and R. Shakarchi, *Real Analysis: Measure Theory, Integration, and Hilbert Spaces*, Princeton University Press, Princeton, NJ, 2009.
- 122. G. W. Stewart and J.-G. Sun, Matrix Perturbation Theory, Academic Press, New York, 1990.
- 123. A. Stuart and A. R. Humphries, *Dynamical Systems and Numerical Analysis*, vol. 2, Cambridge University Press, Cambridge, 1998.
- Y. Susuki, A. Mauroy, and I. Mezic, Koopman resolvent: A Laplace-domain analysis of nonlinear autonomous dynamical systems, SIAM J. Appl. Dyn. Syst. 20 (2021), no. 4, 2013–2036.
- 125. Y. Susuki and I. Mezic, Nonlinear Koopman modes and coherency identification of coupled swing dynamics, IEEE Trans. Power Syst. **26** (2011), no. 4, 1894–1904.
- 126. Y. Susuki, I. Mezić, and T. Hikihara, *Coherent swing instability of power grids*, J. Nonlin. Sci. **21** (2011), no. 3, 403–439.
- 127. E. Tadmor, Filters, mollifiers and the computation of the Gibbs phenomenon, Acta Numer. 16 (2007), 305–378.
- 128. K. Taira, S. L. Brunton, S. T. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, and L. S. Ukeiley, *Modal analysis of fluid flows: an overview*, AIAA J. 55 (2017), no. 12, 4013–4041.
- 129. L. N. Trefethen and M. Embree, Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators, Princeton University Press, Princeton, NJ, 2005.
- L. N. Trefethen and J. Weideman, *The exponentially convergent trapezoidal rule*, SIAM Rev. 56 (2014), no. 3, 385–458.
- L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll, *Hydrodynamic stability without eigenvalues*, Science 261 (1993), no. 5121, 578–584.
- 132. J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz, On dynamic mode decomposition: theory and applications, J. Comput. Dyn. 1 (2014), no. 2, 391–421.
- 133. N. B. Tufillaro, T. Abbott, and J. Reilly, *An Experimental Approach to Nonlinear Dynamics and Chaos*, Addison-Wesley, Redwood City, CA, 1992.
- 134. H. Vandeven, Family of spectral filters for discontinuous problems, J. Sci. Comput. 6 (1991), no. 2, 159–192.
- 135. P. Walters, *An Introduction to Ergodic Theory*, vol. 79, Springer Science & Business Media, Berlin/Heidelberg, 2000.
- 136. E. Wegert, Visual Complex Functions: An Introduction with Phase Portraits, Springer Science & Business Media, Berlin/Heidelberg, 2012.

- 137. P. Welch, The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms, IEEE Trans. Audio Electro. **15** (1967), no. 2, 70–73.
- 138. M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, *A data-driven approximation of the Koopman operator:* extending dynamic mode decomposition, J. Nonlin. Sci. **25** (2015), no. 6, 1307–1346.
- 139. M. O. Williams, C. W. Rowley, and I. G. Kevrekidis, *A kernel-based method for data-driven Koopman spectral analysis*, J. Comput. Dyn. **2** (2015), no. 2, 247.
- B. Woodley and N. Peake, Resonant acoustic frequencies of a tandem cascade. Part 2. Rotating blade rows, J. Fluid Mech. 393 (1999), 241–256.
- 141. E. Yeung, S. Kundu, and N. Hodas, *Learning deep neural network representations for Koopman operators of nonlinear dynamical systems*, 2019 Amer. Contr. Conf., IEEE, pp. 4832–4839.
- 142. G. M. Zaslavsky, Chaos, fractional kinetics, and anomalous transport, Phys. Rep. 371 (2002), no. 6, 461-580.
- 143. Z. Zhao and D. Giannakis, Analog forecasting with dynamics-adapted kernels, Nonlin. 29 (2016), no. 9, 2888.

APPENDIX A: A USEFUL CRITERION FOR mth ORDER KERNELS

The following lemma is used to prove Proposition 3.1 and build mth order rational kernels of the form (5.6). It provides sufficient conditions for a family of integrable functions to be an mth order kernel.

Lemma A.1. Let $\{K_{\epsilon} : \epsilon \in (0,1]\}$ be a family of integrable functions on the periodic interval $[-\pi,\pi]_{per}$ that integrate to 1. Suppose that there exists a constant *C* such that for any integer *n* with $0 < n \le m - 1$,

$$\left| \int_{[-\pi,\pi]_{\text{per}}} K_{\epsilon}(-\theta) e^{in\theta} d\theta - 1 \right| \le C \epsilon^m \log(\epsilon^{-1}), \tag{A.1}$$

and such that

$$|K_{\varepsilon}(\theta)| \le \frac{C\epsilon^m}{(\epsilon + |\theta|)^{m+1}},\tag{A.2}$$

for any $\theta \in [-\pi, \pi]_{\text{per}}$ and $\epsilon \in (0, 1]$. Then $\{K_{\epsilon}\}$ is an *m*th order kernel for $[-\pi, \pi]_{\text{per}}$.

Proof of Lemma A.1. There are three properties that an *m*th order kernel must satisfy (see Definition 3.1). The lemma statement already assumes properties (i) and (iii). To show part (ii) of Definition 3.1, let *n* be an integer and consider θ^n with n < m. We use the convergent power series for the principal branch of the logarithm around 1 to write, for $\theta \in [-\pi/4, \pi/4]$,

$$\theta^n = \sum_{k=1}^{\infty} a_{k,n} (e^{i\theta} - 1)^k = \underbrace{\sum_{k=1}^{m-1} a_{k,n} (e^{i\theta} - 1)^k}_{=p_{n,m}(\theta)} + \theta^m \underbrace{\frac{(e^{i\theta} - 1)^m}{\theta^m} \sum_{k=m}^{\infty} a_{k,n} (e^{i\theta} - 1)^{k-m}}_{=h_{n,m}(\theta)},$$

where $a_{k,n} \in \mathbb{C}$ and the domain of the trigonometric polynomial $p_{n,m}$ is extended to $[-\pi,\pi]_{\text{per}}$. We note that $h_{n,m}(\theta)$ is a bounded continuous function on $[-\pi/4,\pi/4]$. Since $|\int_{[-\pi,\pi]_{\text{per}}} K_{\epsilon}(-\theta)p_{n,m}(\theta)d\theta - p_{n,m}(0)| \leq \epsilon^m \log(\epsilon^{-1})$, we know by (A.2) that

$$\begin{split} \left| \int_{-\pi/4}^{\pi/4} \theta^n K_{\epsilon}(-\theta) \, d\theta \right| &\lesssim \int_0^{\pi/(4\epsilon)} \frac{\epsilon^m \tau^m}{(1+\tau)^{m+1}} \, d\tau + \int_{|\theta| > \pi/4} |K_{\epsilon}(-\theta)| |p_{n,m}(\theta)| \, d\theta + \epsilon^m \log(\epsilon^{-1}) \\ &\lesssim \epsilon^m \log(\epsilon^{-1}) + \int_{|\theta| > \pi/4} \frac{\epsilon^m \|p_{n,m}\|_{L^{\infty}}}{(\epsilon + \pi/4)^{m+1}} \, d\theta \lesssim \epsilon^m \log(\epsilon^{-1}). \end{split}$$

Furthermore, by (A.2) we obtain

$$\left| \int_{[-\pi,-\pi/4)\cup(\pi/4,\pi]} \theta^n K_{\varepsilon}(-\theta) \, d\theta \right| \leq \pi^n (2\pi - \pi/2) \frac{C\epsilon^m}{(\epsilon + \pi/4)^{m+1}} \lesssim \epsilon^m.$$

Property (ii) of Definition 3.1 now follows by the symmetry of (3.10) under the mapping $\theta \mapsto -\theta$.

APPENDIX B: ResDMD FOR COMPUTING SPECTRA AND PSEUDOSPECTRA

In this section, we show how to modify ResDMD in Section 4.2 to rigorously compute the global spectral properties of the Koopman operator \mathcal{K} . Recall that we assume that \mathcal{K} is a closed and densely defined operator. Throughout this section, we make the following assumption about the dictionary $\{\psi_j\}_{j=1}^{N_K}$:

• Assumption: span{ $\psi \in V_{N_{\mathcal{K}}}$: $N_{\mathcal{K}} \in \mathbb{N}$ } forms a core of \mathcal{K} .

This assumption means that the closure of the restriction of \mathcal{K} to span{ $\psi \in V_{N_K} : N_K \in \mathbb{N}$ } is \mathcal{K} . In other words, the assumption allows us to study the infinite-dimensional operator \mathcal{K} , which acts on a possibly strict subdomain $\mathcal{D}(\mathcal{K}) \subset L^2(\Omega, \omega)$, by considering its action on the observables in V_{N_K} . This is vital for capturing the full operator \mathcal{K} , and hence, equivalently, the full dynamics. To simplify our arguments, we also assume that ψ_j is independent of N_K for a fixed j, and that $V_{N_K} \subset V_{N_K+1}$. These last two assumptions can be dropped with minor modifications to our arguments.

We begin by re-writing the quantity $\|(\mathcal{K} - \lambda)^{-1}\|^{-1}$ as a minimal residual. Given a linear operator *T* with domain $\mathcal{D}(T)$, we define the injection modulus $\mathcal{J}(T) = \inf\{\|T\psi\| : \|\psi\| = 1, \psi \in \mathcal{D}(T)\}$. It is well-known that

$$\|(\mathcal{K}-\lambda)^{-1}\|^{-1} = \min\{\mathcal{J}(\mathcal{K}-\lambda), \mathcal{J}(\mathcal{K}^*-\overline{\lambda})\} \le \mathcal{J}(\mathcal{K}-\lambda) = \inf_{g \in \mathcal{D}(\mathcal{K}), g \neq 0} \frac{\|(\mathcal{K}-\lambda)g\|}{\|g\|}$$

with $\mathcal{J}(\mathcal{K} - \lambda) = \mathcal{J}(\mathcal{K}^* - \overline{\lambda})$ if $\lambda \notin \sigma(\mathcal{K})$. Therefore, $\sigma(\mathcal{K}) = \{\lambda : \min\{\mathcal{J}(\mathcal{K} - \lambda), \mathcal{J}(\mathcal{K}^* - \overline{\lambda})\} = 0\}$.

B.1 | The metric space

Before discussing the convergence of Algorithm 3, we must be precise about what convergence means. A reliable algorithm for computing a spectral set should converge locally on compact subsets of \mathbb{C} . In other words, it should converge to the complete spectral set and have no limiting points not in it. We will see that our algorithms achieve this, and we use the following metric space to quantify this convergence.

In general, $\sigma(\mathcal{K})$ is a closed and possibly unbounded subset of \mathbb{C} . We therefore use the Attouch–Wets metric defined by [9]:

$$d_{AW}(C_1, C_2) = \sum_{n=1}^{\infty} 2^{-n} \min\left\{1, \sup_{|x| \le n} |\operatorname{dist}(x, C_1) - \operatorname{dist}(x, C_2)|\right\},\tag{B.1}$$

where $C_1, C_2 \in Cl(\mathbb{C})$ and $Cl(\mathbb{C})$ is the set of closed nonempty²⁵ subsets of \mathbb{C} . This metric generalizes the familiar Hausdorff metric on compact sets to unbounded closed sets and corresponds to local uniform convergence on compact subsets of \mathbb{C} . If the spectral set is bounded, the two topologies and notions of convergence agree. An intuitive characterization of the Attouch–Wets topology we use in our proofs is as follows. For any closed nonempty sets C and $C_n, d_{AW}(C_n, C) \rightarrow 0$ if and only if for any $\delta > 0$ and $B_m(0)$ (closed ball of radius $m \in \mathbb{N}$ about 0), there exists N such that if n > N then $C_n \cap B_m(0) \subset C + B_{\delta}(0)$ and $C \cap B_m(0) \subset C_n + B_{\delta}(0)$.

B.2 | Computing approximate point pseudospectra

We first compute the approximate point pseudospectrum, which is given by

$$\sigma_{\epsilon,\mathrm{ap}}(\mathcal{K}) := \mathrm{cl}(\{\lambda \in \mathbb{C} : \mathcal{J}(\mathcal{K} - \lambda) < \epsilon\}).$$
(B.2)

In other words, $\sigma_{\varepsilon,ap}(\mathcal{K})$ is the closure of the set of all λ such that there exists $g \in D(\mathcal{K})$ of norm 1 with $\|(\mathcal{K} - \lambda)g\| < \varepsilon$. Such an observable g is known as an $(\varepsilon$ -)approximate eigenfunction. Moreover, $\sigma_{\varepsilon}(\mathcal{K}) = \sigma_{\varepsilon,ap}(\mathcal{K}) \cup \{\overline{\lambda} : \lambda \in \sigma_p(\mathcal{K}^*)\}$. Thus $\sigma_{\varepsilon,ap}(\mathcal{K})$ is equivalent to the usual pseudospectrum $\sigma_{\varepsilon}(\mathcal{K})$, up to the eigenvalues of \mathcal{K}^* . In most cases of interest, we have $\sigma_{\varepsilon,ap}(\mathcal{K}) = \sigma_{\varepsilon}(\mathcal{K})$.

To compute $\sigma_{\epsilon,ap}(\mathcal{K})$, we begin by providing approximations of the function $\lambda \mapsto \mathcal{J}(\mathcal{K} - \lambda)$ using the residual in (4.6). The following lemma shows how to do this.

Lemma B.1. Define a relative residual as

$$\pi_{M,N_{K}}(\lambda,\mathcal{K}) = \min_{\boldsymbol{g} \in \mathbb{C}^{N_{K}}} \sqrt{\frac{\sum_{j,k=1}^{N_{K}} \overline{g_{j}} g_{k} \left[(\Psi_{Y}^{*} W \Psi_{Y})_{jk} - \lambda (\Psi_{Y}^{*} W \Psi_{X})_{jk} - \overline{\lambda} (\Psi_{X}^{*} W \Psi_{Y})_{jk} + |\lambda|^{2} (\Psi_{X}^{*} W \Psi_{X})_{jk} \right]}{\sum_{j,k=1}^{N_{K}} \overline{g_{j}} g_{k} (\Psi_{X}^{*} W \Psi_{X})_{jk}}}$$

Then, assuming (4.7) holds,

$$\lim_{M \to \infty} \tau_{M,N_{K}}(\lambda,\mathcal{K}) = \tau_{N_{K}}(\lambda,\mathcal{K}) := \min_{g \in V_{N_{K}}} \sqrt{\frac{\int_{\Omega} \left| [\mathcal{K}g](\boldsymbol{x}) - \lambda g(\boldsymbol{x}) \right|^{2} d\omega(\boldsymbol{x})}{\int_{\Omega} \left| g(\boldsymbol{x}) \right|^{2} d\omega(\boldsymbol{x})}}.$$
(B.4)

Moreover, $\tau_{N_K}(\lambda, \mathcal{K})$ is a nonincreasing function of N_K and converges uniformly down to $\mathcal{J}(\mathcal{K} - \lambda)$ on compact subsets of \mathbb{C} as a function of the spectral parameter λ .

Proof. The limit (B.4) follows trivially from the convergence of matrices in (4.7). Since $V_{N_K} \subset V_{N_K+1}, \tau_{N_K}(\lambda, \mathcal{K})$ is a nonincreasing function of N_K . By definition, we also have that $\tau_{N_K}(\lambda, \mathcal{K}) \geq \mathcal{J}(\mathcal{K} - \lambda)$ for every λ and N_K . Now let $0 < \epsilon < 1$ and choose $g \in \mathcal{D}(\mathcal{K})$ of norm 1 such that $\|(\mathcal{K} - \lambda)g\| \leq \mathcal{J}(\mathcal{K} - \lambda) + \epsilon$. Since span $\{\psi \in V_{N_K} : N_K \in \mathbb{N}\}$ forms a core of \mathcal{K} , there exists some n and $g_n \in V_n$ such that $\|(\mathcal{K} - \lambda)g_n\| \leq \|(\mathcal{K} - \lambda)g\| + \epsilon$ and $\|g - g_n\| \leq \epsilon$. We find that

(B.3)

²⁵ We can add the empty set as an isolated point of this metric space to cover the case that $\sigma(\mathcal{K})$ may be empty. In what follows, this is implicitly done in our algorithms since the output is eventually the empty set in such circumstances.

 $\|g_n\| \ge 1 - \epsilon$ and hence that $\tau_n(\lambda, \mathcal{K}) \le (\mathcal{J}(\mathcal{K} - \lambda) + 2\epsilon)/(1 - \epsilon)$. Since this holds for any $0 < \epsilon < 1$, $\lim_{N_K \to \infty} \tau_{N_K}(\lambda, \mathcal{K}) = \mathcal{J}(\mathcal{K} - \lambda)$. Since $\mathcal{J}(\mathcal{K} - \lambda)$ is Lipschitz continuous in λ and hence continuous, $\tau_{N_K}(\lambda, \mathcal{K})$ converges uniformly down to $\mathcal{J}(\mathcal{K} - \lambda)$ on compact subsets of \mathbb{C} by Dini's theorem.

In other words, in the large data limit, we can compute $\tau_{N_K}(\lambda, \mathcal{K})$, which is at least as large as $\mathcal{J}(\mathcal{K} - \lambda)$, and converges uniformly as our truncation parameter N_K increases. To compute $\tau_{M,N_K}(\lambda, \mathcal{K})$, one can first diagonalize $\Psi_X^*W\Psi_X$ to compute $(\Psi_X^*W\Psi_X)^{-1/2}$ and then the problem reduces to computing a standard singular value decomposition. To turn $\tau_{N_K}(\lambda, \mathcal{K})$ into an approximation of $\sigma_{\epsilon,\mathrm{ap}}(\mathcal{K})$, we search for points where this function is less than ϵ . Suppose that $\mathrm{Grid}(N_K)$ is a sequence of finite grids such that for any $\lambda \in \mathbb{C}$, $\lim_{N_K \to \infty} \mathrm{dist}(\lambda, \mathrm{Grid}(N_K)) = 0$. For example, we could take

$$\operatorname{Grid}(N_K) = \frac{1}{N_K} [\mathbb{Z} + i\mathbb{Z}] \cap \{ z \in \mathbb{C} : |z| \le N_K \}.$$
(B.5)

In practice, one considers a grid of points over the region of interest in the complex plane. We then define

$$\Gamma_{N_{\mathcal{K}}}^{\epsilon}(\mathcal{K}) := \left\{ \lambda \in \operatorname{Grid}(N_{\mathcal{K}}) : \tau_{N_{\mathcal{K}}}(\lambda, \mathcal{K}) < \epsilon \right\}.$$

Our algorithm is summarized in Algorithm 3. Note that because of the nature of the nonlinear pencil in (B.3), one cannot use decompositions such as the generalized Schur decomposition to speed up the computational time. We can now state our convergence theorem, which says that the output of Algorithm 3 lies in $\sigma_{\varepsilon,ap}(\mathcal{K})$ and converges to $\sigma_{\varepsilon,ap}(\mathcal{K})$ as $M \to \infty$.

Theorem B.1. Let $\epsilon > 0$. Then, $\Gamma_{N_{\kappa}}^{\epsilon}(\mathcal{K}) \subset \sigma_{\epsilon, \mathrm{ap}}(\mathcal{K})$ and

$$\lim_{N_K \to \infty} d_{AW} \Big(\Gamma_{N_K}^{\epsilon}(\mathcal{K}), \sigma_{\epsilon, ap}(\mathcal{K}) \Big) = 0.$$
(B.6)

In other words, we recover the full approximate point pseudospectrum without spectral pollution.

Proof. If $\lambda \in \Gamma_{N_K}^{\epsilon}(\mathcal{K})$, then $\mathcal{J}(\mathcal{K}-\lambda) \leq \tau_{N_K}(\lambda, \mathcal{K}) < \epsilon$ by Lemma B.1. It follows that $\Gamma_{N_K}^{\epsilon}(\mathcal{K}) \subset \sigma_{\epsilon, \mathrm{ap}}(\mathcal{K})$. Now suppose that $\sigma_{\epsilon, \mathrm{ap}}(\mathcal{K}) \neq \emptyset$. It follows that there exists $\lambda \in \mathbb{C}$ with $\mathcal{J}(\mathcal{K}-\lambda) < \epsilon$. Since $\mathcal{J}(\mathcal{K}-\lambda)$ is continuous in λ and $\lim_{N_K \to \infty} \operatorname{dist}(\lambda, \operatorname{Grid}(N_K)) = 0$, it follows that $\Gamma_{N_K}^{\epsilon}(\mathcal{K}) \neq \emptyset$ for large N_K . We use the characterization of the Attouch–Wets topology. Suppose that m is large such that $B_m(0) \cap \sigma_{\epsilon, \mathrm{ap}}(\mathcal{K}) \neq \emptyset$. Since $\Gamma_{N_K}^{\epsilon}(\mathcal{K}) \subset \sigma_{\epsilon, \mathrm{ap}}(\mathcal{K})$, we clearly have $\Gamma_{N_K}^{\epsilon}(\mathcal{K}) \cap B_m(0) \subset \sigma_{\epsilon, \mathrm{ap}}(\mathcal{K})$. Hence, to prove (B.6), we must show that given $\delta > 0$, there exists n_0 such that if $N_K > n_0$ then $\sigma_{\epsilon, \mathrm{ap}}(\mathcal{K}) \cap B_m(0) \subset \Gamma_{N_K}^{\epsilon}(\mathcal{K}) + B_{\delta}(0)$. Suppose for a contradiction that this statement is false. Then, there exists $\delta > 0$, $\lambda_{n_j} \in \sigma_{\epsilon, \mathrm{ap}}(\mathcal{K}) \cap B_m(0)$, and $n_j \to \infty$ such that $\operatorname{dist}(\lambda_{n_j}, \Gamma_{n_j}^{\epsilon}(\mathcal{K})) \geq \delta$. Without loss of generality, we can assume that $\lambda_{n_j} \to \lambda \in \sigma_{\epsilon, \mathrm{ap}}(\mathcal{K}) \cap B_m(0)$. There exists some z with $\mathcal{J}(\mathcal{K}-z) < \epsilon$ and $|\lambda-z| \leq \delta/2$. Let $z_{n_j} \in \operatorname{Grid}(n_j)$ such that $|z - z_{n_j}| \leq \operatorname{dist}(z, \operatorname{Grid}(n_j)) + n_j^{-1}$. In particular, note that $\lim_{n_j \to \infty} |z - z_{n_j}| = 0$. Moreover,

$$\tau_{n_j}(z_{n_j},\mathcal{K}) \le \left|\tau_{n_j}(z_{n_j},\mathcal{K}) - \mathcal{J}(\mathcal{K} - z_{n_j})\right| + \left|\mathcal{J}(\mathcal{K} - z) - \mathcal{J}(\mathcal{K} - z_{n_j})\right| + \underbrace{\mathcal{J}(\mathcal{K} - z)}_{<\varepsilon}.$$
 (B.7)

But, $\mathcal{J}(\mathcal{K} - \lambda)$ is continuous in λ and τ_{n_j} converges uniformly to $\mathcal{J}(\mathcal{K} - \cdot)$ on compact subsets of \mathbb{C} by Lemma B.1. It follows that the first two terms on the right-hand side of (B.7) must converge to zero. Hence for large $n_j, \tau_{n_j}(z_{n_j}, \mathcal{K}) < \epsilon$ so that $z_{n_j} \in \Gamma_{n_j}^{\epsilon}(\mathcal{K})$. But $|z_{n_j} - \lambda| \le |z - \lambda| + |z_{n_j} - z| \le \delta/2 + |z - z_{n_i}|$, which is smaller than δ for large n_j , and we reach the desired contradiction. \Box

Theorem B.1 is stated for an algorithm that uses τ_{N_K} . Strictly speaking, τ_{N_K} is only computed as $M \to \infty$ (see Lemma B.1). In practice, this is not an issue since we can often estimate the accuracy of the quadrature rule for the computed inner products by using known convergence rates as $M \to \infty$ (see Section 4.1.3) or by comparing answers for different values of M. We can then combine these estimates and adaptively select M based on N_K so that τ_{N_K} is sufficiently accurate. For example, if τ_{N_K} is approximated to accuracy $\mathcal{O}(N_K^{-1})$ by $\tilde{\tau}_{N_K}$ we can set $\Gamma_{N_K}^{\epsilon}(\mathcal{K}) = \{\lambda \in \operatorname{Grid}(N_K) : \tilde{\tau}_{N_K}(\lambda, \mathcal{K}) + N_K^{-1/2} < \epsilon\}$, where the $N_K^{-1/2}$ term is to ensure that it eventually dominates the quadrature error.

B.3 | Recovering the full spectrum and pseudospectrum

In this subsection, we make the following additional assumption:

• Assumption: span{ $\psi \in V_{N_K} : N_K \in \mathbb{N}$ } forms a core of \mathcal{K}^* .

If we have access to the matrix elements $\langle \mathcal{K}^* \psi_j, \mathcal{K}^* \psi_i \rangle$, then we can apply the techniques described in Section 4 to compute $\tau_{N_{\mathcal{K}}}(\overline{\lambda}, \mathcal{K}^*)$ in the large data limit $M \to \infty$. Define

$$\hat{\Gamma}_{N_{K}}^{\epsilon}(\mathcal{K}) = \left\{ \lambda \in \operatorname{Grid}(N_{K}) : \min\{\tau_{N_{K}}(\lambda, \mathcal{K}), \tau_{N_{K}}(\overline{\lambda}, \mathcal{K}^{*})\} < \epsilon \right\}.$$

The proof of Theorem B.1 can be adapted to show that $\hat{\Gamma}_{N_{K}}^{\epsilon}(\mathcal{K}) \subset \sigma_{\epsilon}(\mathcal{K})$, as well as the convergence $\lim_{N_{K}\to\infty} d_{AW}(\hat{\Gamma}_{N_{K}}^{\epsilon}(\mathcal{K}), \sigma_{\epsilon}(\mathcal{K})) = 0$. Unfortunately, in general, one does not have access to approximations of $\langle \mathcal{K}^{*}\psi_{j}, \mathcal{K}^{*}\psi_{i} \rangle$. Nevertheless, one can overcome this.

To compute $\sigma_{\epsilon}(\mathcal{K})$, we note that because of Theorem B.1, it is enough to compute $\sigma_{\epsilon,ap}(\mathcal{K}^*)$ since $\sigma_{\epsilon}(\mathcal{K}) = \sigma_{\epsilon,ap}(\mathcal{K}) \cup \{\overline{\lambda} : \lambda \in \sigma_{\epsilon,ap}(\mathcal{K}^*)\}$. Let $N_1, N_2 \in \mathbb{N}$, with $N_1 \ge N_2$, and consider the matrices Ψ_X and Ψ_Y computed using N_1 dictionary functions. Assume that (4.7) holds and let

$$L_{N_1}(\lambda) = \lim_{M \to \infty} \Psi_X^* W \Psi_Y (\Psi_X^* W \Psi_X)^{-1} \Psi_Y^* W \Psi_X - \lambda \Psi_Y^* W \Psi_X - \overline{\lambda} \Psi_X^* W \Psi_Y + |\lambda|^2 \Psi_X^* W \Psi_X \in \mathbb{C}^{N_1 \times N_1}$$

Note that this matrix can be obtained from the given data in the large data limit $M \to \infty$. Let \mathcal{P}_S denote the orthogonal projection onto a subspace $S \subset L^2(\Omega, \omega)$. If $g = \sum_{j=1}^{N_2} \psi_j g_j \in V_{N_2}$, then

$$\left\langle \mathcal{P}_{V_{N_1}}(\mathcal{K}^* - \overline{\lambda})g, \mathcal{P}_{V_{N_1}}(\mathcal{K}^* - \overline{\lambda})g \right\rangle = \sum_{j,k=1}^{N_2} \overline{g_j}g_k[L_{N_1}(\lambda)]_{jk}.$$
 (B.8)

In other words, this inner product is given by a square $N_2 \times N_2$ truncation of the matrix $L_{N_1}(\lambda) \in \mathbb{C}^{N_1 \times N_1}$. The following lemma uses the fact that (B.8) provides an increasingly more accurate approximation of the inner product $\langle (\mathcal{K}^* - \overline{\lambda})g, (\mathcal{K}^* - \overline{\lambda})g \rangle$, as $N_1 \to \infty$.

Lemma B.2. Define $v_{N_1,N_2}(\lambda, \mathcal{K}) = \min_{g \in V_{N_2}, ||g||=1} \sqrt{\sum_{j,k=1}^{N_2} \overline{g_j} g_k [L_{N_1}(\lambda)]_{jk}}$. Then, $v_{N_1,N_2}(\lambda, \mathcal{K})$ is a nondecreasing function of N_1 and, for fixed N_2 , converges uniformly up to $\tau_{N_2}(\overline{\lambda}, \mathcal{K}^*)$ on compact subsets of \mathbb{C} as a function of λ as $N_1 \to \infty$.

61

Proof. It is clear that $v_{N_1,N_2}(\lambda, \mathcal{K})$ is a nondecreasing function of N_1 . Since $\mathcal{P}_{V_{N_1}}$ converges strongly to the identity and $(\mathcal{K}^* - \overline{\lambda})\mathcal{P}^*_{V_{N_2}}$ acts on a finite-dimensional subspace, we also have that $\lim_{N_1 \to \infty} \|(\mathcal{P}_{V_{N_1}} - I)(\mathcal{K}^* - \overline{\lambda})\mathcal{P}^*_{V_{N_2}}\| = 0$. The uniform convergence of $v_{N_1,N_2}(\lambda, \mathcal{K})$ up to $\tau_{N_2}(\overline{\lambda}, \mathcal{K}^*)$ on compact subsets of \mathbb{C} now follows from Dini's theorem.

The following theorem shows that we have inclusion and convergence.

Theorem B.2. Let $\epsilon > 0$ and define $\hat{\Gamma}_{N_2,N_1}^{\epsilon}(\mathcal{K}) = \{\lambda \in \operatorname{Grid}(N_2) : \upsilon_{N_1,N_2}(\lambda,\mathcal{K}) + 1/N_2 \le \epsilon\}$. Then $\lim_{N_1 \to \infty} \hat{\Gamma}_{N_2,N_1}^{\epsilon}(\mathcal{K}) = \hat{\Gamma}_{N_2}^{\epsilon}(\mathcal{K}) = \{\lambda \in \operatorname{Grid}(N_2) : \tau_{N_2}(\overline{\lambda},\mathcal{K}^*) + 1/N_2 \le \epsilon\} =: \hat{\Gamma}_{N_2}^{\epsilon}(\mathcal{K})$. Moreover, we have $\hat{\Gamma}_{N_2}^{\epsilon}(\mathcal{K}) \subset \sigma_{\epsilon,\operatorname{ap}}(\mathcal{K}^*)$ and $\lim_{N_2 \to \infty} d_{\operatorname{AW}}(\hat{\Gamma}_{N_2}^{\epsilon}(\mathcal{K}), \{\overline{\lambda} : \lambda \in \sigma_{\epsilon,\operatorname{ap}}(\mathcal{K}^*)\}) = 0$.

Proof. We prove that $\hat{\Gamma}_{N_2}^{\epsilon}(\mathcal{K}) = \lim_{N_1 \to \infty} \hat{\Gamma}_{N_2,N_1}^{\epsilon}(\mathcal{K})$ exists and is given by $\{\lambda \in \operatorname{Grid}(N_2) : \tau_{N_2}(\overline{\lambda}, \mathcal{K}^*) + 1/N_2 \leq \epsilon\}$. The rest of the proof carries over from Theorem B.1 and its proof with minor modifications. (The additional term N_2^{-1} is needed to ensure convergence to the pseudospectrum in the second limit - it is added so that we can use " \leq " instead of "<" in the definition of $\hat{\Gamma}_{N_2,N_1}^{\epsilon}$ to ensure the existence of the first limit via the following argument.) Fix $\lambda \in \operatorname{Grid}(N_2)$. If $\tau_{N_2}(\overline{\lambda}, \mathcal{K}^*) + N_2^{-1} \leq \epsilon$, then by Lemma B.2, $v_{N_1,N_2}(\lambda, \mathcal{K}) + N_2^{-1} \leq \epsilon$. It follows that $\hat{\Gamma}_{N_2}^{\epsilon}(\mathcal{K}) \subset \hat{\Gamma}_{N_2,N_1}^{\epsilon}(\mathcal{K})$. If $\tau_{N_2}(\overline{\lambda}, \mathcal{K}^*) + N_2^{-1} > \epsilon$, then since $\lim_{N_1 \to \infty} v_{N_1,N_2}(\lambda, \mathcal{K}) = \tau_{N_2}(\overline{\lambda}, \mathcal{K}^*)$, $v_{N_1,N_2}(\lambda, \mathcal{K}) + N_2^{-1} > \epsilon$ for large N_1 . The theorem now follows.

Combining Theorem B.1 and Theorem B.2, we have the following corollary that shows how the full spectrum and pseudospectrum can be recovered.

Corollary B.1. For any $\epsilon > 0$,

$$\lim_{N_2\to\infty}\lim_{N_1\to\infty}\left[\Gamma_{N_2}^{\epsilon}(\mathcal{K})\cup\hat{\Gamma}_{N_2,N_1}^{\epsilon}(\mathcal{K})\right]=\sigma_{\epsilon}(\mathcal{K}).$$

Additionally,

$$\lim_{\varepsilon \downarrow 0} \lim_{N_2 \to \infty} \lim_{N_1 \to \infty} \left[\Gamma^{\varepsilon}_{N_2}(\mathcal{K}) \cup \hat{\Gamma}^{\varepsilon}_{N_2,N_1}(\mathcal{K}) \right] = \sigma(\mathcal{K})$$

where the convergence holds in the Attouch-Wets topology.

As a final remark, the reader may notice that a few of our algorithms require us to take several parameters successively to infinity. These limits do not generally commute, and it may be impossible to rewrite them in fewer limits or develop a different algorithm that uses fewer limits. This is a generic feature of infinite-dimensional spectral problems [29] and has given rise to the Solvability Complexity Index [35, 60]. We do not go into the details, but there are many open questions on the foundations of computing spectral properties of Koopman operators.

APPENDIX C: COMPUTING RATIONAL KERNEL CONVOLUTIONS WITH ERROR CONTROL

This section shows that the required convolutions with rational kernels can be computed rigorously with error control. Recall from (5.7) that there are two types of inner products that we need to compute: (i) $\langle g, (\mathcal{K} - \lambda)^{-1}g \rangle$, and (ii) $\langle (\mathcal{K} - \lambda)^{-1}g, \mathcal{K}^*g \rangle$ for some observable g. Here, we show that the inner products can be computed with error control via the right-hand sides of (5.14) and (5.15) by adaptively selecting the truncation size N_K for a fixed smoothing parameter ϵ . Without loss of generality, we assume that g is normalized so that ||g|| = 1 and we denote $\tilde{g}_{N_K} = \Psi \boldsymbol{a}$. To compute the error bounds in the large data limit, we also use the fact that

$$h_{N_{K}} = \sum_{j=1}^{N_{K}} \left[(\Psi_{X}^{*} W \Psi_{Y} - \lambda \Psi_{X}^{*} W \Psi_{X})^{-1} \Psi_{X}^{*} W \Psi_{X} \boldsymbol{a} \right]_{j} \boldsymbol{\psi}_{j} = \Psi \boldsymbol{h} \approx (\mathcal{P}_{V_{N}} \mathcal{K} \mathcal{P}_{V_{N}}^{*} - \lambda I_{N})^{-1} g_{N_{K}}.$$

The following theorem shows that we can bound the error of using the right-hand sides of (5.14) and (5.15) to approximate the desired inner products.

Theorem C.1. Let $\lambda \in \mathbb{C}$ with $|\lambda| > 1$. Consider the setup of Section 5.3 with ||g|| = 1 and let

$$\begin{split} \delta_1(N_K) &= \|g - g_{N_K}\|, \\ \delta_2(N_K) &= \|(\mathcal{K} - \lambda)(\mathcal{P}_{V_{N_K}}\mathcal{K}\mathcal{P}_{V_{N_K}}^* - \lambda I_{N_K})^{-1}g_{N_K} - g_{N_K}\|, \\ \delta_3(N_K) &= \|(\mathcal{P}_{V_{N_K}}\mathcal{K}\mathcal{P}_{V_{N_K}}^* - \lambda I_{N_K})^{-1}g_{N_K}\|. \end{split}$$

Then, $\lim_{N_K\to\infty} \delta_1(N_K) = \lim_{N_K\to\infty} \delta_2(N_K) = 0$ and $\delta_3(N_K)$ remains bounded as $N_K \to \infty$. Moreover,

$$\left| \left\langle g_{N_{K}}, (\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1} g_{N_{K}} \right\rangle - \left\langle g, (\mathcal{K} - \lambda)^{-1} g \right\rangle \right| \leq \frac{\delta_{1}(N_{K}) + \delta_{2}(N_{K})}{|\lambda| - 1} + \delta_{1}(N_{K})\delta_{3}(N_{K}),$$
$$\left| \left\langle (\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1} g_{N_{K}}, \mathcal{K}^{*} g_{N_{K}} \right\rangle - \left\langle (\mathcal{K} - \lambda)^{-1} g, \mathcal{K}^{*} g \right\rangle \right| \leq \frac{\delta_{1}(N_{K}) + \delta_{2}(N_{K})}{|\lambda| - 1} + \delta_{1}(N_{K})\delta_{3}(N_{K}).$$

Assuming (4.7), $\delta_2(N_K)$ and $\delta_3(N_K)$ can be computed in the large data limit via

$$\delta_2(N_K)^2 = \lim_{M \to \infty} \boldsymbol{h}^* \Psi_X^* W \Psi_X \boldsymbol{h} \cdot \operatorname{res}(\lambda, h_{N_K})^2 - 2\operatorname{Re}\left(\boldsymbol{a}^* (\Psi_X^* W \Psi_Y - \lambda \Psi_X^* W \Psi_X) \boldsymbol{h}\right) + \boldsymbol{a}^* \Psi_X^* W \Psi_X \boldsymbol{a},$$
(C.1)

$$\delta_3(N_K)^2 = \lim_{M \to \infty} \boldsymbol{h}^* \Psi_X^* W \Psi_X \boldsymbol{h}.$$
(C.2)

Proof. First, by assumption, $\lim_{N_K \to \infty} \delta_1(N_K) = 0$. Using this, together with Theorem 5.1 and the fact that \mathcal{K} is bounded, it follows that $\lim_{N_K \to \infty} \delta_2(N_K) = 0$. Similarly, $\delta_3(N_K)$ remains bounded as $N_K \to \infty$. Since $|\lambda| > 1$ and $||\mathcal{K}|| = 1$, the arguments in Section 5.2.2 show that $||(\mathcal{K} - \lambda)^{-1}|| \le 1/(|\lambda| - 1)$. It follows that

$$\begin{split} \|(\mathcal{P}_{V_{N_{K}}}\mathcal{K}\mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1}g_{N_{K}} - (\mathcal{K} - \lambda)^{-1}g\| &\leq \frac{\|(\mathcal{K} - \lambda)(\mathcal{P}_{V_{N_{K}}}\mathcal{K}\mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1}g_{N_{K}} - g\|}{|\lambda| - 1} \\ &\leq \frac{\|(\mathcal{K} - \lambda)(\mathcal{P}_{V_{N_{K}}}\mathcal{K}\mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1}g_{N_{K}} - g_{N_{K}}\| + \|g - g_{N_{K}}\|}{|\lambda| - 1} \\ &\leq \frac{\delta_{1}(N_{K}) + \delta_{2}(N_{K})}{|\lambda| - 1}. \end{split}$$

We now re-write the inner product

$$\begin{split} \left\langle g_{N_{K}}, (\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1} g_{N_{K}} \right\rangle = \left\langle g, (\mathcal{K} - \lambda)^{-1} g \right\rangle + \left\langle g_{N_{K}} - g, (\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1} g_{N_{K}} \right\rangle \\ + \left\langle g, (\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1} g_{N_{K}} - (\mathcal{K} - \lambda)^{-1} g \right\rangle. \end{split}$$

Since ||g|| = 1, it follows by the Cauchy–Schwarz inequality that

$$\begin{split} \left| \left\langle g_{N_{K}}, (\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1} g_{N_{K}} \right\rangle - \left\langle g, (\mathcal{K} - \lambda)^{-1} g \right\rangle \right| &\leq \frac{\delta_{1}(N_{K}) + \delta_{2}(N_{K})}{|\lambda| - 1} \\ &+ \delta_{1}(N_{K}) \Big\| (\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1} g_{N_{K}} \Big\|. \end{split}$$

Similarly, since $\|\mathcal{K}^*\| \leq 1$, we obtain

$$\begin{split} \left| \left\langle (\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1} g_{N_{K}}, \mathcal{K}^{*} g_{N_{K}} \right\rangle - \left\langle (\mathcal{K} - \lambda)^{-1} g, \mathcal{K}^{*} g \right\rangle \right| &\leq \frac{\delta_{1}(N_{K}) + \delta_{2}(N_{K})}{|\lambda| - 1} \\ &+ \delta_{1}(N_{K}) \left\| (\mathcal{P}_{V_{N_{K}}} \mathcal{K} \mathcal{P}_{V_{N_{K}}}^{*} - \lambda I_{N_{K}})^{-1} g_{N_{K}} \right\|. \end{split}$$

This proves the upper bounds of the theorem. The equalities in (C.1) and (C.2) follow from expanding the inner products and from the definition of h.

From Theorem C.1 we see that if $\delta_1(N_K)$ is known, then we can compute the inner products $\langle g, (\mathcal{K} - \lambda)^{-1}g \rangle$ and $\langle (\mathcal{K} - \lambda)^{-1}g, \mathcal{K}^*g \rangle$ to any desired accuracy by adaptively selecting N_K . Moreover, the error bounds in Theorem C.1 can be computed in $\mathcal{O}(N_K^2)$ operations without the generalized Schur decomposition. In practice, we increase N_K , keeping track of the computed inner product until convergence is apparent.