
Sharp Spectral Rates for Koopman Operator Learning

Vladimir R. Kostic

Istituto Italiano di Tecnologia
University of Novi Sad
vladimir.kostic@iit.it

Karim Lounici

CMAE-Ecole Polytechnique
karim.lounici@polytechnique.edu

Pietro Novelli

Istituto Italiano di Tecnologia
pietro.novelli@iit.it

Massimiliano Pontil

Istituto Italiano di Tecnologia
University College London
massimiliano.pontil@iit.it

Abstract

Nonlinear dynamical systems can be handily described by the associated Koopman operator, whose action evolves every observable of the system forward in time. Learning the Koopman operator and its spectral decomposition from data is enabled by a number of algorithms. In this work we present for the first time non-asymptotic learning bounds for the Koopman eigenvalues and eigenfunctions. We focus on time-reversal-invariant stochastic dynamical systems, including the important example of Langevin dynamics. We analyze two popular estimators: Extended Dynamic Mode Decomposition (EDMD) and Reduced Rank Regression (RRR). Our results critically hinge on novel minimax estimation bounds for the operator norm error, that may be of independent interest. Our spectral learning bounds are driven by the simultaneous control of the operator norm error and a novel metric distortion functional of the estimated eigenfunctions. The bounds indicates that both EDMD and RRR have similar variance, but EDMD suffers from a larger bias which might be detrimental to its learning rate. Our results shed new light on the emergence of spurious eigenvalues, an issue which is well known empirically. Numerical experiments illustrate the implications of the bounds in practice.

1 Introduction

Recently, researchers have emphasized the utmost importance of developing physically-informed machine learning models that prioritize interpretability and foster physical insight and intuition, see for example [22] and references therein. One technique highlighted in these works is the Koopman operator regression framework to learn and interpret nonlinear dynamical systems see, e.g. [8, 26] and references therein. A key component of this approach is the Koopman Mode Decomposition (KMD), which decomposes complex dynamical systems into simpler, coherent structures. When ordinary least squares are used to learn Koopman operator from data, estimated KMD is known as the Dynamic Mode Decomposition (DMD) [36]. Koopman operator estimators and their modal decomposition find many applications, including fluid dynamics, molecular kinetics and robotics [7, 19].

The Koopman operator returns the expected value of observables of the system in the future given the present, and one relies on estimators of this operator to in turn estimate its spectral decomposition that leads to the estimation of KMD. Our goal is to study the statistical properties of the eigenvalues and eigenfunctions of the Koopman operator estimators via two mainstream algorithms: Principal Component Regression (PCR) and Reduced Rank Regression (RRR) studied in [21, 24]. PCR encompasses as particular cases the popular Extended Dynamic Mode Decomposition (EDMD), which is the de-facto estimator in the data-driven dynamical system literature [see 26, 45, and references

therein]. Both PCR and RRR are kernel-based algorithms that, given a dataset of observations of the dynamical system, implement a strategy to approximate the action of the Koopman operator on a reproducing kernel Hilbert space (RKHS) [3, 38].

We present for the first time non-asymptotic learning bounds on the distance between the Koopman eigenvalues and eigenfunctions and those estimated by either PCR or RRR. We show that the eigenvalues produced by such algorithms are *biased* estimators of the true Koopman eigenvalues, with PCR incurring larger bias. Our results critically hinge on novel estimation bounds for the operator norm error, that may be of independent interest, leading to minimax optimal bounds for finite-rank Koopman operators. Moreover, we introduce the novel notion of metric distortion, which characterizes how the norm of eigenfunctions vary when moving from the RKHS in which learning takes place to the underlying ambient space where the Koopman operator is properly defined. We show that both the operator norm error and metric distortion are needed in order to estimate the operator spectra and our bounds can be used to explain the well-known spuriousness phenomena in eigenvalue estimation [14], namely, the scenario in which the estimated eigenvalues are not related to the true ones, despite small operator norm error.

Contributions and Organization. We make the following contributions: **i)** We introduce the notion of metric distortion and show that it has to be used alongside the operator norm error to derive Koopman spectra estimation error bounds (Theorem 1); **ii)** We establish the first sharp estimation bound for the operator norm error (Theorem 2); **iii)** We establish spectral learning rates (Thms. 3 and 4) for both PCR and RRR; **iv)** We propose how to use the results entailed by Theorem 4 to detect the presence of spurious eigenvalues from data.

The paper is organized as follows. In Section 2 we recall the notion of Koopman operator, its spectral decomposition, and review PCR and RRR estimators. Section 3 describes the estimation problem and outline our main results. Section 4 presents our approach to bound eigenvalue and eigenvector estimation errors. Section 5 gives sharp upper bounds for the operator norm error. Section 6 presents our spectral learning bounds. Finally, Section 7 illustrates the implications of the bounds in practice, and is designed to provide practitioners with the tools to benchmark the performance of algorithms in real scenarios.

2 Background

Dynamical Systems and Koopman Operator. In this work we study Markovian dynamical systems, that is collections of random variables $\{X_t : t \in \mathbb{N}\}$, where X_t represents the *state* at time t , taking values in some space \mathcal{X} . We focus on time-homogeneous (i.e. autonomous) systems hosting an invariant measure π for which the *Koopman operator* [24, 27]

$$(A_\pi f)(x) := \mathbb{E}[f(X_{t+1}) | X_t = x], \quad x \in \mathcal{X} \quad (1)$$

is a well defined bounded linear operator on $L_\pi^2(\mathcal{X})$, the space of square integrable functions on \mathcal{X} relative to measure π . In the field of stochastic processes, (1) is also known as the *transfer operator* and returns the expected value of f in the future given the present. This operator is self-adjoint (i.e. $A_\pi = A_\pi^*$) whenever dynamics is time-reversal invariant w.r.t. π , which is satisfied by many stochastic processes in the physical sciences.

Example 1 (Langevin Dynamics). *Let $\mathcal{X} = \mathbb{R}^d$ and let $\beta > 0$. The (overdamped) Langevin equation driven by a potential $U : \mathbb{R}^d \rightarrow \mathbb{R}$ is given by $dX_t = -\nabla U(X_t)dt + \sqrt{2\beta^{-1}}dW_t$, where W_t is a Wiener process. The invariant measure of this process is the Boltzman distribution $\pi(dx) \propto e^{-\beta U(x)}dx$, and the associated Koopman operator is self-adjoint.*

The Langevin equation models a wealth of phenomena, such as the evolution of chemical and biological systems at thermal equilibrium [16], the mechanism regulating cell size in bacteria [1], chemical reactions [25], the dynamics of synapses [11, 42], stock market fluctuations [6] and many more. Furthermore, when $U(x) = \theta \|x\|^2/2$ ($\theta > 0$), the Langevin equation reduces to the celebrated Ornstein–Uhlenbeck process [34, Chapter 6].

The operator (1) evolves every observable of the system forward in time. Since it is bounded and linear, it admits a *spectral decomposition*, which plays a central role in the analysis and interpretation of the dynamical system [26], as well as (nonlinear) control [2]. As in [46], to study the spectral decomposition we further assume that A_π is a *compact* operator, which rules out the presence of

continuous and residual spectrum components and leads to

$$A_\pi = \sum_{i \in \mathbb{N}} \mu_i f_i \otimes f_i, \quad (2)$$

where $(\mu_i, f_i)_{i \in \mathbb{N}} \subseteq \mathbb{R} \times L_\pi^2(\mathcal{X})$ are Koopman eigenpairs, i.e. $A_\pi f_i = \mu_i f_i$. Moreover, $\lim_{i \rightarrow \infty} \mu_i = 0$ and $\{f_i\}_{i \in \mathbb{N}}$ form a complete orthonormal system of $L_\pi^2(\mathcal{X})$. In the context of molecular dynamics, the leading eigenvalues and their eigenfunctions are key in the study of long-term dynamics and so-called *meta-stable* states [see, e.g., 40].

Koopman Operator Regression in RKHS. Throughout the paper we let \mathcal{H} be an RKHS and let $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be the associated kernel function. We let $\phi : \mathcal{X} \rightarrow \mathcal{H}$ be a *feature map* [38] such that $k(x, x') = \langle \phi(x), \phi(x') \rangle$ for all $x, x' \in \mathcal{X}$. We consider RKHSs satisfying $\mathcal{H} \subset L_\pi^2(\mathcal{X})$ [38, Chapter 4.3], so that PCR and RRR approximate $A_\pi : L_\pi^2(\mathcal{X}) \rightarrow L_\pi^2(\mathcal{X})$ with an operator $G : \mathcal{H} \rightarrow \mathcal{H}$. Notice that despite $\mathcal{H} \subset L_\pi^2(\mathcal{X})$, the two spaces have different metric structures, that is for all $f, g \in \mathcal{H} \subset L_\pi^2(\mathcal{X})$, one in general has $\langle f, g \rangle_{\mathcal{H}} \neq \langle f, g \rangle_{L_\pi^2(\mathcal{X})}$. In order to handle this ambiguity, we introduce the *injection operator* $S : \mathcal{H} \rightarrow L_\pi^2(\mathcal{X})$ such that for all $f \in \mathcal{H}$, the object Sf is the element of $L_\pi^2(\mathcal{X})$ which is pointwise equal to $f \in \mathcal{H}$, but endowed with the appropriate $L_\pi^2(\mathcal{X})$ norm. With this in mind, the Koopman operator restricted to \mathcal{H} is simply $A_\pi S$, which is then estimated by SG for some $G \in \text{HS}(\mathcal{H})$. We will measure the operator norm error, $\|A_\pi S - SG\|$. This is in contrast to the more frequently used Hilbert-Schmidt (HS) norm.

Koopman operator regression estimators are supervised learning algorithms to learn the Koopman operator, in which input and output data are consecutive states of the system (X_t, X_{t+1}) for some $t \in \mathbb{N}$. Since the Markov process is time-homogeneous and stationary, the joint probability distribution of (X_t, X_{t+1}) is the same for every $t \in \mathbb{N}$ and we denote it by ρ . Furthermore, stationarity also implies that $X_t \sim \pi$ for all $t \in \mathbb{N}$. Given a dataset¹ $\mathcal{D}_n := (x_i, y_i)_{i=1}^n$ of consecutive states, PCR and RRR are two different strategies to minimize, under a fixed-rank constraint, the mean square error

$$\hat{\mathcal{R}}(G) := \frac{1}{n} \sum_{i \in [n]} \|\phi(y_i) - G^* \phi(x_i)\|^2, \quad (3)$$

where $G \in \text{HS}(\mathcal{H})$, the space of Hilbert-Schmidt operator acting on \mathcal{H} . PCR and RRR estimators are expressed as functions of the *input* and *cross* empirical covariances, defined respectively as

$$\hat{C} = \frac{1}{n} \sum_{i \in [n]} \phi(x_i) \otimes \phi(x_i), \quad \text{and} \quad \hat{T} = \frac{1}{n} \sum_{i \in [n]} \phi(x_i) \otimes \phi(y_i).$$

Likewise, the population risk is $\mathcal{R}(G) = \mathbb{E}_{(X,Y) \sim \rho} \|\phi(Y) - G^* \phi(X)\|^2$, and the population covariance and cross covariance are $C = \mathbb{E}_{X \sim \pi} \phi(X) \otimes \phi(X)$, and $T = \mathbb{E}_{(X,Y) \sim \rho} \phi(X) \otimes \phi(Y)$, respectively. We note that by the *reproducing kernel property* one finds that $C = S^* S$; see e.g. [38].

Two Important Estimators. We next briefly recall two operator regression estimators that we study in this paper. The *Principal Component Regression* (PCR) estimator works by first projecting the input data into the r -dimensional principal subspace of the covariance matrix \hat{C} , and then ordinary least squares are solved for such projected data, yielding the estimator [see e.g. 24]

$$\hat{G}_{r,\gamma}^{\text{PCR}} = \llbracket \hat{C}^{-1} \rrbracket_r \hat{T}. \quad (4)$$

Here $\hat{C}_\gamma := \hat{C} + \gamma I_{\mathcal{H}}$ and $\llbracket \cdot \rrbracket_r$ denotes the r -truncated SVD. The population counterpart is $G_{r,\gamma}^{\text{PCR}} = \llbracket C_\gamma^{-1} \rrbracket_r T$, where $C_\gamma := C + \gamma I_{\mathcal{H}}$. Note, however, that the empirical PCR estimator does *not* minimize the empirical risk (3) under the low-rank constraint.

The *Reduced Rank Regression* (RRR) algorithm, in contrast, is the *exact* minimizer of (3) under fixed rank constraint. Specifically, RRR is defined as $\hat{G}_{r,\gamma}^{\text{RRR}} := \arg \min \{ \hat{\mathcal{R}}(G) + \gamma \|G\|_{\text{HS}}^2 : G \in \text{B}_r(\mathcal{H}) \}$, where the *regularization* term $\gamma \|G\|_{\text{HS}}^2$ is added to ensure stability, and $\text{B}_r(\mathcal{H})$ denotes the set of bounded operators on \mathcal{H} that have rank at most r . The closed form solution of the empirical RRR estimator is [24]

$$\hat{G}_{r,\gamma}^{\text{RRR}} = \hat{C}_\gamma^{-1/2} \llbracket \hat{C}_\gamma^{-1/2} \hat{T} \rrbracket_r, \quad (5)$$

while the population counterpart is given by $G_{r,\gamma}^{\text{RRR}} = C_\gamma^{-1/2} \llbracket C_\gamma^{-1/2} T \rrbracket_r$.

Once either the PCR or RRR estimators are fitted, their spectral decomposition is a proxy for the spectral decomposition of the Koopman operator A_π . Theorem 2 in [24] shows how such a decomposition can be calculated via the kernel trick for both $\hat{G}_{r,\gamma}^{\text{PCR}}$ and $\hat{G}_{r,\gamma}^{\text{RRR}}$.

¹For simplicity we consider the i.i.d. setting, however our forthcoming analysis is directly applicable to sample trajectories following [24].




Figure 1: PCR vs. RRR in estimating the largest eigenvalues of the 1D Ornstein–Uhlenbeck process with three different kernels over 50 independent trials. Vertical lines correspond to Koopman eigenvalues. The *good* kernel is such that its \mathcal{H} corresponds to the leading eigenspace of the Koopman operator, while the other two are spans of scaled and permuted eigenfunctions for which the distortion with respect to the original metric structure of A_π introduce slow (*bad* kernel) and fast (*ugly* kernel) spectral decay of the covariance.

3 The Problem and Main Result in a Nutshell

In this section we introduce the spectral estimation problem, outline our main results in a distilled form, and discuss some important implications. Recall the definition of Koopman operator (1) and its spectral decomposition (2). Given a rank r estimator $\hat{G} \in B_r(\mathcal{H})$ of A_π , we let $(\hat{\lambda}_i, \hat{\psi}_i)_{i=1}^r$ be its spectral decomposition, satisfying $\hat{G}\hat{\psi}_i = \hat{\lambda}_i \hat{\psi}_i$. We aim to study how well a nonzero eigenvalue $\hat{\lambda}_i$ of \hat{G} estimates its *closest* Koopman eigenvalue $\mu_{j(i)}$, where

$$j(i) = \operatorname{argmin}_{j \in \mathbb{N}} |\hat{\lambda}_i - \mu_j|. \quad (6)$$

Moreover we wish to compare $\hat{\psi}_i$ with the corresponding true Koopman eigenfunction. To this end, we embed $\hat{\psi}_i$ in $L_\pi^2(\mathcal{X})$ by means of the operator S and define the normalized estimated eigenfunction

$$\hat{f}_i = S\hat{\psi}_i / \|S\hat{\psi}_i\|. \quad (7)$$

One of the key quantities studied in this work is the eigenvalue estimation error $|\hat{\lambda}_i - \mu_{j(i)}|$, $i \in [r]$. Recalling that A_π is compact and self-adjoint, the classical Davis-Kahan result [17] implies that the eigenvalue estimation error $|\hat{\lambda}_i - \mu_{j(i)}|$ also bounds the quality of the eigenfunction approximation as

$$\|\hat{f}_i - f_{j(i)}\|^2 \leq \frac{2|\hat{\lambda}_i - \mu_{j(i)}|}{[\operatorname{gap}_{j(i)}(A_\pi) - |\hat{\lambda}_i - \mu_{j(i)}|]_+} \quad (8)$$

where $\operatorname{gap}_j(A_\pi) = \min_{\ell \neq j} |\mu_\ell - \mu_j|$ is the distance between μ_j and its closest Koopman eigenvalue.

Let $\sigma_j(\cdot)$ denotes the j -th singular value of an operator. To give a flavour of our results, here we report spectral bounds for the Gaussian kernel. In this case, Theorem 3 below gives a high probability bound on the estimation error $|\hat{\lambda}_i - \mu_{j(i)}|$, that is of order

$$\mathcal{O}\left(\frac{\sigma_{r+1}(A_\pi S)}{\sigma_r(A_\pi S)} + \frac{1}{\sqrt{n}}\right) \text{ for } \hat{G}_{r,\gamma}^{\text{RRR}}, \text{ and } \mathcal{O}\left(\frac{\sigma_{r+1}(S)}{[\sigma_r(A_\pi S) - \sigma_{r+1}(S)]_+} + \frac{1}{\sqrt{n}}\right) \text{ for } \hat{G}_{r,\gamma}^{\text{PCR}}.$$

If the Koopman operator has finite rank then $\sigma_{r+1}(A_\pi S) = 0$, the RRR estimator is unbiased, and its error goes to zero at the rate $1/\sqrt{n}$. Otherwise, recalling that $\sigma_{r+1}(S)$ is the square root of the $(r+1)$ -th eigenvalue of the kernel operator [38, Chapter 4.5], if \mathcal{H} is infinite dimensional $\sigma_{r+1}(S) > 0$, i.e. PCR has a strictly positive bias. In general, the presence of a bias in the estimated eigenvalues may result in the appearance of *spurious eigenvalues*. This phenomenon for PCR is well documented in practice, see e.g. [13, 14, 26, 28]. In Figure 1 we illustrate such an effect on a simple dynamical system discussed both in Example 3 and in Section 7.

4 Approach

The core of our analysis is Theorem 1. It reveals that in order to derive spectral estimation bounds for the Koopman operator, it is not enough to study the excess risk in the HS norm. Indeed, our spectral

bounds are determined by both the *operator norm error* of the Koopman estimator

$$\mathcal{E}(\widehat{G}) := \|A_\pi S - S\widehat{G}\|, \quad \widehat{G} \in \text{HS}(\mathcal{H}) \quad (9)$$

and the *metric distortion* between \mathcal{H} and $L_\pi^2(\mathcal{X})$,

$$\eta(h) := \|h\| / \|Sh\|, \quad h \in \mathcal{H}. \quad (10)$$

Note that since $Sh \in L_\pi^2(\mathcal{X})$ is just an equivalence class of a function h , $\|Sh\|$ is simply $L_\pi^2(\mathcal{X})$ -norm of h , and hence the metric distortion can be written, with a slight abuse of notation, as $\eta(h) := \|h\|_{\mathcal{H}} / \|h\|_{L_\pi^2(\mathcal{X})}$. While the (HS norm) error was studied before [see 29, and references therein], little is known about operator norm error bounds. Moreover, the metric distortion is, to the best of our knowledge, a novel quantity in the spectral analysis of Koopman operator.

Theorem 1. *Let A_π be a self-adjoint compact operator and let $r \in \mathbb{N}$. Then, for every empirical estimator $\widehat{G} \in B_r(\mathcal{H})$ and every $i \in [r]$*

$$|\widehat{\lambda}_i - \mu_{j(i)}| \leq \eta(\widehat{\psi}_i) \mathcal{E}(\widehat{G}), \quad \text{and} \quad \|\widehat{f}_i - f_{j(i)}\|^2 \leq \frac{2\eta(\widehat{\psi}_i) \mathcal{E}(\widehat{G})}{[\text{gap}_{j(i)}(A_\pi) - \eta(\widehat{\psi}_i) \mathcal{E}(\widehat{G})]_+}. \quad (11)$$

Proof Sketch. First, note that for compact self-adjoint operators $|\widehat{\lambda}_i - \mu_{j(i)}| \leq \|(A_\pi - \widehat{\lambda}_i I)^{-1}\|^{-1}$. So, following the reasoning of [24, Theorem 1] and observing that $\|(A_\pi S - S\widehat{G})\widehat{\psi}_i\| / \|S\widehat{\psi}_i\| \leq \mathcal{E}(\widehat{G})\eta(\widehat{\psi}_i)$ gives the right hand side of the first equation in (11). Next, since additionally $\|(A_\pi S - S\widehat{G})\widehat{\psi}_i\| / \|S\widehat{\psi}_i\| \leq \mathcal{E}(\widehat{G})\eta(\widehat{\psi}_i)$, we can apply the Davis-Kahan spectral perturbation result for compact self-adjoint operators (Proposition 2, Appendix C) to bound $\sin(\widehat{\theta})$, where $\widehat{\theta}_i := \angle(\widehat{f}_i, f_{j(i)})$. The claim then follows since $\|\widehat{f}_i - f_{j(i)}\|^2 \leq 2(1 - \cos(\widehat{\theta}_i)) \leq 2\sin(\widehat{\theta}_i)$. The full proof can be found in Appendix C. \square

Note that the error (9), at least for universal kernels, can be made arbitrary small, see [24, Proposition 1]. Still, the metric distortion may dominate the error and, since the bound (11) is tight, one may have that the operator is well estimated in norm, but the estimated eigenpairs are far from the true ones. This phenomenon is at the origin of *spurious eigenvalues*. The proposed way to detect them for *deterministic* systems in [14] is to check if eigenvalue equations are satisfied empirically, which, however, is not useful for *stochastic* systems, see Rem. 4 of Appendix C.

Spuriousness may also originate from poor conditioning of the true eigenvalues, i.e. when the angle between true left and right eigenfunctions is small. Here, however, we assume $A_\pi = A_\pi^*$, so that we restrict ourselves to the case in which the *only source of spuriousness is due to the learning method*.

While we defer the discussion of the operator norm error to the next section, the following result bound the metric distortion; the proof can be found in Appendix C.

Proposition 1. *Let $\widehat{G} \in B_r(\mathcal{H})$. For all $i \in [r]$ the metric distortion of $\widehat{\psi}_i$ can be tightly bounded as*

$$1 / \sqrt{\|C\|} \leq \eta(\widehat{\psi}_i) \leq \min(|\widehat{\lambda}_i| \text{cond}(\widehat{\lambda}_i), \|\widehat{G}\|) / \sigma_{\min}^+(S\widehat{G}), \quad (12)$$

where $\text{cond}(\widehat{\lambda}_i) := \|\widehat{\xi}_i\| \|\widehat{\psi}_i\| / |\langle \widehat{\psi}_i, \widehat{\xi}_i \rangle|$ is the condition number of $\widehat{\lambda}_i$, and $\widehat{\xi}_i$ is its left eigenfunction.

The upper bound (12) depends on the estimator's eigenvalues and their conditioning. Notice that while the true eigenvalues of A_π have condition number one, the conditioning of the estimated ones depends on the choice of the kernel. Moreover, the upper bound can be controlled by tuning the estimator rank. Since the bound is tight (see Rem. 3 in Appendix C), the metric distortion can grow with the rank of the estimator, further motivating the use of low-rank estimators of A_π in practice, see [26].

We end this section by introducing an empirical estimator of the metric distortion $\eta(\widehat{\psi}_i)$, given by

$$\widehat{\eta}_i := \|\widehat{\psi}_i\| / \sqrt{\langle \widehat{C}\widehat{\psi}_i, \widehat{\psi}_i \rangle}. \quad (13)$$

Proposition 4 of Appendix C shows that $\widehat{\eta}_i$ can be efficiently computed and report upper bounds for concentration around its mean. The empirical metric distortion (13), used in conjunction with the spectral bounds in Theorem 4 below, provides a proxy to assess the reliability of the PCR and RRR estimators and can be successfully used as novel model selection criterion. We refer the reader to the second and third experiment in Section 7 for concrete use-cases.

5 Controlling the Operator Norm Error

The HS norm error of the PCR estimator was already studied, either in the "well-specified" setting, i.e. when there exists $G_{\mathcal{H}} \in \text{HS}(\mathcal{H})$ such that $A_{\pi}S = SG_{\mathcal{H}}$, i.e. $G_{\mathcal{H}}$ is π -a.e. Koopman operator [12, Theorem B.10]. On the other hand, KRR estimator is studied also in the "misspecified setting" [29]. But, up to our knowledge, the operator norm error has not yet been studied. To analyse these learning rates, we make the following assumptions:

(RC) *Regularity of A_{π} .* For some $\alpha \in (0, 2]$ there exists $a > 0$ such that $TT^* \preceq a^2 C^{1+\alpha}$;

(BK) *Boundedness.* There exists $c_{\mathcal{H}} > 0$ such that $\text{ess sup}_{x \sim \pi} \|\phi(x)\|^2 \leq c_{\mathcal{H}}$, i.e. $\phi \in L_{\pi}^{\infty}(\mathcal{X}, \mathcal{H})$;

(SD) *Spectral Decay.* There exists $\beta \in (0, 1]$ and $b > 0$ such that $\lambda_j(C) \leq b j^{-1/\beta}$, for all $j \in J$.

While we keep assumptions **(BK)** and **(SD)** as in [18, 29], assumption **(RC)** is, up to our knowledge, novel. The rationale behind it is that for $\alpha = 1$ **(RC)** is equivalent to $\text{Im}(A_{\pi}S) \subseteq \text{Im}(S)$, in which case there exists a bounded π -a.e. Koopman operator $G_{\mathcal{H}}: \mathcal{H} \rightarrow \mathcal{H}$ [24]. On the other hand, as $\alpha \rightarrow 0$ **(RC)** becomes closer to $\text{Im}(A_{\pi}S) \subseteq \text{cl}(\text{Im}(S))$ which is always satisfied for universal kernels since $\text{cl}(\text{Im}(S)) = L_{\pi}^2(\mathcal{X})$ [38, Chapter 4]. Importantly, as the next example shows, **(RC)** is weaker condition than the usual regularity conditions; see Appendix D.1 for a detailed discussion.

Example 2. Let X be an \mathcal{X} -valued random variable with law π . Consider the Markov chain $(X_t)_{t \in \mathbb{N}}$ such that $X_t = X$ for all $t \in \mathbb{N}$. Then π is an invariant measure and $A_{\pi} = I_{L_{\pi}^2(\mathcal{X})}$ is the identity map on $L_{\pi}^2(\mathcal{X})$. Clearly, **(RC)** holds for all $\alpha \in (0, 1]$. On the other hand, since $A_{\pi}S = SG_{\mathcal{H}}$ for bounded operator $G_{\mathcal{H}} = I_{\mathcal{H}} \notin \text{HS}(\mathcal{H})$, HS-norm learning rates derived in [29] do not apply.

In order to study the error of any empirical finite rank estimator \hat{G} we rely on the error decomposition

$$\mathcal{E}(\hat{G}) \leq \underbrace{\|A_{\pi}S - SG_{\gamma}\|}_{\text{regularization bias}} + \underbrace{\|S(G_{\gamma} - G)\|}_{\text{rank reduction bias}} + \underbrace{\|S(G - \hat{G})\|}_{\text{estimator's variance}}, \quad (14)$$

where $G_{\gamma} := C_{\gamma}^{-1}T$ is the minimizer of the full (i.e. without rank constraint), Tikhonov regularized, HS norm error, and G is the population version of the empirical estimator \hat{G} .

While the last two terms in the r.h.s. of (14) depend of the estimator of choice, the first term depends only on the choice of \mathcal{H} and the regularity of A_{π} w.r.t. \mathcal{H} . In this work we focus on the classical kernel-based learning of the Koopman operator [20, 24, 29], where one chooses a universal kernel [38, Chapter 4] for which $\text{Im}(A_{\pi}S) \subseteq \text{cl}(\text{Im}(S))$, and controls the regularization bias with a regularity condition. For details see Rem. 7 of Appendix D.2.

The second source of bias and the estimator's variance in our error decomposition depends on the choice of the low rank estimator. While throughout this section we consider **(RC)** for $\alpha \in [1, 2]$, we discuss extensions of our results to $\alpha < 1$ in Appendix D.5.

Theorem 2. Assume the operator A_{π} satisfies $\sigma_r(A_{\pi}S) > \sigma_{r+1}(A_{\pi}S) \geq 0$ for some $r \in \mathbb{N}$. Let **(SD)** and **(RC)** hold for some $\beta \in (0, 1]$ and $\alpha \in [1, 2]$, respectively, and let $\text{cl}(\text{Im}(S)) = L_{\pi}^2(\mathcal{X})$. Let

$$\gamma \asymp n^{-\frac{1}{\alpha+\beta}} \text{ and } \varepsilon_n^* := n^{-\frac{\alpha}{2(\alpha+\beta)}}. \quad (15)$$

Let $\delta \in (0, 1)$. Then, there exists a constant $c > 0$, depending only on \mathcal{H} , such that for large enough $n \geq r$, with probability at least $1 - \delta$ in the i.i.d. draw of \mathcal{D}_n from ρ

$$\mathcal{E}(\hat{G}) \leq \begin{cases} \sigma_{r+1}(A_{\pi}S) + c\varepsilon_n^* \ln \delta^{-1} & \text{if } \hat{G} = \hat{G}_{r,\gamma}^{\text{RRR}}, \end{cases} \quad (16a)$$

$$\begin{cases} \sigma_{r+1}(S) + c\varepsilon_n^* \ln \delta^{-1} & \text{if } \hat{G} = \hat{G}_{r,\gamma}^{\text{PCR}} \text{ and } \sigma_r(S) > \sigma_{r+1}(S). \end{cases} \quad (16b)$$

Proof Sketch. The regularization bias is bounded by $a\gamma^{\frac{\alpha}{2}}$ by Proposition 5 of Appendix D.2. For the RRR estimator, the rank reduction bias is upper bounded by $\sigma_{r+1}(A_{\pi}S)$, while for PCR by $\sigma_{r+1}(S)$. The bounds on the variance terms critically rely on the well-known perturbation result for spectral projectors reported in Proposition 3, Appendix A. This result is then chained to two versions of the Bernstein inequality in separable Hilbert spaces. The first one is Pinelis-Sakhanenko's inequality and the second is Minsker's inequality extended to self-adjoint HS-operators, Props. 9 and 11 in Appendix D.3.1, respectively. These inequalities provide high probability bounds for the norms

of $C_\gamma^{-1/2}(\hat{C} - C)$ and $C_\gamma^{-1/2}(\hat{C} - C)C_\gamma^{-1/2}$, as well as $C_\gamma^{-1/2}(\hat{T} - T)$ and $C_\gamma^{-1/2}(\hat{T} - T)C_\gamma^{-1/2}$. Combining the bias due to regularization and variance terms, for both estimators we obtain the balancing equation $\gamma^{\frac{\alpha}{2}} = \gamma^{-\frac{\beta}{2}} n^{-\frac{1}{2}} \ln \delta^{-1}$, which yields the optimal choice of γ and the rates. \square

We stress that the number of samples in the previous theorem depends on the problem's complexity, expressed in the constants $c_{\text{RRR}} = \frac{1}{\sigma_r^2(A_\pi S) - \sigma_{r+1}^2(A_\pi S)}$, and $c_{\text{PCR}} = \frac{1}{\sigma_r(S) - \sigma_{r+1}(S)}$. Namely, the better the separation of singular values, the smaller number of needed samples. Furthermore, analyzing the bounds (16a) and (16b), we see that faster spectral decay is, in general, preferable. For example, for the Gaussian kernel β can be chosen arbitrarily small, yielding the rate $n^{-1/2}$. On the other hand, kernels with slow spectral decay for which $\beta = 1$ can give slower rates between $n^{-1/4}$ and $n^{-1/3}$. Finally, from the variance bounds for RRR and PCR, c.f. Appendix D.3, one can specify constants. Namely, in the slower regime when $\varepsilon_n^* > n^{-1/2}$ we have that $c = a + 7.2 \log(10) \sqrt{2c_{\mathcal{H}}} (1 + ac_{\mathcal{H}}^{(\alpha-1)/2}) (\sqrt{c_{\mathcal{H}}} \wedge \frac{b^{\beta/2}}{\sqrt{1-\beta}})$, while in fastest regime $\varepsilon_n^* = n^{-1/2}$, there is a significant difference between RRR and PCR since c should be multiplied with the constants c_{RRR} and c_{PCR} , respectively.

As argued in Section 3, the bound (16a) indicates that for rank r Koopman operators the error converges to zero w.r.t. the number of training samples, while the bias of PCR is strictly positive. Hence, in order to ensure small error for PCR, high values of the rank parameter might be necessary. To theoretically explain this effect, in Theorem 6 of Appendix D.4 we give also lower bounds of operator norm error for the RRR and PCR estimators showing that $\mathcal{E}(\hat{G}_{r,\gamma}^{\text{RRR}})$ always concentrates around $\sigma_{r+1}(A_\pi S)$, while the concentration of $\mathcal{E}(\hat{G}_{r,\gamma}^{\text{PCR}})$ around $\sigma_{r+1}(S)$ depends on the *irreducible risk* of the learning problem. To illustrate the tightness of the error concentration bounds we present Example 3 (see also Appendix D.4).

Example 3. Let $\mathcal{X} = \mathbb{R}$. Consider the 1D equidistant sampling of the Ornstein–Uhlenbeck process, obtained by integrating the Langevin equation of Example 1 with $\beta = 1$ and $U(x) = x^2/2$, given by $X_t = e^{-1}X_{t-1} + \sqrt{1-e^{-2}}\epsilon_t$, where $\{\epsilon_t\}_{t \geq 1}$ are i.i.d. standard Gaussians. For this process it is well-known [34] that π is $\mathcal{N}(0, 1)$ and that A_π admits a spectral decomposition $(\mu_i, f_i)_{i \in \mathbb{N}}$ in terms of Hermite polynomials. We study the family of kernel functions $k_{\Pi, \nu}(x, x') := \sum_{i \in \mathbb{N}} \mu_{\Pi(i)}^{2\nu} f_i(x) f_i(x')$, where Π is a permutation of the indices of the eigenvalues and ν is a scaling factor. The rationale behind this class of kernels is that by varying Π and ν one morphs the original metric structure of A_π in a way which is harder and harder to revert when learning from finite sets of data. In particular, for any target rank r , setting $\nu := 1/r^2$ and Π to the permutation such that $i \mapsto 2r - i + 1$ ($i \leq r$), $i \mapsto i - r$ ($r + 1 \leq i \leq 2r$) and $i \mapsto i$ elsewhere, elementary algebra and our concentration bounds give

$$|\mathcal{E}(\hat{G}_{r,\gamma}^{\text{PCR}}) - e^{-1/r}| \lesssim n^{-1/2} \ln \delta^{-1}, \quad |\mathcal{E}(\hat{G}_{r,\gamma}^{\text{RRR}}) - e^{-r}| \lesssim n^{-1/2} \ln \delta^{-1}.$$

We refer the reader to Figure 1 and to Section 7 for a numerical implementation of this example.

We conclude this section with remarks on the tightness of our statistical analysis of operator norm error. Since discussed results are not the main focus of the paper, we present them in Appendix D.5.

Remark 1 (Lower bound). The rate $\varepsilon_n^* = n^{-\frac{\alpha}{2(\alpha+\beta)}}$ guaranteed by (16a) matches the minimax lower bound for the operator norm error when learning finite rank A_π . Formal statement and its proof is given in Theorem 7 of Appendix D.5.

Remark 2 (Extension to misspecified setting). The optimal rates for HS-norm error of the KRR estimator are developed in [29] under a stronger condition than **(RC)**. In Theorem 9 of Appendix D.6 we extended this analysis to PCR and RRR estimators, deriving the optimal operator norm rates that also cover cases when Koopman operator cannot be properly defined as bounded operator on the chosen RKHS space \mathcal{H} .

6 Spectral Learning Rates

Collecting all the previous results, we are now ready to present our spectral learning rates for the two estimators in a general form. For brevity, we focus on two different type of bounds in which (i) we analyse the uniform bound for the whole estimated spectra, and (ii) we express the estimators' bias in *empirical form* to provide an insight into spuriousness of eigenvalues. Moreover, we present only

eigenvalue estimation bounds, noting that the eigenfunction estimation bounds readily follow from (8). The complete results are presented in detail in Appendix E.

Theorem 3. *Let A_π be a compact self-adjoint operator. Under the assumptions of Theorem 2, there exists a constant $c > 0$, depending only on \mathcal{H} , such that for every $\delta \in (0, 1)$, for every large enough $n \geq r$ and every $i \in [r]$ with probability at least $1 - \delta$ in the i.i.d. draw of \mathcal{D}_n from ρ*

$$|\hat{\lambda}_i - \mu_{j(i)}| \leq \begin{cases} \frac{2\sigma_{r+1}(A_\pi S)}{\sigma_r(A_\pi S)} + c\varepsilon_n^* \ln \delta^{-1} & \text{if } \hat{G} = \hat{G}_{r,\gamma}^{\text{RRR}}, \\ \frac{2\sigma_{r+1}(S)}{[\sigma_r(A_\pi S) - \sigma_{r+1}^\alpha(S)]_+} + c\varepsilon_n^* \ln \delta^{-1} & \text{if } \hat{G} = \hat{G}_{r,\gamma}^{\text{PCR}}. \end{cases} \quad (17)$$

The uniform eigenvalue learning rates for RRR and PCR estimators, differ in the estimator's bias. While the PCR bias has a factor $\sigma_{r+1}(S)$ in the numerator, RRR has $\sigma_{r+1}(A_\pi S) \leq \sigma_{r+1}(S)$. The striking difference happens when the Koopman operator is of finite rank. Then, assuming that r is properly chosen, RRR estimator has no bias, and c_{RRR} is typically moderate. On the other hand, PCR's bias can be potentially large, depending of the choice of the kernel, and choosing higher rank increases c_{PCR} , thus requiring larger sample sizes. Therefore, even in well-conditioned problems (self-adjoint operator) the *spurious eigenvalues may arise purely from the learning method*. To facilitate detection of such occurrences, we further provide an empirical estimator of the bias of both methods and illustrate their use experimentally in Section 7.

Theorem 4. *Under the assumptions of Theorems 2 and 3, there exists a constant $c > 0$, depending only on \mathcal{H} , such that for large enough $n \geq r$ and every $i \in [r]$ with probability at least $1 - \delta$ in the i.i.d. draw of \mathcal{D}_n from ρ*

$$|\hat{\lambda}_i - \mu_{j(i)}| \leq \begin{cases} \hat{\eta}_i \sigma_{r+1}(\hat{C}^{-1/2} \hat{T}) + c\varepsilon_n^* \ln \delta^{-1}, & \hat{G} = \hat{G}_{r,\gamma}^{\text{RRR}}, \\ \hat{\eta}_i \sqrt{\sigma_{r+1}(\hat{C})} + c\varepsilon_n^* \ln \delta^{-1}, & \hat{G} = \hat{G}_{r,\gamma}^{\text{PCR}}. \end{cases} \quad (18)$$

We remark that when A_π is of finite rank r , the bound above for the RRR estimator reduces to

$$|\hat{\lambda}_i - \mu_i| \leq c\varepsilon_n^* \ln \delta^{-1} \quad \text{and} \quad \|\hat{f}_i - f_i\|^2 \leq \frac{2c\varepsilon_n^* \ln \delta^{-1}}{[\text{gap}_i(\hat{G}_{r,\gamma}^{\text{RRR}}) - 3c\varepsilon_n^* \ln \delta^{-1}]_+},$$

see Cor. 1 in Appendix E. Hence, in this case RRR algorithm can learn all the eigenvalues and eigenfunctions of A_π with rate $\varepsilon_n^* = n^{-\frac{\alpha}{2(\alpha+\beta)}}$. On the other hand, even in this case, the bounds for the PCR estimator do not guarantee unbiased estimation of Koopman eigenvalues and eigenfunctions.

Choosing γ and r . The bias term $\sigma_{r+1}(A_\pi S)/\sigma_r(A_\pi S)$ appearing in (17) represents the theoretical limit when estimating eigenvalues using RRR. It reflects the capacity of the RKHS to detect the separation of the leading r Koopman eigenvalues from the rest of its spectra. If A_π has infinite rank and slowly decaying eigenvalues, estimating the leading ones becomes challenging, since increasing r leads to smaller operator norm error, but larger bias. Luckily, in many practical problems there is a separation of time-scales in the dynamics and the above ratio can be controlled by choosing r appropriately. While we do not have access to $A_\pi S$, we can still choose r via the empirical operator $\hat{C}^{-1/2} \hat{T}$, see Proposition 20 of Appendix D.4. Note also that the optimal γ depends on α which is typically unknown. In practice, one can implement a standard grid-search CV procedure for time series to tune this parameter.

Spectral Bias as a Tool for Model Selection. In equation (18), the data dependent quantities $\hat{s}_i(\hat{G}_{r,\gamma}^{\text{RRR}}) := \hat{\eta}_i \sigma_{r+1}(\hat{C}^{-1/2} \hat{T})$ and $\hat{s}_i(\hat{G}_{r,\gamma}^{\text{PCR}}) := \hat{\eta}_i \sigma_{r+1}(\hat{C})$ represent the *empirical spectral biases* of RRR and PCR estimators of the Koopman operator, respectively. When they are small enough, the spectral estimation error is dominated by the same variance term, which decreases as the number of samples grows. Therefore, given a number of different kernels, we propose to select the best one (w.r.t. spectral estimation) by choosing the smallest spectral bias. This is illustrated in the Alenine Dipeptide example of following section.

Normal operators. Since Davis-Kahan theorem [17] also holds for normal operators, the results in this section apply whenever $A_\pi A_\pi^* = A_\pi^* A_\pi$. While in this case Koopman eigenfunctions remain orthogonal in $L_\pi^2(\mathcal{X})$, the eigenvalues are in general complex. On the other hand, extension beyond normal compact operators asks for involved spectral perturbation analysis and a new statistical learning theory.




Figure 2: Estimated eigenfunctions $\hat{\psi}_i$ of a Langevin dynamics vs. ground truth. The average empirical biases \hat{s}_i , $i \in [4]$ are discussed at the end of Section 6. The results correspond to 50 independent estimations on 2000 training points each. PCR and RRR estimators were fitted with the same parameters: Gaussian kernel of length scale 0.175, $\gamma = 10^{-5}$ and $r = 4$.

7 Experiments

We illustrate various aspects of our theory with simple experiments. They have been implemented in Python using the library Kooplearn (available at <https://github.com/CSML-IIT-UCL/kooplearn>) to fit the PCR and RRR estimators. Full details are in Appendix F.

Learning the Spectrum of the Ornstein–Uhlenbeck Process. In this experiment we designed three different kernel functions (the “good”, the “bad” and the “ugly”) to illustrate how an unseemly kernel choice can induce catastrophic biases in the estimation of Koopman eigenvalues. We focus on the uniformly sampled Ornstein-Uhlenbeck (OU) process, discussed in Example 3, relying on the spectral decomposition of its Koopman operator $(\mu_i, f_i)_{i \in \mathbb{N}}$ to design the three kernel functions. The *good* kernel is just the sum of the leading $T = 53$ terms of the spectral decomposition of A_π , i.e. $k_{\text{good}}(x, y) := \sum_{i=1}^T \mu_i f_i(x) f_i(y)$. The associated RKHS coincides with the leading eigenspace of A_π , and *no deformation of the metric structure* takes place, so that the injection map $S: \mathcal{H} \hookrightarrow L_\pi^2(\mathcal{X})$ is a partial isometry. The *bad* kernel is defined according to the construction presented in Example 3 for $\nu = 1/r^2$ where r is the rank of the estimator. For this kernel, the introduced bias is innocuous for RRR, but lethal for PCR. Finally, the *ugly* kernel corresponds to $\nu = r^2$, introducing large quotients $\sigma_{r+1}(A_\pi S)/\sigma_r(A_\pi S)$ and $\sigma_{r+1}(S)/\sigma_r(S)$, and, hence, an irreparable bias in both estimators.

Figure 1 depicts the distribution of the eigenvalues estimated by PCR and RRR over 50 independent simulations, against the ground truth. For both algorithms each simulation is comprised of 20000 training points, the regularization is $\gamma = 10^{-4}$ and the rank is $r = 3$. The three largest eigenvalues of A_π are correctly estimated by both algorithms for k_{good} and by RRR for k_{bad} . On the contrary, the distribution of the eigenvalues for k_{ugly} (and k_{bad} for PCR) does not concentrate around any true eigenvalue of A_π , signaling the presence of *spurious eigenvalues* in the estimation.




Figure 3: Forecasting RMSE on the Alanine Dipeptide dataset for 19 different RRR estimators, each corresponding to a different kernel, which show how the best model, according to the empirical spectral bias metric, also attains the best forecasting performances by a large margin.

A Realistic Example: Langevin Dynamics. Because of its ubiquitous use in modelling real systems, we now study a numerical implementation of the Langevin dynamics Example 1 with $\beta = 1$ and a potential $U(x) = 4(x^8 + 0.8e^{-80x^2} + 0.2e^{-80(x-0.5)^2} + 0.5e^{-40(x+0.5)^2})$ that is a mixture of three Gaussians barriers at $x \in \{-0.5, 0, 0.5\}$ and a smooth “bounding” term $\propto x^8$ constraining most of the equilibrium distribution in the interval $[-1, 1]$, see [37]. In Figure 2, for $i \in [r]$, we compare the \hat{f}_i estimated by PCR and RRR against the ground truth f_i . The visible difficulty of PCR compared to RRR in estimating eigenfunctions is nicely explained by larger values of the empirical bias for PCR, which we report in the upper part of the figure. The reference eigenpairs of A_π have been obtained by diagonalizing a finely discretized approximation of the infinitesimal generator (see Appendix A).

Spectral Bias and Model Selection: the Case of Alanine Dipeptide. In this example we show that minimizing the first term on the r.h.s. of (18) over a validation dataset, is also a good criterion for Koopman model selection. We use a realistic simulation of the small molecule Alanine Dipeptide already discussed in [24, 44]. We trained 19 RRR estimators each corresponding to a different kernel and then we evaluated the forecasting RMSE on 2000 initial conditions drawn from a test dataset. In Figure 3 we report these errors, highlighting the model with the smallest average empirical spectral bias (18) evaluated on 5000 validation points.

8 Conclusion

We established minimax optimal rates for the operator norm error in the Koopman regression problem, which we then used to derive sharp estimation bounds for eigenvalues and eigenfunctions of the Koopman operator associated with a time-invariant Markov chain. We considered two important estimators that implement either principal component regression (PCR) or reduced rank regression (RRR) to learn a linear operator on a reproducing kernel Hilbert space. Our bounds indicate that RRR may be advantageous over PCR (also known as EDMD, the de-facto estimator in the data-driven dynamical system literature) which may exhibit a larger estimation bias. This ultimately depends on the choice of the kernel, which significantly impacts the rate. A bad choice of the kernel could also introduce spurious eigenvalues, a phenomena which has been observed in the literature and which is now explained by our theory. Finally, we proposed a method to detect spuriousness in practice, which can be used also as a kernel selection tool. A limitation of this work is that it applies to compact normal operators only. While many real dynamical systems involve such operators, in the future our analysis may be extended using more sophisticated spectral perturbation theory.

Acknowledgements. This work was supported in part from the PNRR MUR Project PE000013 CUP J53C22003010006 “Future Artificial Intelligence Research (FAIR)”, funded by the European Union – NextGenerationEU, and EU Project ELIAS under grant agreement No. 101120237.

References

- [1] Amir, A. (2014). Cell size regulation in bacteria. *Physical Review Letters*, 112(20).
- [2] Arbabi, H. and Mezić, I. (2017). Ergodic theory, dynamic mode decomposition, and computation of spectral properties of the koopman operator. *SIAM Journal on Applied Dynamical Systems*, 16(4):2096–2126.
- [3] Aronszajn, N. (1950). Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3):337–404.
- [4] Bakry, D., Gentil, I., and Ledoux, M. (2014). *Analysis and Geometry of Markov Diffusion Operators*. Springer.
- [5] Bonati, L., Piccini, G., and Parrinello, M. (2021). Deep learning the slow modes for rare events sampling. *Proceedings of the National Academy of Sciences*, 118(44).
- [6] Bouchaud, J.-P. and Cont, R. (1998). A langevin approach to stock market fluctuations and crashes. *The European Physical Journal B*, 6(4):543–550.
- [7] Bruder, D., Fu, X., Gillespie, R. B., Remy, C. D., and Vasudevan, R. (2021). Data-Driven Control of Soft Robots Using Koopman Operator Theory. *IEEE Transactions on Robotics*, 37(3):948–961.

- [8] Brunton, S. L., Budišić, M., Kaiser, E., and Kutz, J. N. (2022). Modern Koopman Theory for Dynamical Systems. *SIAM Review*, 64(2):229–340.
- [9] Budišić, M., Mohr, R., and Mezić, I. (2012). Applied Koopmanism. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 22(4):047510.
- [10] Caponnetto, A. and De Vito, E. (2007). Optimal rates for the regularized least-squares algorithm. *Foundations of Computational Mathematics*, 7(3):331–368.
- [11] Choquet, D. and Triller, A. (2013). The dynamic synapse. *Neuron*, 80(3):691–703.
- [12] Ciliberto, C., Rosasco, L., and Rudi, A. (2020). A general framework for consistent structured prediction with implicit loss embeddings. *Journal of Machine Learning Research*, 21(98):1–67.
- [13] Colbrook, M. J., Roman, B., and Hansen, A. C. (2019). How to compute spectra with error control. *Physical Review Letters*, 122(25).
- [14] Colbrook, M. J. and Townsend, A. (2021). Rigorous data-driven computation of spectral properties of koopman operators for dynamical systems.
- [15] Da Prato, G. and Zabczyk, J. (1996). *Ergodicity for Infinite Dimensional Systems*. London Mathematical Society Lecture Note Series. Cambridge University Press.
- [16] Davidchack, R. L., Ouldridge, T. E., and Tretyakov, M. V. (2015). New langevin and gradient thermostats for rigid body dynamics. *The Journal of Chemical Physics*, 142(14):144114.
- [17] Davis, C. and Kahan, W. M. (1970). The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46.
- [18] Fischer, S. and Steinwart, I. (2020). Sobolev norm learning rates for regularized least-squares algorithms. *Journal of Machine Learning Research*, 21(205):1–38.
- [19] Folkestad, C., Chen, Y., Ames, A. D., and Burdick, J. W. (2021). Data-Driven Safety-Critical Control: Synthesizing Control Barrier Functions with Koopman Operators. *IEEE Control Systems Letters*, 5(6):2012–2017.
- [20] Grünewälder, S., Lever, G., Gretton, A., Baldassarre, L., Patterson, S., and Pontil, M. (2012). Conditional mean embeddings as regressors. In *ICML*. icml.cc / Omnipress.
- [21] Héas, P. and Herzet, C. (2021). Low-rank dynamic mode decomposition: An exact and tractable solution. *Journal of Nonlinear Science*, 32(1).
- [22] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., and et al. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440.
- [23] Klus, S., Nüske, F., Koltai, P., Wu, H., Kevrekidis, I., Schütte, C., and Noé, F. (2018). Data-driven model reduction and transfer operator approximation. *Journal of Nonlinear Science*, 28(3):985–1010.
- [24] Kostic, V., Novelli, P., Maurer, A., Ciliberto, C., Rosasco, L., and Pontil, M. (2022). Learning dynamical systems via Koopman operator regression in reproducing kernel hilbert spaces. In *Advances in Neural Information Processing Systems*.
- [25] Kramers, H. (1940). Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4):284–304.
- [26] Kutz, J. N., Brunton, S. L., Brunton, B. W., and Proctor, J. L. (2016). *Dynamic Mode Decomposition*. Society for Industrial and Applied Mathematics.
- [27] Lasota, A. and Mackey, M. C. (1994). *Chaos, Fractals, and Noise*, volume 97 of *Applied Mathematical Sciences*. Springer New York.
- [28] Lewin, M. and Séré, É. (2009). Spectral pollution and how to avoid it. *Proceedings of the London Mathematical Society*, 100(3):864–900.

- [29] Li, Z., Meunier, D., Mollenhauer, M., and Gretton, A. (2022). Optimal rates for regularized conditional mean embedding learning. In *Advances in Neural Information Processing Systems*.
- [30] Mardt, A., Pasquali, L., Wu, H., and Noé, F. (2018). VAMPnets for deep learning of molecular kinetics. *Nature Communications*, 9(1).
- [31] Mauroy, A., Mezić, I., and Susuki, Y., editors (2020). *The Koopman Operator in Systems and Control*. Springer.
- [32] Meyn, S. P. and Tweedie, R. L. (1993). *Markov Chains and Stochastic Stability*. Communications and Control Engineering. Springer London.
- [33] Minsker, S. (2017). On some extensions of Bernstein’s inequality for self-adjoint operators. *Statistics & Probability Letters*, 127:111–119.
- [34] Pavliotis, G. A. (2014). *Stochastic Processes and Applications*. Springer New York.
- [35] Reed, M. and Simon, B. (1980). *I: Functional Analysis*. Academic Press.
- [36] Rowley, C. W., Mezić, I., Bagheri, S., Schlatter, P., and Henningson, D. S. (2009). Spectral analysis of nonlinear flows. *J. Fluid Mech.*, 641:115–127.
- [37] Schwantes, C. R. and Pande, V. S. (2015). Modeling Molecular Kinetics with tICA and the Kernel Trick. *Journal of Chemical Theory and Computation*, 11(2):600–608.
- [38] Steinwart, I. and Christmann, A. (2008). *Support Vector Machines*. Springer New York.
- [39] Tropp, J. A. (2012). User-friendly tail bounds for sums of random matrices. Technical report.
- [40] Tuckerman, M. (2010). *Statistical Mechanics: Theory and Molecular Simulation*. Oxford university press.
- [41] Varadhan, S. R. S. (2007). Additional remarks. In *Courant Lecture Notes*, pages 121–121. American Mathematical Society, Providence, Rhode Island.
- [42] Ventriglia, F. and Maio, V. D. (2000). A brownian model of glutamate diffusion in excitatory synapses of hippocampus. *Biosystems*, 58(1-3):67–74.
- [43] Wang, Z., Luo, Y., Li, Y., Zhu, J., and Schölkopf, B. (2022). Spectral representation learning for conditional moment models.
- [44] Wehmeyer, C. and Noé, F. (2018). Time-lagged autoencoders: Deep learning of slow collective variables for molecular kinetics. *The Journal of Chemical Physics*, 148(24):241703.
- [45] Williams, M. O., , Rowley, C. W., and Kevrekidis, I. G. (2015). A kernel-based method for data-driven Koopman spectral analysis. *Journal of Computational Dynamics*, 2(2):247–265.
- [46] Wu, H. and Noé, F. (2019). Variational Approach for Learning Markov Processes from Time Series Data. *Journal of Nonlinear Science*, 30(1):23–66.
- [47] Zabczyk, J. (2020). *Mathematical Control Theory: An Introduction*. Systems & Control: Foundations & Applications. Springer International Publishing.
- [48] Zwald, L. and Blanchard, G. (2005). On the Convergence of Eigenspaces in Kernel Principal Component Analysis. In *Advances in Neural Information Processing Systems*.

Supplementary Material

The supplementary material is organized as follows.

- Appendix A contains additional background on stochastic processes with self-adjoint Koopman operator, on Markov processes and spectral theory. Additionally, it contains a notation table.
- Appendix B discusses learning of the Koopman operator with kernel-based methods.
- Appendix C contains details for the content presented in Section 4, notably proving the key perturbation result of Theorem 1.
- Appendix D contains details of the content presented in Section 5. In particular, in Appendix D.1 we discuss the main assumptions and their relationship with the existing literature, in Appendix D.2 we prove the bounds of different bias terms: bias due to RKHS, bias due to estimator of choice and bias in the effective rank estimation, in Appendix D.3 we prove the bounds of the corresponding variance terms, and in Appendix D.4 we show the error bounds for the RRR and PCR estimators under stronger **(RC)** condition, while in Appendix D.6 we extend them and prove the matching lower bound in Theorem 7 in Appendix D.5.
- In Appendix E we prove the spectral learning rates of Section 6 in more detailed form.
- Finally, in Appendix F we provide more details on the experimental section, as well as present additional experiments.

A Notation and Background

First, we briefly recall the basic notions related to Markov chains and Koopman operators and refer to [27, 31, 32] for further details.

Let $\mathbf{X} := \{X_t : t \in \mathbb{N}\}$ be a family of random variables with values in a measurable space $(\mathcal{X}, \Sigma_{\mathcal{X}})$, called state space. We call \mathbf{X} a *Markov chain* if $\mathbb{P}\{X_{t+1} \in B \mid X_{[t]}\} = \mathbb{P}\{X_{t+1} \in B \mid X_t\}$. Further, we call \mathbf{X} *time-homogeneous* if there exists $p: \mathcal{X} \times \Sigma_{\mathcal{X}} \rightarrow [0, 1]$, called *transition kernel*, such that, for every $(x, B) \in \mathcal{X} \times \Sigma_{\mathcal{X}}$ and every $t \in \mathbb{N}$,

$$\mathbb{P}\{X_{t+1} \in B \mid X_t = x\} = p(x, B).$$

A large class of Markov chains consists of those who possess an *invariant measure* π satisfying $\pi(B) = \int_{\mathcal{X}} \pi(dx) p(x, B)$, $B \in \Sigma_{\mathcal{X}}$, see e.g. [15]. For those, we can consider the space of square integrable functions on \mathcal{X} relative to the measure π , denoted as $L^2_{\pi}(\mathcal{X})$, and define *Markov transfer operator*, i.e. *Stochastic Koopman operator*, $A_{\pi}: L^2_{\pi}(\mathcal{X}) \rightarrow L^2_{\pi}(\mathcal{X})$

$$A_{\pi}f(x) := \int_{\mathcal{X}} p(x, dy) f(y) = \mathbb{E}[f(X_{t+1}) \mid X_t = x], \quad f \in L^2_{\pi}(\mathcal{X}), x \in \mathcal{X}. \quad (19)$$

Since it is easy to see that $\|A_{\pi}\| = 1$, we conclude that the Markov transfer operator is a bounded linear operator.

This work focuses on the Markov chains that originate from a dynamical system that is time-reversal invariant, which as a consequence has that Koopman operator on the $L^2_{\pi}(\mathcal{X})$ space is self-adjoint. Since many microscopical equations of motion in both classical and quantum physics are time-reversal invariant, learning self-adjoint Koopman operators is of paramount importance in the field of machine learning for physical sciences.

Next, we discuss the (overdamped) Langevin equation

$$dX_t = -\nabla U(X_t)dt + \sqrt{2\beta^{-1}}dW_t,$$

where W_t is a Wiener process. For any $f \in C^2(\mathbb{R}^d)$ we let $u(x, t) := \mathbb{E}[f(X_t) \mid X_0 = x]$. As showed in Section 6.3, pp. 95-96 of Ref. [41], $u(x, t)$ is the solution of the *backward Kolmogorov equation*

$$\partial_t u = \mathcal{L}u \quad u(x, 0) = f(x), \quad (20)$$

where $(\mathcal{L}f)(x) := \beta^{-1}\nabla^2 f(x) - \nabla U(x) \cdot \nabla f(x)$. A straightforward calculation shows that \mathcal{L} is a self-adjoint operator with respect to the scalar product $\langle f, g \rangle := \int_{\mathcal{X}} f(x)g(x)\pi(dx)$ where $\pi(dx)$ is

notation	meaning	notation	meaning
\wedge	minimum	\vee	maximum
$[\cdot]$	set $\{1, 2, \dots, \cdot\}$	$[\cdot]_+$	nonegative part of a number
\mathcal{X}	state space of the Markov chain	$(X_t)_{t \in \mathbb{N}}$	time-homogeneous Markov chain
p	transition kernel of the Markov chain	π	invariant measure of the Markov chain
$L^2_\pi(\mathcal{X})$	L2 space of functions on \mathcal{X} w.r.t. measure π	A_π	Koopman operator on $L^2_\pi(\mathcal{X})$
$\text{Ker}(\cdot)$	null space of an operator	$\text{Im}(\cdot)$	range of an operator
$\text{cl}(\cdot)$	closure of a subspace	$\text{tr}(\cdot)$	trace of an operator
$\sigma_i(\cdot)$	i -th singular value of an operator	$\lambda_i(\cdot)$	i -th eigenvalue of an operator
$\llbracket \cdot \rrbracket_r$	r -truncated SVD of an operator	I	identity operator
$k(x, y)$	kernel	ϕ	canonical feature map
\mathcal{H}	reproducing kernel Hilbert space	S	canonical injection $\mathcal{H} \hookrightarrow L^2_\pi(\mathcal{X})$
$\text{HS}(\mathcal{H}, \mathcal{G})$	space of Hilbert-Schmidt operators $\mathcal{H} \rightarrow \mathcal{G}$	$\text{B}_r(\mathcal{H})$	set of rank- r Hilbert-Schmidt operators on \mathcal{H}
$\ A\ $	operator norm of an operator A	$\ A\ _{\text{HS}}$	Hilbert-Schmidt norm of operator A
Z	restriction of the Koopman operator to \mathcal{H}	g_p	conditional mean embedding
σ_j	j -th singular value of S	J	countable index set of singular values of S
ℓ_j	j -th left singular function of S	h_j	j -th right singular function of S
$\mathbb{1}$	function in $L^2_\pi(\mathcal{X})$ with the constant output 1	γ	regularization parameter
\mathcal{R}	true risk	\mathcal{E}	operator norm error
\mathcal{E}_{HS}	excess risk, i.e. HS norm error	\mathcal{R}_0	irreducible risk
\mathcal{D}_n	dataset $(x_i, y_i)_{i \in [n]}$	$\widehat{\mathcal{R}}$	empirical risk
\widehat{S}	sampling operator of the inputs	\widehat{Z}	sampling operator of the outputs
G	population Koopman estimator in $\text{HS}(\mathcal{H})$	\widehat{G}	empirical Koopman estimator in $\text{HS}(\mathcal{H})$
G_γ	population KRR estimator	\widehat{G}_γ	empirical KRR estimator
$G_{r,\gamma}^{\text{PCR}}$	population PCR estimator	$\widehat{G}_{r,\gamma}^{\text{PCR}}$	empirical PCR estimator
$G_{r,\gamma}^{\text{RRR}}$	population RRR estimator	$\widehat{G}_{r,\gamma}^{\text{RRR}}$	empirical RRR estimator
C	covariance operator	\widehat{C}	empirical covariance operator
C_γ	regularized covariance operator	\widehat{C}_γ	regularized empirical covariance operator
T	cross-covariance operator	\widehat{T}	empirical cross-covariance operator
K	input kernel matrix	L	output kernel Gramm matrix
K_γ	regularized input kernel matrix	M	cross-kernel matrix
B	operator $C_\gamma^{-1/2}T$	\widehat{B}	empirical operator $\widehat{C}_\gamma^{-1/2}\widehat{T}$
P	spectral projector	\widehat{P}	empirical spectral projector
η	metric distortion	$\widehat{\eta}$	empirical metric distortion
μ	Koopman eigenvalue	$\widehat{\lambda}$	eigenvalue of the empirical estimator
f	Koopman eigenfunction in $L^2_\pi(\mathcal{X})$	\widehat{f}	empirical eigenfunction in $L^2_\pi(\mathcal{X})$
$\widehat{\psi}$	right empirical eigenfunction	$\widehat{\xi}$	left empirical eigenfunction
$\text{gap}_J(\cdot)$	spectral gap of an operator w.r.t. indices J	cond	eigenvalue condition number
$c_{\mathcal{H}}$	boundness constant	$P_{\mathcal{H}}$	orthogonal projector in $L^2_\pi(\mathcal{X})$ onto $\text{Im}(S)$
α	regularity parameter	a	regularity constant
β	spectral decay parameter	b	spectral decay constant
τ	embedding parameter	c_τ	embedding constant

Table 1: Summary of used notations.

the Boltzmann (invariant) distribution $\pi(dx) := Z^{-1}e^{-\beta U(x)}dx$ for the process X_t (Z being just a normalizing constant). Technically, \mathcal{L} is the infinitesimal generator of the semigroup associated to the Markov process X_t on $L^2_\pi(\mathcal{X})$ and completely defines [4] the Koopman operator via the unique solution of (20). As usual, the infinitesimal generator is well defined on a dense domain of $L^2_\pi(\mathcal{X})$, that is $C^2(\mathbb{R}^d)$.

A.1 Spectral Decomposition

Recalling that for a bounded linear operator A on some Hilbert space \mathcal{H} the *resolvent set* of the operator A is defined as $\text{Res}(A) := \{\lambda \in \mathbb{C} : A - \lambda I \text{ is bijective}\}$, and its *spectrum* $\text{Sp}(A) := \mathbb{C} \setminus \{\text{Res}(A)\}$, let $\lambda \subseteq \text{Sp}(A)$ be isolated part of spectra, i.e. both λ and $\mu := \text{Sp}(A) \setminus \lambda$ are closed in $\text{Sp}(A)$. Then, the *Riesz spectral projector* $P_\lambda : \mathcal{H} \rightarrow \mathcal{H}$ is defined by

$$P_\lambda := \frac{1}{2\pi} \int_\Gamma (zI - A)^{-1} dz, \quad (21)$$

where Γ is any contour in the resolvent set $\text{Res}(A)$ with λ in its interior and separating λ from μ . Indeed, we have that $P_\lambda^2 = P_\lambda$ and $\mathcal{H} = \text{Im}(P_\lambda) \oplus \text{Ker}(P_\lambda)$ where $\text{Im}(P_\lambda)$ and $\text{Ker}(P_\lambda)$ are both invariant under A and $\text{Sp}(A|_{\text{Im}(P_\lambda)}) = \lambda$, $\text{Sp}(A|_{\text{Ker}(P_\lambda)}) = \mu$. Moreover, $P_\lambda + P_\mu = I$ and $P_\lambda P_\mu = P_\mu P_\lambda = 0$.

Finally if A is *compact* operator, then the Riesz-Schauder theorem, see e.g. [35], assures that $\text{Sp}(T)$ is a discrete set having no limit points except possibly $\lambda = 0$. Moreover, for any nonzero $\lambda \in \text{Sp}(T)$, then λ is an *eigenvalue* (i.e. it belongs to the point spectrum) of finite multiplicity, and, hence, we can deduce the spectral decomposition in the form

$$A = \sum_{\lambda \in \text{Sp}(A)} \lambda P_\lambda, \quad (22)$$

where geometric multiplicity of λ , $r_\lambda := \text{rank}(P_\lambda)$, is bounded by the algebraic multiplicity of λ . If additionally A is normal operator, i.e. $AA^* = A^*A$, then $P_\lambda = P_\lambda^*$ is orthogonal projector for each $\lambda \in \text{Sp}(A)$ and $P_\lambda = \sum_{i=1}^{r_\lambda} \psi_i \otimes \psi_i$, where ψ_i are normalized eigenfunctions of A corresponding to λ and r_λ is both algebraic and geometric multiplicity of λ .

We conclude this section with well-known perturbation bounds for eigenfunctions and spectral projectors of self-adjoint compact operators.

Proposition 2 ([17]). *Let A be compact self-adjoint operator on a separable Hilbert space \mathcal{H} . Given a pair $(\hat{\lambda}, \hat{f}) \in \mathbb{C} \times \mathcal{H}$ such that $\|\hat{f}\| = 1$, let μ be the eigenvalue of A that is closest to $\hat{\lambda}$ and let f be its normalized eigenfunction. If $\hat{g} := \min\{|\hat{\lambda} - \lambda| : \lambda \in \text{Sp}(A) \setminus \{\mu\}\} > 0$, then $\sin(\angle(\hat{f}, f)) \leq \|A\hat{f} - \hat{\lambda}\hat{f}\|/\hat{g}$.*

Proposition 3 ([48]). *Let A and \hat{A} be two compact operators on a separable Hilbert space. For nonempty index set $J \subset \mathbb{N}$ let*

$$\text{gap}_J(A) := \min\{|\lambda_i(A) - \lambda_j(A)| : i \in \mathbb{N} \setminus J, j \in J\}$$

denote the spectral gap w.r.t J and let P_J and \hat{P}_J be the corresponding spectral projectors of A and \hat{A} , respectively. If A is self-adjoint and for some $\|A - \hat{A}\| < \text{gap}_J(A)$, then

$$\|P_J - \hat{P}_J\| \leq \frac{\|A - \hat{A}\|}{\text{gap}_J(A)}.$$

A.2 Koopman Operator and Mode Decomposition

The main reason for the use of (stochastic) *Koopman operator* in dynamical systems lies in the fact that its linearity can be exploited to compute a spectral decomposition. Indeed, in many situations, and notably for compact Koopman operators, there exist scalars $\mu_i \in \mathbb{C}$, called Koopman eigenvalues, and observables $\psi_i \in L^2_\pi(\mathcal{X}) \setminus \{0\}$, called Koopman eigenfunctions, such that $A_\pi f_i = \mu_i f_i$. Then, the dynamical system can be decomposed into superposition of simpler signals that can be used in different tasks such as system identification and control, see e.g. [8]. This becomes particularly

elegant when A_π is compact, then for every observable $f \in L_\pi^2(\mathcal{X})$ there exist corresponding scalars $\gamma_i^f \in \mathbb{C}$ known as Koopman modes of f , such that

$$A_\pi^t f(x) = \mathbb{E}[f(X_t) | X_0 = x] = \sum_{j \in \mathbb{N}} \mu_j^t \gamma_j^f f_j(x), \quad x \in \mathcal{X}, t \in \mathbb{N}. \quad (23)$$

This formula is known as *Koopman Mode Decomposition* (KMD) [2, 9]. It decomposes the expected dynamics observed by f into *stationary* modes γ_j^f that are combined with *temporal changes* governed by eigenvalues μ_j and *spatial changes* governed by the eigenfunctions f_j .

KMD is closely related to general theory of spectral decomposition for bounded linear operators, and in particular the Riesz decomposition theorem. Namely, KMD of a compact self-adjoint Koopman operator can be stated as

$$A_\pi^t f(x) = \mathbb{E}[f(X_t) | X_0 = x] = \sum_{j \in \mathbb{N}} \mu_j^t \langle f_j, f \rangle f_j(x), \quad f \in L_\pi^2(\mathcal{X}), x \in \mathcal{X}, t \in \mathbb{N}. \quad (24)$$

B Kernel-Based Learning of the Koopman Operator

In many practical scenarios the transition kernel p , hence A_π , is unknown, but data from one or multiple system trajectories are available. In such situations a learning framework called Koopman operator regression was proposed in [24] to estimate Koopman operator on $L_\pi^2(\mathcal{X})$ using reproducing kernel Hilbert spaces (RKHS). More precisely, let \mathcal{H} be an RKHS with kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ [3] and let $\phi : \mathcal{X} \rightarrow \mathcal{H}$ be an associated feature map, such that $k(x, y) = \langle \phi(x), \phi(y) \rangle$ for all $x, y \in \mathcal{X}$. We assume that $k(x, x) \leq c_{\mathcal{H}} < \infty$, π -almost surely. This ensures that $\mathcal{H} \subseteq L_\pi^2(\mathcal{X})$ and the injection operator $S : \mathcal{H} \rightarrow L_\pi^2(\mathcal{X})$ given by $(Sf)(x) = f(x)$, $x \in \mathcal{X}$ is a well defined Hilbert-Schmidt operator [10, 38]. Then, the Koopman operator restricted to \mathcal{H} is given by

$$Z := A_\pi S : \mathcal{H} \rightarrow L_\pi^2(\mathcal{X}).$$

Since operator Z is, unlike A_π , Hilbert-Schmidt, one can aim to approximate Z by minimizing the following risk $\mathcal{R}(G) = \mathbb{E}_{x \sim \pi} \sum_{i \in \mathbb{N}} \mathbb{E}[(h_i(X_{t+1}) - (Gh_i)(X_t))^2 | X_t = x]$ over Hilbert-Schmidt operators $G \in \text{HS}(\mathcal{H})$, where $(h_i)_{i \in \mathbb{N}}$ is an orthonormal basis of \mathcal{H} . Moreover, there is a bias-variance decomposition of the risk $\mathcal{R}(G) = \mathcal{R}_0 + \mathcal{E}_{\text{HS}}(G)$, where

$$\mathcal{R}_0 = \|S\|_{\text{HS}}^2 - \|Z\|_{\text{HS}}^2 \geq 0 \quad \text{and} \quad \mathcal{E}_{\text{HS}}(G) = \|Z - SG\|_{\text{HS}}^2, \quad (25)$$

are the irreducible risk (i.e. the variance term in the classical bias-variance decomposition) and the excess risk, respectively. This can be equivalently expressed in the terms of embedded dynamics in RKHS as:

$$\underbrace{\mathbb{E}_{(x,y) \sim \rho} \|\phi(y) - G^* \phi(x)\|^2}_{\mathcal{R}(G)} = \underbrace{\mathbb{E}_{(x,y) \sim \rho} \|g_p(x) - \phi(y)\|^2}_{\mathcal{R}_0} + \underbrace{\mathbb{E}_{x \sim \pi} \|g_p(x) - G^* \phi(x)\|^2}_{\mathcal{E}_{\text{HS}}(G)}, \quad (26)$$

where the regression function $g_p : \mathcal{X} \rightarrow \mathcal{H}$ is defined as $g_p(x) := \mathbb{E}[\phi(X_{t+1}) | X_t = x] = \int_{\mathcal{X}} p(x, dy) \phi(y)$, $x \in \mathcal{X}$, and is known as the *conditional mean embedding* (CME) of the conditional probability p into \mathcal{H} . It was also shown that using universal kernels one can approximate the restriction of Koopman arbitrarily well, i.e. excess risk can be made arbitrarily small $\inf_{G \in \text{HS}(\mathcal{H})} \mathcal{E}_{\text{HS}}(G) = 0$.

Therefore, to develop estimators one can consider the problem of minimizing the Tikhonov regularized risk

$$\min_{G \in \text{HS}(\mathcal{H})} \mathcal{R}^\gamma(G) := \mathcal{R}(G) + \gamma \|G\|_{\text{HS}}^2, \quad (27)$$

where $\gamma > 0$. Denoting the covariance matrix as $C := S^* S = \mathbb{E}_{x \sim \pi} \phi(x) \otimes \phi(x)$ and cross-covariance matrix $T := S^* Z = \mathbb{E}_{(x,y) \sim \rho} \phi(x) \otimes \phi(y)$, where $\rho(dx, dy) := \pi(dx) p(x, dy)$ is the joint probability measure of two consecutive states of the Markov chain, and regularized covariance as $C_\gamma := C + \gamma I_{\mathcal{H}}$, one easily shows that $G_\gamma := C_\gamma^{-1} T$ is the unique solution of (27) which is known as the Kernel Ridge Regression (KRR) estimator of A_π .

To approximate the leading eigenvalues of the Koopman operator low rank estimators have been also considered. Notably, Principal Component Regression (PCR) estimator given by $\llbracket C \rrbracket_r^\dagger T$, where $\llbracket \cdot \rrbracket_r$ denotes the r -truncated SVD of the Hilbert-Schmidt operator. However, it is observed that both

KRR and PCR estimators can fail in estimating well the leading Koopman eigenvalues. To mitigate this, Reduced Rank Regression (RRR) estimator has been introduced in [24] as the optimal one that solves (27) with an additional rank constraint by minimizing over the class of rank- r HS operators $B_r(\mathcal{H}) := \{G \in \text{HS}(\mathcal{H}) \mid \text{rank}(G) \leq r\}$, where $1 \leq r < \infty$, i.e.

$$C_\gamma^{-1/2} \llbracket C_\gamma^{-1/2} T \rrbracket_r = \arg \min_{G \in B_r(\mathcal{H})} \mathcal{R}^\gamma(G). \quad (28)$$

Now, assuming that data $\mathcal{D}_n = \{(x_i, y_i)\}_{i \in [n]}$ is collected, the estimators are typically obtained via the regularized empirical risk $\widehat{\mathcal{R}}^\gamma(G) := \frac{1}{n} \sum_{i \in [n]} \|\phi(y_i) - G^* \phi(x_i)\|^2 + \gamma \|G\|_{\text{HS}}^2$ minimization (RERM). Introducing the sampling operators for data \mathcal{D}_n and RKHS \mathcal{H} by

$$\widehat{S}: \mathcal{H} \rightarrow \mathbb{R}^n \quad \text{s.t. } f \mapsto \frac{1}{\sqrt{n}} [f(x_i)]_{i \in [n]} \quad \text{and} \quad \widehat{Z}: \mathcal{H} \rightarrow \mathbb{R}^n \quad \text{s.t. } f \mapsto \frac{1}{\sqrt{n}} [f(y_i)]_{i \in [n]},$$

and their adjoints by

$$\widehat{S}^*: \mathbb{R}^n \rightarrow \mathcal{H} \quad \text{s.t. } w \mapsto \frac{1}{\sqrt{n}} \sum_{i \in [n]} w_i \phi(x_i) \quad \text{and} \quad \widehat{Z}^*: \mathbb{R}^n \rightarrow \mathcal{H} \quad \text{s.t. } w \mapsto \frac{1}{\sqrt{n}} \sum_{i \in [n]} w_i \psi(y_i),$$

we obtain $\widehat{\mathcal{R}}^\gamma(G) = \|\widehat{Z} - \widehat{S}G\|_{\text{HS}}^2 + \gamma \|G\|_{\text{HS}}^2$.

In the following we also use the empirical covariance operators defined as

$$\widehat{C} := \widehat{S}^* \widehat{S} = \frac{1}{n} \sum_{i \in [n]} \phi(x_i) \otimes \phi(x_i) \quad \text{and} \quad \widehat{D} := \widehat{Z}^* \widehat{Z} = \frac{1}{n} \sum_{i \in [n]} \phi(y_i) \otimes \phi(y_i), \quad (29)$$

empirical cross-covariance operator

$$\widehat{T} := \widehat{S}^* \widehat{Z} = \frac{1}{n} \sum_{i \in [n]} \phi(x_i) \otimes \phi(y_i), \quad (30)$$

and kernel Gramm matrices

$$K := \widehat{S} \widehat{S}^* = \frac{1}{n} [k(x_i, x_j)]_{i,j \in [n]} \in \mathbb{R}^{n \times n} \quad \text{and} \quad L := \widehat{Z} \widehat{Z}^* = \frac{1}{n} [k(y_i, y_j)]_{i,j \in [n]} \in \mathbb{R}^{n \times n}. \quad (31)$$

Additionally, we let $\widehat{C}_\gamma := \widehat{C} + \gamma I_{\mathcal{H}}$ be the regularized empirical covariance and $K_\gamma := K + \gamma I_n$ the regularized kernel Gram matrix. Then we obtain the empirical estimators of the Koopman operator on an RKHS that correspond to the population ones: empirical KRR estimator $\widehat{G}_\gamma := \widehat{C}_\gamma^{-1} \widehat{T}$, empirical PCR estimator $\llbracket \widehat{C} \rrbracket_r^\dagger \widehat{T}$, and empirical RRR estimator $\widehat{C}_\gamma^{-1/2} \llbracket \widehat{C}_\gamma^{-1/2} \widehat{T} \rrbracket_r$.

Noting that all of the empirical estimators above are of the form $\widehat{G} = \widehat{S} U_r V_r^\top \widehat{Z}$, where $U_r, V_r \in \mathbb{R}^{n \times r}$ and $r \in [n]$, see [24], we conclude this section with the result on how their spectral decompositions can be computed in an infinite dimensional RKHS.

Theorem 5 ([24]). *Let $1 \leq r \leq n$ and $\widehat{G} = \widehat{S} U_r V_r^\top \widehat{Z}$, where $U_r, V_r \in \mathbb{R}^{n \times r}$. If $V_r^\top M U_r \in \mathbb{R}^{r \times r}$, for $M = n^{-1} [k(y_i, x_j)]_{i,j \in [n]}$, is full rank and non-defective, the spectral decomposition $(\widehat{\lambda}_i, \widehat{\xi}_i, \widehat{\psi}_i)_{i \in [r]}$ of \widehat{G} can be expressed in terms of the spectral decomposition $(\widehat{\lambda}_i, \widehat{u}_i, \widehat{v}_i)_{i \in [r]}$ of $V_r^\top M U_r$ as $\widehat{\xi}_i = \widehat{\lambda}_i \widehat{Z}^* V_r \widehat{u}_i / |\widehat{\lambda}_i|$ and $\widehat{\psi}_i = \widehat{S}^* U_r \widehat{v}_i$, for all $i \in [r]$.*

C Approach

In this section, we prove key perturbation result and discuss the properties of the metric distortion. We conclude this section with the approximation bound for arbitrary estimator $G \in B_r(\mathcal{H})$ that is the basis of the statistical bounds that follow. This result is a direct consequence of [24] and Davis-Khan spectral perturbation result for compact self-adjoint operators, [17].

Theorem 1. *Let A_π be a self-adjoint compact operator and let $r \in \mathbb{N}$. Then, for every empirical estimator $\widehat{G} \in B_r(\mathcal{H})$ and every $i \in [r]$*

$$|\widehat{\lambda}_i - \mu_{j(i)}| \leq \eta(\widehat{\psi}_i) \mathcal{E}(\widehat{G}), \quad \text{and} \quad \|\widehat{f}_i - f_{j(i)}\|^2 \leq \frac{2\eta(\widehat{\psi}_i) \mathcal{E}(\widehat{G})}{[\text{gap}_{j(i)}(A_\pi) - \eta(\widehat{\psi}_i) \mathcal{E}(\widehat{G})]_+}. \quad (11)$$

Proof. We first remark that

$$\|(A_\pi - \hat{\lambda}_i I_{L_\pi^2(\mathcal{X})})^{-1}\|^{-1} \leq \|(A_\pi S - S\hat{G})\hat{\psi}_i\|/\|S\hat{\psi}_i\| \leq \mathcal{E}(\hat{G})\eta(\hat{\psi}_i).$$

Then, from the first inequality, using that A_π is normal, we obtain the first bound in (11). So, observing that for every $\mu \in \text{Sp}(A_\pi) \setminus \{\mu_i\}$,

$$|\hat{\lambda}_i - \mu| \geq |\mu_i - \mu| - |\hat{\lambda}_i - \mu_i| \geq |\mu_i - \mu| - \mathcal{E}(\hat{G})\eta(\hat{\psi}_i),$$

we conclude that

$$\min\{|\hat{\lambda}_i - \mu| \mid \mu \in \text{Sp}(A_\pi) \setminus \{\mu_i\}\} \geq \text{gap}_i(A_\pi) - \mathcal{E}(\hat{G})\eta(\hat{\psi}_i).$$

So, applying Proposition 2, we obtain

$$\begin{aligned} \sin(\angle(\hat{f}_i, f_i)) &\leq \frac{\|A_\pi \hat{f}_i - \hat{\lambda}_i \hat{f}_i\|}{[\text{gap}_i(A_\pi) - \mathcal{E}(\hat{G})\eta(\hat{\psi}_i)]_+} \leq \frac{\|(A_\pi S - S\hat{G})\hat{\psi}_i\|/\|S\hat{\psi}_i\|}{[\text{gap}_i(A_\pi) - \mathcal{E}(\hat{G})\eta(\hat{\psi}_i)]_+} \\ &\leq \frac{\mathcal{E}(\hat{G})\eta(\hat{\psi}_i)}{[\text{gap}_i(A_\pi) - \mathcal{E}(\hat{G})\eta(\hat{\psi}_i)]_+}. \end{aligned}$$

Since, clearly $\|\hat{f}_i - f_i\|^2 \leq 2(1 - \cos(\angle(\hat{f}_i, f_i))) \leq 2\sin(\angle(\hat{f}_i, f_i))$, the proof of the second bound in (11) is completed. \square

The following result shows how for the finite rank estimators one can control the metric distortion.

Proposition 1. *Let $\hat{G} \in \text{B}_r(\mathcal{H})$. For all $i \in [r]$ the metric distortion of $\hat{\psi}_i$ can be tightly bounded as*

$$1/\sqrt{\|C\|} \leq \eta(\hat{\psi}_i) \leq \min(|\hat{\lambda}_i| \text{cond}(\hat{\lambda}_i), \|\hat{G}\|) / \sigma_{\min}^+(S\hat{G}), \quad (12)$$

where $\text{cond}(\hat{\lambda}_i) := \|\hat{\xi}_i\| \|\hat{\psi}_i\| / |\langle \hat{\psi}_i, \hat{\xi}_i \rangle|$ is the condition number of $\hat{\lambda}_i$, and $\hat{\xi}_i$ is its left eigenfunction.

Proof. First, we have that $\hat{\psi}_i = \hat{\lambda}_i^{-1} \hat{G} \hat{\psi}_i = \hat{\lambda}_i^{-1} \hat{G} \hat{G}^\dagger \hat{G} \hat{\psi}_i$. But, then $g_i := \hat{G}^\dagger \hat{G} \hat{\psi}_i \in \text{Ker}(\hat{G})^\perp$ and $\hat{\lambda}_i \hat{\psi}_i = \hat{G} g_i$.

Next, recall that

$$\inf_{g \in \text{Ker}(C^{1/2} \hat{G})^\perp} \frac{\|C^{1/2} \hat{G} g\|}{\|g\|} = \sigma_{\min}^+(C^{1/2} \hat{G}) = \sigma_r(S\hat{G}).$$

So, since $\text{Im}(\hat{G}^* C^{1/2}) \subseteq \text{Im}(\hat{G}^*)$, then $g_i \in \text{Ker}(\hat{G})^\perp \subseteq \text{Ker}(C^{1/2} \hat{G})^\perp$, and we conclude

$$\sigma_r(S\hat{G})\|g_i\| \leq \|C^{1/2} \hat{G} \hat{\psi}_i\| = |\hat{\lambda}_i| \|S\hat{\psi}_i\|.$$

Therefore, since

$$\|g_i\|^2 = \langle \hat{\psi}_i, g_i \rangle = \cos(\angle(\hat{\psi}_i, \text{Ker}(\hat{G})^\perp)) \|\hat{\psi}_i\| \|g_i\| = \cos(\angle(\hat{\psi}_i, \text{Im}(\hat{G}^*))) \|\hat{\psi}_i\| \|g_i\|,$$

we have that

$$\frac{|\hat{\lambda}_i| \|S\hat{\psi}_i\|}{\sigma_r(S\hat{G})} \geq \|g_i\| = \cos(\angle(\hat{\psi}_i, \text{Im}(\hat{G}^*))) \|\hat{\psi}_i\| \geq \left(|\cos(\angle(\hat{\psi}_i, \hat{\xi}_i))| \wedge |\hat{\lambda}_i| \|\hat{G}\|^{-1} \right) \|\hat{\psi}_i\|$$

where the last inequality holds since $\hat{\xi}_i \in \text{Im}(\hat{G}^*)$ and

$$\|g_i\| = |\hat{\lambda}_i|^{-1} \|\hat{G} g_i\| = |\hat{\lambda}_i|^{-1} \|\hat{G} \hat{\psi}_i\| \leq |\hat{\lambda}_i|^{-1} \|\hat{G}\| \|\hat{\psi}_i\|.$$

We remark that this inequality becomes equality when the eigenvalue $\hat{\lambda}_i$ is simple. Finally, noticing that $|\cos(\angle(\hat{\psi}_i, \hat{\xi}_i))| = \text{cond}(\hat{\lambda}_i)$ we have

$$\eta(\hat{\psi}_i) \leq \frac{|\hat{\lambda}_i| \text{cond}(\hat{\lambda}_i) \wedge \|\hat{G}\|}{\sigma_r(S\hat{G})},$$

and application of Weyl's inequality to the denominator completes the proof. \square

Remark 3. We remark that in Example 3 of the main text, for the choice of kernel with $\Pi: i \mapsto i$, $i \in \mathbb{N}$, after some basic algebra, for the r -th eigenpair (λ_r, ψ_r) of $G_{r,\gamma}^{\text{RRR}}$ we obtain that $\eta(\psi_r) = \lambda_r \text{cond}(\psi_r) / \sigma_r(SG_{r,\gamma}^{\text{RRR}})$. This makes inequality (12) tight. Moreover, for small enough $\gamma > 0$ and $r = 1$, $|\lambda_r - \mu_r| = \eta(\psi_r)\mathcal{E}(G_{r,\gamma}^{\text{RRR}})$, where . Hence, (11) is also tight.

Next result provides the reasoning for using empirical metric distortion given by (13) in the main body.

Proposition 4. Given $r \in \mathbb{N}$, let $(\hat{\lambda}, \hat{\psi}_i)_{i=1}^r$ be nonzero eigenpairs of $\hat{G} = \hat{S}^* U_r V_r^\top \hat{Z} \in \text{B}_r(\mathcal{H})$. If $(\hat{v}_i)_{i \in [r]}$ are eigenvectors of the non-defective matrix $V_r^\top M U_r \in \mathbb{R}^{r \times r}$, for $M = n^{-1}[k(y_i, x_j)_{i,j \in [n]}]$, then for every $i \in [r]$

$$\hat{\eta}_i = \frac{\|\hat{\psi}_i\|}{\|\hat{S}\hat{\psi}_i\|} = \sqrt{\frac{\hat{v}_i^* U_r^\top K U_r \hat{v}_i}{\|K U_r \hat{v}_i\|^2}}, \quad (32)$$

and

$$\left| \hat{\eta}_i - \eta(\hat{\psi}_i) \right| \leq \left(\eta(\hat{\psi}_i) \wedge \hat{\eta}_i \right) \eta(\hat{\psi}_i) \hat{\eta}_i \|\hat{C} - C\|. \quad (33)$$

Proof. First, note that (32) follows directly from Theorem 5. Next, since for every $i \in [r]$,

$$(\hat{\eta}_i)^{-2} - (\eta(\hat{\psi}_i))^{-2} = \frac{\langle \hat{\psi}_i, (\hat{C} - C)\hat{\psi}_i \rangle}{\|\hat{\psi}_i\|^2} \leq \|\hat{C} - C\|,$$

we obtain

$$\left| \hat{\eta}_i^{-1} - (\eta(\hat{\psi}_i))^{-1} \right| \leq \frac{|\hat{\eta}_i^{-2} - (\eta(\hat{\psi}_i))^{-2}|}{(\eta(\hat{\psi}_i))^{-1} \vee \hat{\eta}_i^{-1}} \leq \left(\eta(\hat{\psi}_i) \wedge \hat{\eta}_i \right) \|\hat{C} - C\|.$$

□

Remark 4. We remark that for deterministic dynamical systems, to check if an eigenpair $(\hat{\lambda}_i, \hat{\psi}_i)$ is spurious authors in [14, 26] suggest to check if the $\hat{\psi}_i(y_i) \approx \hat{\lambda}_i \hat{\psi}_i(x_i)$ on a training set \mathcal{D}_n . Clearly, one should check the same on the validation set in order to assure that over-fitting does not occur. It is interesting to note that such strategies rely on the empirical estimate $\|(\hat{Z} - \hat{S}\hat{G})\hat{\psi}_i\|$, while our analysis aims to give high-probability finite sample guarantees via the bounds on metric distortion and operator norm error.

D Controlling the Operator Norm Error

D.1 Main Assumptions

We start by observing that $S \in \text{HS}(\mathcal{H}, L_\pi^2(\mathcal{X}))$, according to the spectral theorem for positive self-adjoint operators, has an SVD, i.e. there exists at most countable positive sequence $(\sigma_j)_{j \in J}$, where $J := \{1, 2, \dots\} \subseteq \mathbb{N}$, and ortho-normal systems $(\ell_j)_{j \in J}$ and $(h_j)_{j \in J}$ of $\text{cl}(\text{Im}(S))$ and $\text{Ker}(S)^\perp$, respectively, such that $S h_j = \sigma_j \ell_j$ and $S^* \ell_j = \sigma_j h_j$, $j \in J$.

Now, given $\alpha \geq 0$, let us define scaled injection operator $S_\alpha: \mathcal{H} \rightarrow L_\pi^2(\mathcal{X})$ as

$$S_\alpha := \sum_{j \in J} \sigma_j^\alpha \ell_j \otimes h_j. \quad (34)$$

Clearly, we have that $S = S_1$, while $\text{Im } S_0 = \text{cl}(\text{Im}(S))$. Next, we equip $\text{Im}(S_\alpha)$ with a norm $\|\cdot\|_\alpha$ to build an interpolation space.

$$[\mathcal{H}]_\alpha := \left\{ f \in \text{Im}(S_\alpha) \mid \|f\|_\alpha^2 := \sum_{j \in J} \sigma_j^{-2\alpha} \langle f, \ell_j \rangle^2 < \infty \right\}.$$

We remark that for $\alpha = 1$ the space $[\mathcal{H}]_\alpha$ is just an RKHS \mathcal{H} seen as a subspace of $L_\pi^2(\mathcal{X})$. Moreover, we have the following injections

$$[\mathcal{H}]_{\alpha_1} \hookrightarrow [\mathcal{H}]_1 \hookrightarrow [\mathcal{H}]_{\alpha_2} \hookrightarrow [\mathcal{H}]_0 = L_\pi^2(\mathcal{X}),$$

where $\alpha_1 \geq 1 \geq \alpha_2 \geq 0$.

In addition, from **(BK)** we also have that RKHS \mathcal{H} can be embedded into $L_\pi^\infty(\mathcal{X})$, i.e. for some $\tau \in (0, 1]$

$$[\mathcal{H}]_1 \hookrightarrow [\mathcal{H}]_\tau \hookrightarrow L_\pi^\infty(\mathcal{X}) \hookrightarrow L_\pi^2(\mathcal{X}),$$

Now, according to [18], if $S_{\tau,\infty}: [\mathcal{H}]_\tau \hookrightarrow L_\pi^\infty(\mathcal{X})$ denotes the injection operator, its boundedness implies the polynomial decay of the singular values of S , i.e. $\sigma_j^2(S) \lesssim j^{-1/\tau}$, $j \in J$, and the following condition is assured

(KE) *Kernel embedding property*: there exists $\tau \in [\beta, 1]$ such that

$$c_\tau := \|S_{\tau,\infty}\|^2 = \text{ess sup}_{x \sim \pi} \sum_{j \in J} \sigma_j^{2\tau} |\ell_j(x)|^2 < +\infty. \quad (35)$$

In what follows we discuss how our novel assumption **(RC)** compares to the existing ones, quantifying how miss-specified the learning problem is. We consider the following assumption made in [29] to analyze the HS error of CME

(SRC) *Source condition from [29]*: for for some $\alpha \in (0, 2]$

$$\text{Im}(Z) \subseteq \text{Im}(S_\alpha) \text{ and } G_\mathcal{H}^\alpha := S_\alpha^\dagger T \in \text{HS}(\mathcal{H}).$$

Remark 5 (Finite-dimensional RKHS). *When \mathcal{H} is finite dimensional, all spaces $[\mathcal{H}]_\alpha$ are finite dimensional. Hence, $\text{Im}(Z) \subset \text{Im}(S)$ implies also $\text{Im}(Z) \subset \text{Im}(S_\alpha)$ for every $\alpha > 0$. Moreover, we can set τ arbitrary close to zero.*

Remark 6 **(SRC)** vs. **(RC)**. *According to [47, Theorem 2.2], the condition **(RC)** is equivalent to $\text{Im}(Z) \subseteq \text{Im}(S_\alpha)$, i.e. $G_\mathcal{H}^\alpha := S_\alpha^\dagger T$ is bounded operator on \mathcal{H} and $Z = S_\alpha G_\mathcal{H}^\alpha$. Hence, if **(SRC)** holds for some $\alpha \in (0, 2]$, then **(RC)** holds, too. Indeed, we have that $T = S^* Z = S^* S_\alpha G_\mathcal{H}^\alpha$, and, thus, $TT^* \preceq \|G_\mathcal{H}^\alpha\|^2 C^{1+\alpha}$. On the other hand, for the Koopman operator $A_\pi = I_{L_\pi^2(\mathcal{X})}$, of Example 2, while $G_\mathcal{H}^1 = I_\mathcal{H}$ implies **(RC)** for at least one $\alpha \in [1, 2]$, one can show that **(SRC)** doesn't hold for any $\alpha > 0$, c.f. [29, Appendix D].*

D.2 Bounding the Bias

Recalling the decomposition

$$\underbrace{\mathcal{E}(\hat{G}) := \|Z - S\hat{G}\|}_{\text{operator norm error}} \leq \underbrace{\|Z - SG_\gamma\|}_{\text{bias due to regularization}} + \underbrace{\|S(G_\gamma - G)\|}_{\text{bias due to rank reduction}} + \underbrace{\|S(G - \hat{G})\|}_{\text{variance of the estimator}}. \quad (36)$$

we first prove the bound of the first term.

Proposition 5. *Let $G_\gamma = C_\gamma^{-1}T$ for $\gamma > 0$, and $P_\mathcal{H}: L_\pi^2(\mathcal{X}) \rightarrow L_\pi^2(\mathcal{X})$ be the orthogonal projector onto $\text{cl}(\text{Im}(S))$. If the assumptions **(BK)**, **(SD)** and **(RC)** hold, then $\|G_\gamma\| \leq ac_\mathcal{H}^{(\alpha-1)/2}$ for $\alpha \in [1, 2]$, $\|G_\gamma\| \leq a\gamma^{(\alpha-1)/2}$ for $\alpha \in (0, 1]$, and*

$$\|A_\pi S - SG_\gamma\| \leq a\gamma^{\frac{\alpha}{2}} + \|(I - P_\mathcal{H})A_\pi S\|. \quad (37)$$

Proof. Recalling that $P_\mathcal{H} := \sum_{j \in J} \ell_j \otimes \ell_j$, start by denoting the orthogonal projectors on the subspace of k leading left singular functions of S as $P_k := \sum_{j \in [k]} \ell \otimes \ell$, respectively. Next, observe that

$$\begin{aligned} Z - SG_\gamma &= (I_{L_\pi^2(\mathcal{X})} - SC_\gamma^{-1}S^*)Z = (I_{L_\pi^2(\mathcal{X})} - (SS^* + \gamma I_\mathcal{H})^{-1}SS^*)Z \\ &= \gamma(SS^* + \gamma I_\mathcal{H})^{-1}Z = \left(\sum_{j \in J} \frac{\gamma}{\sigma_j^2 + \gamma} \ell_j \otimes \ell_j \right) Z = \left(\sum_{j \in J} \frac{\gamma}{(\sigma_j^2 + \gamma)\sigma_j} \ell_j \otimes h_j \right) T. \end{aligned}$$

Therefore, for every $k \in J$

$$\|P_k(Z - SG_\gamma)\|^2 = \left\| \left(\sum_{j \in [k]} \frac{\gamma}{(\sigma_j^2 + \gamma)\sigma_j} \ell_j \otimes h_j \right) TT^* \left(\sum_{j \in [k]} \frac{\gamma}{(\sigma_j^2 + \gamma)\sigma_j} h_j \otimes \ell_j \right) \right\|,$$

which, due to **(RC)**, implies that

$$\|P_k(Z - SG_\gamma)\| \leq a \left\| \sum_{j \in [k]} \frac{\gamma \sigma_j^\alpha}{\sigma_j^2 + \gamma} \ell_j \otimes \ell_j \right\|.$$

On the other hand,

$$\sum_{j \in [k]} \frac{\gamma \sigma_j^\alpha}{\sigma_j^2 + \gamma} \ell_j \otimes \ell_j = \gamma^{\frac{\alpha}{2}} \sum_{j \in [k]} \frac{(\sigma_j^2 \gamma^{-1})^{\frac{\alpha}{2}}}{\sigma_j^2 \gamma^{-1} + 1} \ell_j \otimes \ell_j \preceq \gamma^{\frac{\alpha}{2}} \sum_{j \in [k]} \ell_j \otimes \ell_j,$$

where the inequality holds due to $x^s \leq x + 1$ for all $x \geq 0$ and $s \in [0, 1]$. Since the norm of the projector equals one, we get $\|P_k(Z - SG_\gamma)\| \leq a\gamma^{\frac{\alpha}{2}}$.

Next, observe that

$$\begin{aligned} \|(P_{\mathcal{H}} - P_k)(Z - SG_\gamma)\|^2 &= \left\| \sum_{j \in J \setminus [k]} \frac{\gamma^2}{(\sigma_j^2 + \gamma)^2} (Z^* \ell_j) \otimes (Z^* \ell_j) \right\|^2 \leq \sum_{j \in J \setminus [k]} \frac{\gamma^2}{(\sigma_j^2 + \gamma)^2} \|Z^* \ell_j\|^2 \\ &\leq \sum_{j \in J \setminus [k]} \frac{\gamma^2 \sigma_j^{2\alpha}}{(\sigma_j^2 + \gamma)^2} \leq \sum_{j \in J \setminus [k]} \sigma_j^{2\alpha} \end{aligned}$$

So, using triangular inequality, for every $k \in J$ we have

$$\|P_{\mathcal{H}}(Z - SG_\gamma)\| \leq \|P_k(Z - SG_\gamma)\| + \|(P_{\mathcal{H}} - P_k)(Z - SG_\gamma)\| \leq a\gamma^{\frac{\alpha}{2}} + \sum_{j \in J \setminus [k]} (\sigma_j^{2\beta})^{\frac{\alpha}{\beta}},$$

and, hence, letting $k \rightarrow \infty$ we obtain $\|P_{\mathcal{H}}Z - SG_\gamma\| \leq a\gamma^{\frac{\alpha}{2}}$. Hence, (37) follows from triangular inequality.

To estimate the $\|G_\gamma\|$, note that **(RC)** implies $\|G_\gamma\| \leq a \|C_\gamma^{-1} C^{\frac{1+\alpha}{2}}\|$ and considering two cases. First, if **(RC)** holds for some $\alpha \in [1, 2]$, then, clearly $\|G_\gamma\| \leq a c_{\mathcal{H}}^{(\alpha-1)/2}$. On the other hand, if $\alpha \in (0, 1]$, then

$$\frac{\sigma_j^{1+\alpha}}{\sigma_j^2 + \gamma} = \gamma^{-1} \frac{(\sigma_j^2 \gamma^{-1})^{\frac{1+\alpha}{2}}}{\sigma_j^2 \gamma^{-1} + 1} \leq \gamma^{\frac{\alpha-1}{2}},$$

and, thus, $\|G_\gamma\| \leq a\gamma^{(\alpha-1)/2}$.

□

Remark 7. Inequality (37) says that the regularization bias is comprised of a term depending on the choice of γ , and on a term depending on the “alignment” between \mathcal{H} and $\text{Im}(A_\pi S)$. The term $\|(I - P_{\mathcal{H}})A_\pi S\|$ can be set to zero by two different approaches. One is choose a kernel which in some way minimizes $\|(I - P_{\mathcal{H}})A_\pi S\|$. Another is to choose a universal kernel [38, Chapter 4], for which $\text{Im}(A_\pi S) \subseteq \text{cl}(\text{Im}(S))$. While the former approach is common to several methods using finite-dimensional kernels, see e.g. [5, 30, 43, 46], the latter is classical in kernel-based learning of the Koopman operator, [20, 24, 29].

In order to proceed with bounding the bias due to rank reduction for both considered estimators, we first provide auxiliary result.

Proposition 6. Let $B := C_\gamma^{-1/2}T$, let **(RC)** hold for some $\alpha \in (0, 2]$. Then for every $j \in J$

$$\sigma_j^2(Z) - a^2 c_{\mathcal{H}}^{\alpha/2} \gamma^{\alpha/2} \leq \sigma_j^2(B) \leq \sigma_j^2(Z). \quad (38)$$

Proof. Start by observing that

$$B^*B = Z^*SC_\gamma^{-1}S^*Z = Z^*Z - \gamma Z^*(SS^* + \gamma I_{L_\pi^2(\mathcal{X})})^{-1}Z,$$

implies that

$$Z^*Z - \sum_{j \in J} \frac{\gamma}{\sigma_j^2 + \gamma} (Z^*\ell_j) \otimes (Z^*\ell_j) = B^*B \preceq Z^*Z.$$

Next, similarly to the above, for every $k \in J$, we have

$$\begin{aligned} \left\| \sum_{j \in [k]} \frac{\gamma}{\sigma_j^2 + \gamma} (Z^*\ell_j) \otimes (Z^*\ell_j) \right\| &\leq a^2 \left\| \sum_{j \in [k]} \frac{\sigma_j^{2\alpha}}{\sigma_j^2 \gamma^{-1} + 1} \ell_j \otimes \ell_j \right\| \\ &= a^2 \left\| \sum_{j \in [k]} \frac{(\sigma_j^2 \gamma^{-1})^{\alpha/2} \sigma_j^\alpha \gamma^{\alpha/2}}{\sigma_j^2 \gamma^{-1} + 1} \ell_j \otimes \ell_j \right\| \leq a^2 \gamma^{\alpha/2} \|C\|^{\alpha/2}, \end{aligned}$$

and

$$\left\| \sum_{j \in J \setminus [k]} \frac{\gamma}{\sigma_j^2 + \gamma} (Z^*\ell_j) \otimes (Z^*\ell_j) \right\| \leq a^2 \sum_{j \in J \setminus [k]} \frac{\gamma}{\sigma_j^2 + \gamma} \sigma_j^{2\alpha} \leq a^2 \sum_{j \in J \setminus [k]} (\sigma_j^{2\beta})^{\alpha/\beta}.$$

So, as before, letting $k \rightarrow \infty$ we get the result \square

Now, the bounds for the rank reduction bias of the two estimator follow.

Proposition 7 (RRR). *Let (RC) hold for some $\alpha \in (0, 2]$. Then the bias of $G_{r,\gamma}^{\text{RRR}}$ due to rank reduction is bounded as*

$$\sigma_{r+1}(Z) - a c_{\mathcal{H}}^{\alpha/4} \gamma^{\alpha/4} - 2a \gamma^{(1 \wedge \alpha)/2} \leq \|S(G_\gamma - G_{r,\gamma}^{\text{RRR}})\| \leq \sigma_{r+1}(Z). \quad (39)$$

Proof. Observe that

$$\|S(G_\gamma - G_{r,\gamma}^{\text{RRR}})\| \leq \|C_\gamma^{1/2}(G_\gamma - G_{r,\gamma}^{\text{RRR}})\| = \|B - \llbracket B \rrbracket_r\| = \sigma_{r+1}(B) \leq \sigma_{r+1}(Z)$$

while

$$\begin{aligned} \|S(G_\gamma - G_{r,\gamma}^{\text{RRR}})\| &\geq \|C_\gamma^{1/2}(G_\gamma - G_{r,\gamma}^{\text{RRR}})\| - \gamma^{1/2} \|G_\gamma - G_{r,\gamma}^{\text{RRR}}\| \\ &\geq \sigma_{r+1}(Z) - a \|C\|^{\alpha/4} \gamma^{\alpha/4} - 2a \gamma^{(1 \wedge \alpha)/2}. \end{aligned}$$

\square

Proposition 8 (PCR). *Let (RC) hold for some $\alpha \in (0, 2]$. Then the bias of $G_{r,\gamma}^{\text{PCR}}$ due to rank reduction is bounded as*

$$\sigma_{r+1}(S) - \sqrt{\mathcal{R}_0^{r+1}} - a \gamma^{\alpha/2} \leq \|S(G_\gamma - G_{r,\gamma}^{\text{PCR}})\| \leq \sigma_{r+1}(S), \quad (40)$$

where $\mathcal{R}_0^{r+1} := \langle (SS^* - ZZ^*)\ell_{r+1}, \ell_{r+1} \rangle \geq 0$ is the irreducible risk restricted to the $(r+1)$ -st left singular function of S .

Proof. Let P_r denote the orthogonal projector onto the subspace of leading r eigenfunctions of C . Then the upper bound is easily obtained as

$$\|S(G_\gamma - G_{r,\gamma}^{\text{PCR}})\| = \|C^{1/2}(I - P_r)G_\gamma\| \leq \sigma_{r+1}(S) \|(I - P_r)G_\gamma\| \leq \sigma_{r+1}(S).$$

Next, observe that

$$\begin{aligned} \|S(G_\gamma - G_{r,\gamma}^{\text{PCR}})\|^2 &= \|Z^*S(I - P_r)CC_\gamma^{-2}S^*Z\| = \left\| \sum_{j \geq r+1} \frac{\sigma_j^4}{(\sigma_j^2 + \gamma)^2} (Z^*\ell_j) \otimes (Z^*\ell_j) \right\| \\ &\geq \left\| \frac{\sigma_{r+1}^4}{(\sigma_{r+1}^2 + \gamma)^2} (Z^*\ell_{r+1}) \otimes (Z^*\ell_{r+1}) \right\| = \frac{\sigma_{r+1}^4}{(\sigma_{r+1}^2 + \gamma)^2} \|Z^*\ell_{r+1}\|^2. \end{aligned}$$

But, since $\|Z^* \ell_{r+1}\| = \|Th_{r+1}\|/\sigma_{r+1} \leq a\|C^{(1+\alpha)/2}h_{r+1}\|/\sigma_{r+1} = a\sigma_{r+1}^\alpha$, we obtain

$$\begin{aligned} \|S(G_\gamma - G_{r,\gamma}^{\text{PCR}})\| &\geq \left(1 - \frac{\gamma}{\sigma_{r+1}^2 + \gamma}\right) \|Z^* \ell_{r+1}\| \geq \|Z^* \ell_{r+1}\| - a \frac{\gamma \sigma_{r+1}}{\sigma_{r+1}^2 + \gamma} \\ &= \|Z^* \ell_{r+1}\| - \frac{(\sigma_{r+1}^2 \gamma^{-1})^{\alpha/2}}{\sigma_{r+1}^2 \gamma^{-1} + 1} \geq \|Z^* \ell_{r+1}\| - a \gamma^{\alpha/2}. \end{aligned}$$

Finally,

$$\|Z^* \ell_{r+1}\|^2 = \langle ZZ^* \ell_{r+1}, \ell_{r+1} \rangle = \langle SS^* \ell_{r+1}, \ell_{r+1} \rangle - \mathcal{R}_0^{r+1},$$

which completes the proof \square

We observe that \mathcal{R}_0^{r+1} measures the variance of the Koopman operator A_π over the $(r+1)$ -st left singular function of S . Hence, it is immediate that $\mathcal{R}_0^{r+1} \leq \mathcal{R}_0$, and the previous result indicates that when Koopman operator has small irreducible risk the bias introduced by PCR is indeed of the order $\sigma_{r+1}(S)$. Namely, the bound in (40) is sharp provided that $\sigma_{r+1}(S) \geq (1+c)\sqrt{\mathcal{R}_0^{r+1}}$ for some absolute constant $c > 0$. This means that the learning rate of PCR can be significantly worse than that of RRR in bias dominating scenarios which can occur when the kernel is somehow "misaligned" with the Koopman operator (that is $\sigma_{r+1}(S) \gg \sigma_{r+1}(Z) \vee n^{-\frac{\alpha}{2(\alpha+\beta)}}$). We have provided one such example in Example 3.

D.3 Bounding the Variance

D.3.1 Concentration Inequalities

All the statistical bounds we present will rely on two versions of Bernstein inequality. The first one is Pinelis and Sakhanenko inequality for random variables in a separable Hilbert space, see [10, Proposition 2].

Proposition 9. *Let A_i , $i \in [n]$ be i.i.d copies of a random variable A in a separable Hilbert space with norm $\|\cdot\|$. If there exist constants $L > 0$ and $\sigma > 0$ such that for every $m \geq 2$ $\mathbb{E}\|A\|^m \leq \frac{1}{2}m!L^{m-2}\sigma^2$, then with probability at least $1 - \delta$*

$$\left\| \frac{1}{n} \sum_{i \in [n]} A_i - \mathbb{E}A \right\| \leq \frac{4\sqrt{2}}{\sqrt{n}} \log \frac{2}{\delta} \sqrt{\sigma^2 + \frac{L^2}{n}} \quad (41)$$

On the other hand, we recall that in [33], a dimension-free version of the non-commutative Bernstein inequality for finite-dimensional symmetric matrices is proposed (see also Theorem 7.3.1 in [39] for an easier to read and slightly improved version) as well as an extension to self-adjoint Hilbert-Schmidt operators on a separable Hilbert spaces.

Proposition 10. *Let A_i , $i \in [n]$ be i.i.d copies of a Hilbert-Schmidt operator A on the separable Hilbert space. Let $\|A\| \leq c$ almost surely, $\mathbb{E}A = 0$ and let $\mathbb{E}[A^2] \preceq V$ for some trace class operator V . Then with probability at least $1 - \delta$*

$$\left\| \frac{1}{n} \sum_{i \in [n]} A_i \right\| \leq \frac{2c}{3n} \mathcal{L}_A(\delta) + \sqrt{\frac{2\|V\|}{n} \mathcal{L}_A(\delta)}, \quad (42)$$

where

$$\mathcal{L}_A(\delta) := \log \frac{4}{\delta} + \log \frac{\text{tr}(V)}{\|V\|}$$

We use the same strategy in combination with a standard dilation method to extend a deviation inequality on rectangular matrices of [39] (Corollary 7.3.2) to Hilbert-Schmidt operators on a separable Hilbert space.

Proposition 11. Let $A_i, i \in [n]$ be i.i.d copies of a Hilbert-Schmidt operator A on the separable Hilbert space. Let $\|A\| \leq c$ almost surely and let $\mathbb{E}[AA^*] \preceq V$ and $\mathbb{E}[A^*A] \preceq V'$ for some trace class operators V and V' . Then with probability at least $1 - \delta$

$$\left\| \frac{1}{n} \sum_{i \in [n]} A_i - \mathbb{E}A \right\| \leq \frac{4c}{3n} \mathcal{L}_A(\delta) + \sqrt{\frac{2 \max\{\|V\|, \|V'\|\}}{n}} \mathcal{L}_A(\delta), \quad (43)$$

where

$$\mathcal{L}_A(\delta) := \log \frac{4}{\delta} + \log \frac{\text{tr}(V + V')}{\max\{\|V\|, \|V'\|\}}$$

Proof. Let

$$B_i = \begin{bmatrix} 0 & A_i \\ A_i^* & 0 \end{bmatrix} \text{ and } B = \begin{bmatrix} 0 & A \\ A^* & 0 \end{bmatrix},$$

then $\|B\| = \|A\|$, $\frac{1}{n} \sum_{i \in [n]} B_i - \mathbb{E}B = \frac{1}{n} \sum_{i \in [n]} A_i - \mathbb{E}A$ and

$$\mathbb{E}(B - \mathbb{E}B)^2 \preceq \mathbb{E}B^2 = \begin{bmatrix} \mathbb{E}[AA^*] & 0 \\ 0 & \mathbb{E}[A^*A] \end{bmatrix} \preceq \begin{bmatrix} V & 0 \\ 0 & V' \end{bmatrix} =: V''.$$

Moreover, $\|B - \mathbb{E}B\| \leq \|B\| + \sqrt{\mathbb{E}\|B\|^2} \leq 2c$. Applying Proposition 10 we complete the proof. \square

Proposition 12. Given $\delta > 0$, with probability in the i.i.d. draw of $(x_i, y_i)_{i=1}^n$ from ρ , it holds that

$$\mathbb{P}\{\|\hat{T} - T\| \leq \varepsilon_n(\delta)\} \wedge \mathbb{P}\{\|\hat{C} - C\| \leq \varepsilon_n(\delta)\} \geq 1 - \delta,$$

where

$$\varepsilon_n(\delta) := \frac{4c_{\mathcal{H}}}{3n} \mathcal{L}(\delta) + \sqrt{\frac{2\|C\|}{n}} \mathcal{L}(\delta) \quad \text{and} \quad \mathcal{L}(\delta) := \log \frac{4 \text{tr}(C)}{\delta \|C\|}. \quad (44)$$

Proof. Proof follows directly from Proposition 11 applied to operators $\phi(x_i) \otimes \phi(x_i)$ and $\phi(x_i) \otimes \phi(y_i)$, respectively using the fact that $C = \mathbb{E} \phi(x_i) \otimes \phi(x_i) = \mathbb{E} \phi(y_i) \otimes \phi(y_i)$ and $T = \mathbb{E} \phi(x_i) \otimes \phi(y_i)$. \square

Proposition 13. Let (KE) hold for $\tau \in [\beta, 1]$. Given $\delta > 0$, with probability in the i.i.d. draw of $(x_i, y_i)_{i=1}^n$ from ρ , it holds that

$$\mathbb{P}\left\{\|C_\gamma^{-1/2}(\hat{C} - C)C_\gamma^{-1/2}\| \leq \varepsilon_n^1(\gamma, \delta)\right\} \wedge \mathbb{P}\left\{\|C_\gamma^{-1/2}(\hat{T} - T)C_\gamma^{-1/2}\| \leq \varepsilon_n^1(\gamma, \delta)\right\} \geq 1 - \delta, \quad (45)$$

where

$$\varepsilon_n^1(\gamma, \delta) := \frac{4c_\tau}{3n\gamma^\tau} \mathcal{L}^1(\gamma, \delta) + \sqrt{\frac{2c_\tau}{n\gamma^\tau}} \mathcal{L}^1(\gamma, \delta), \quad (46)$$

where

$$\mathcal{L}^1(\gamma, \delta) := \ln \frac{4}{\delta} + \ln \frac{\text{tr}(C_\gamma^{-1}C)}{\|C_\gamma^{-1}C\|},$$

Moreover,

$$\mathbb{P}\left\{\|C_\gamma^{1/2}\hat{C}_\gamma^{-1}C_\gamma^{1/2}\| \leq \frac{1}{1 - \varepsilon_n^1(\gamma, \delta)}\right\} \geq 1 - \delta. \quad (47)$$

Proof. The idea is to apply Proposition 11 for operator $A = \xi(x) \otimes \xi(x)$, where $\xi(x) := C_\gamma^{-1/2} \phi(x)$. To that end, observe that for every $\tau > 0$ we have that

$$\begin{aligned} \|\xi(x)\|^2 &= \sum_{j \in J} \langle C_\gamma^{-1/2} \phi(x), h_j \rangle^2 = \sum_{j \in J} \frac{1}{\sigma_j^2 + \gamma} \langle \phi(x), h_j \rangle^2 = \sum_{j \in J} \frac{\sigma_j^{2(1-\tau)}}{\sigma_j^2 + \gamma} \frac{\langle \phi(x), h_j \rangle^2}{\sigma_j^2} \sigma_j^{2\tau} \\ &= \gamma^{-\tau} \sum_{j \in J} \frac{(\sigma_j^2 \gamma^{-1})^{1-\tau}}{\sigma_j^2 \gamma^{-1} + 1} \frac{|h_j(x)|^2}{\sigma_j^2} \sigma_j^{2\tau} \leq \gamma^{-\tau} \sum_{j \in J} \frac{|(Sh_j)(x)|^2}{\sigma_j^2} \sigma_j^{2\tau} = \gamma^{-\tau} \sum_{j \in J} |\ell_j(x)|^2 \sigma_j^{2\tau} \end{aligned}$$

So, due to (35), we obtain $\|A\| \leq \|\xi\|_\infty^2 \leq \gamma^{-\tau} c_\tau$. On the other hand, since

$$\mathbb{E}_{x \sim \pi} (\xi(x) \otimes \xi(x))^2 \preceq \|\xi\|_\infty^2 \mathbb{E}_{x \sim \pi} \xi(x) \otimes \xi(x) = \|\xi\|_\infty^2 C_\gamma^{-1/2} C C_\gamma^{-1/2}.$$

Next, we observe that

$$\|I_{\mathcal{H}} - C_\gamma^{-1/2} \hat{C}_\gamma C_\gamma^{-1/2}\| = \|C_\gamma^{-1/2} (C - \hat{C}) C_\gamma^{-1/2}\| \leq \varepsilon_n^1(\gamma, \delta).$$

Thus, for $\varepsilon_n^1(\gamma, \delta)$ smaller than one, it follows that

$$\|C_\gamma^{1/2} \hat{C}_\gamma^{-1} C_\gamma^{1/2}\| = \|(C_\gamma^{-1/2} \hat{C}_\gamma C_\gamma^{-1/2})^{-1}\| \leq \frac{1}{1 - \|I_{\mathcal{H}} - C_\gamma^{-1/2} \hat{C}_\gamma C_\gamma^{-1/2}\|},$$

and the proof is completed. \square

Proposition 14. *Let (KE) hold for $\tau \in [\beta, 1]$. Given $\delta > 0$, with probability in the i.i.d. draw of $(x_i, y_i)_{i=1}^n$ from ρ , it holds*

$$\mathbb{P} \left\{ \|C_\gamma^{-1/2} (\hat{C} - C)\|_{\text{HS}} \leq \varepsilon_n^2(\gamma, \delta) \right\} \wedge \mathbb{P} \left\{ \|C_\gamma^{-1/2} (\hat{T} - T)\|_{\text{HS}} \leq \varepsilon_n^2(\gamma, \delta) \right\} \geq 1 - \delta,$$

where

$$\varepsilon_n^2(\gamma, \delta) := 4 \sqrt{2 c_{\mathcal{H}}} \ln \frac{2}{\delta} \sqrt{\frac{\text{tr}(C_\gamma^{-1} C)}{n} + \frac{c_\tau}{n^2 \gamma^\tau}}. \quad (48)$$

Proof. First, recall that $\text{HS}(\mathcal{H})$ equipped with $\|\cdot\|_{\text{HS}}$ is separable Hilbert space. Hence, we will apply Proposition 9 for $A = \xi(x) \otimes \phi(y)$, where $\xi(x) := C_\gamma^{-1/2} \phi(x)$. To that end, observe that

$$\mathbb{E}[\|A\|_{\text{HS}}^m] = \mathbb{E}[\|\xi(x)\|^m \|\phi(y)\|^m] \leq \|\xi\|_\infty^{m-2} \|\phi\|_\infty^m \mathbb{E}[\|\xi(x)\|^2] \quad (49)$$

$$= \|\xi\|_\infty^{m-2} \|\phi\|_\infty^m \text{tr}(C_\gamma^{-1} C) \leq \frac{1}{2} m! \left(\gamma^{-\tau/2} \sqrt{c_\tau c_{\mathcal{H}}} \right)^{m-2} \left(\sqrt{c_{\mathcal{H}}} \text{tr}(C_\gamma^{-1} C) \right)^2. \quad (50)$$

\square

Proposition 15. *Let (KE) hold for $\tau \in [\beta, 1]$. Given $\delta > 0$, with probability in the i.i.d. draw of $(x_i, y_i)_{i=1}^n$ from ρ , it holds*

$$\mathbb{P} \left\{ \|C_\gamma^{-1} (\hat{C} - C)\|_{\text{HS}} \leq \varepsilon_n^3(\gamma, \delta) \right\} \wedge \mathbb{P} \left\{ \|C_\gamma^{-1} (\hat{T} - T)\|_{\text{HS}} \leq \varepsilon_n^3(\gamma, \delta) \right\} \geq 1 - \delta,$$

where

$$\varepsilon_n^3(\gamma, \delta) := 4 \sqrt{2 c_{\mathcal{H}}} \ln \frac{2}{\delta} \sqrt{\frac{\text{tr}(C_\gamma^{-1} C)}{n} + \frac{c_\tau}{n^2 \gamma^{\tau+1}}}. \quad (51)$$

Proof. Similar to the above, we apply Proposition 9 for $A = \xi(x) \otimes \phi(y)$, where $\xi(x) := C_\gamma^{-1} \phi(x)$. Hence,

$$\begin{aligned} \|\xi(x)\|^2 &= \sum_{j \in J} \langle C_\gamma^{-1} \phi(x), h_j \rangle^2 = \sum_{j \in J} \frac{1}{(\sigma_j^2 + \gamma)^2} \langle \phi(x), h_j \rangle^2 = \sum_{j \in J} \left(\frac{\sigma_j^{(1-\tau)}}{\sigma_j^2 + \gamma} \right)^2 \frac{\langle \phi(x), h_j \rangle^2}{\sigma_j^2} \sigma_j^{2\tau} \\ &= \gamma^{-(1+\tau)} \sum_{j \in J} \left(\frac{(\sigma_j^2 \gamma^{-1})^{(1-\tau)/2}}{\sigma_j^2 \gamma^{-1} + 1} \right)^2 |\ell_j(x)|^2 \sigma_j^{2\tau} \leq \gamma^{-(1+\tau)} c_\tau, \end{aligned}$$

completes the proof. \square

Next, we develop concentration bounds of some key quantities used to build RRR and PCR empirical estimators.

D.3.2 Variance and Norm of KRR Estimator

Proposition 16. Let (RC), (SD) and (KE) hold for some $\alpha \in [1, 2]$, $\beta \in (0, 1]$ and $\tau \in [\beta, 1]$. Given $\delta > 0$ if $\varepsilon_n^1(\gamma, \delta) < 1$, then with probability at least $1 - \delta$ in the i.i.d. draw of $(x_i, y_i)_{i=1}^n$ from ρ

$$\mathbb{P} \left\{ \|C_\gamma^{1/2}(\hat{G}_\gamma - G_\gamma)\| \leq \frac{(1 + a c_{\mathcal{H}}^{(\alpha-1)/2}) \varepsilon_n^2(\gamma, \delta/3)}{1 - \varepsilon_n^1(\gamma, \delta/3)} \right\} \geq 1 - \delta,$$

and

$$\mathbb{P} \left\{ \|\hat{G}_\gamma\| \leq 1 + \frac{2 \varepsilon_n^3(\gamma, \delta/3)}{1 - \varepsilon_n^3(\gamma, \delta/3)} \right\} \geq 1 - \delta.$$

Proof. Note that $C_\gamma^{1/2}(\hat{G}_\gamma - G_\gamma) = C_\gamma^{1/2}(\hat{C}_\gamma^{-1}\hat{T} - C_\gamma^{-1}T)$, and, hence,

$$\begin{aligned} C_\gamma^{1/2}(\hat{G}_\gamma - G_\gamma) &= C_\gamma^{1/2}\hat{C}_\gamma^{-1}(\hat{T} - \hat{C}_\gamma C_\gamma^{-1}T \pm T) \\ &= C_\gamma^{1/2}\hat{C}_\gamma^{-1}C_\gamma^{1/2} \left(C_\gamma^{-1/2}(\hat{T} - T) - C_\gamma^{-1/2}(\hat{C} - C)C_\gamma^{-1}T \right). \end{aligned} \quad (52)$$

Thus, taking the norm and using $\|C_\gamma^{-1}T\| \leq a \sigma_1^{\alpha-1}(S)$ with the Propositions 14 and 13 we prove the first bound. For the second one, we use

$$\hat{G}_\gamma - G_\gamma = C_\gamma^{-1}(\hat{T} - T) - C_\gamma^{-1}(\hat{C} - C)\hat{G}_\gamma$$

with Proposition 15 to obtain

$$\|\hat{G}_\gamma\| - 1 \leq \|\hat{G}_\gamma\| - \|G_\gamma\| \leq \|\hat{G}_\gamma - G_\gamma\| \leq \varepsilon_n^3(\gamma, \delta/2)(1 + \|\hat{G}_\gamma\|),$$

which completes the proof. \square

D.3.3 Variance of Singular Values

Proposition 17. Let (RC), (SD) and (KE) hold for some $\alpha \in [1, 2]$, $\beta \in (0, 1]$ and $\tau \in [\beta, 1]$. Let $B := C_\gamma^{-1/2}T$ and $\hat{B} := \hat{C}_\gamma^{-1/2}\hat{T}$. Given $\delta > 0$ if $\varepsilon_n^1(\gamma, \delta/5) < 1$, then with probability at least $1 - \delta$ in the i.i.d. draw of $(x_i, y_i)_{i=1}^n$ from ρ

$$\|\hat{B}^* \hat{B} - B^* B\| \leq (c^2 - 1) \varepsilon_n(\delta/5) + c^2 \frac{(\varepsilon_n^2(\gamma, \delta/5))^2}{1 - \varepsilon_n^1(\gamma, \delta/3)}, \quad (53)$$

where $c := 1 + a c_{\mathcal{H}}^{(\alpha-1)/2}$. Consequently, for every $i \in [n]$

$$|\sigma_i^2(\hat{B}) - \sigma_i^2(B)| \leq (c^2 - 1) \varepsilon_n(\delta/5) + c^2 \frac{(\varepsilon_n^2(\gamma, \delta/5))^2}{1 - \varepsilon_n^1(\gamma, \delta/5)}. \quad (54)$$

Proof. We start from the Weyl's inequalities for the square of singular values

$$|\sigma_i^2(\hat{B}) - \sigma_i^2(B)| \leq \|\hat{B}^* \hat{B} - B^* B\|, \quad i \in [n].$$

But, since,

$$\hat{B}^* \hat{B} - B^* B = \hat{T}^* \hat{C}_\gamma^{-1} \hat{T} - T^* C_\gamma^{-1} T = (\hat{T} - T)^* \hat{C}_\gamma^{-1} \hat{T} + T^* C_\gamma^{-1} (\hat{T} - T) + T^* (\hat{C}_\gamma^{-1} - C_\gamma^{-1}) \hat{T}$$

denoting $M = C_\gamma^{-1/2}(\hat{T} - T)$, $N = C_\gamma^{-1/2}(\hat{C} - C)$ and $R := C_\gamma^{1/2}(\hat{G}_\gamma - G_\gamma)$, we have

$$\begin{aligned} \hat{B}^* \hat{B} - B^* B &= B^* M + M^* C_\gamma^{1/2} \hat{G}_\gamma - B^* N \hat{G}_\gamma = B^* M + (M^* C_\gamma^{1/2} - B^* N)(\hat{G}_\gamma \pm G_\gamma) \\ &= B^* M + M^* B - B^* N G_\gamma + (M^* - B^* N C_\gamma^{-1/2}) R \\ &= (G_\gamma)^* (\hat{T} - T) + (\hat{T} - T) G_\gamma - (G_\gamma)^* (\hat{C} - C) G_\gamma + (M^* + (G_\gamma)^* N^*) R. \end{aligned}$$

Therefore, recalling that, due to (52), $R = C_\gamma^{1/2} \hat{C}_\gamma^{-1} C_\gamma^{1/2} (M - N G_\gamma)$, we conclude

$$\begin{aligned} \hat{B}^* \hat{B} - B^* B &= (G_\gamma)^* (\hat{T} - T) + (\hat{T} - T) G_\gamma - (G_\gamma)^* (\hat{C} - C) G_\gamma \\ &\quad + (M - N G_\gamma)^* C_\gamma^{1/2} \hat{C}_\gamma^{-1} C_\gamma^{1/2} (M - N G_\gamma). \end{aligned} \quad (55)$$

Now, applying Propositions 12, 13 and 14 we obtain (53), and, therefore, (54) follows. \square

Remark that to bound singular values we can rely on the fact

$$|\sigma_i(\hat{B}) - \sigma_i(B)| = \frac{|\sigma_i^2(\hat{B}) - \sigma_i^2(B)|}{\sigma_i(\hat{B}) + \sigma_i(B)} \leq \frac{|\sigma_i^2(\hat{B}) - \sigma_i^2(B)|}{\sigma_i(\hat{B}) \vee \sigma_i(B)}.$$

D.3.4 Variance of RRR Estimator

Recalling the notation $B := C_\gamma^{-1/2}T$ and $\hat{B} := \hat{C}_\gamma^{-1/2}\hat{T}$, let denote P_r and \hat{P}_r denote the orthogonal projector onto the subspace of leading r right singular vectors of B and \hat{B} , respectively. Then we have $\llbracket B \rrbracket_r = B P_r$ and $\llbracket \hat{B} \rrbracket_r = \hat{B} \hat{P}_r$, and, hence $G_{r,\gamma}^{\text{RRR}} = G_\gamma P_r$ and $\hat{G}_{r,\gamma}^{\text{RRR}} = \hat{G}_\gamma \hat{P}_r$.

Proposition 18. *Let (RC), (SD) and (KE) hold for some $\alpha \in [1, 2]$, $\beta \in (0, 1]$ and $\tau \in [\beta, 1]$. Given $\delta > 0$ and $\gamma > 0$, if $\varepsilon_n^1(\gamma, \delta) < 1$, then with probability at least $1 - \delta$ in the i.i.d. draw of $(x_i, y_i)_{i=1}^n$ from ρ ,*

$$\|S(G_{r,\gamma}^{\text{RRR}} - \hat{G}_{r,\gamma}^{\text{RRR}})\| \leq \frac{c \varepsilon_n^2(\gamma, \delta/5)}{1 - \varepsilon_n^1(\gamma, \delta/5)} + \frac{\sigma_1(B)}{\sigma_r^2(B) - \sigma_{r+1}^2(B)} \frac{(c^2 - 1) \varepsilon_n(\delta/5) + c^2 (\varepsilon_n^2(\gamma, \delta/5))^2}{(1 - \varepsilon_n^1(\gamma, \delta/5))^2}, \quad (56)$$

where $c := 1 + a c_{\mathcal{H}}^{(\alpha-1)/2}$

Proof. Start by observing that $\|S(G_{r,\gamma}^{\text{RRR}} - \hat{G}_{r,\gamma}^{\text{RRR}})\| \leq \|C_\gamma^{1/2}(G_{r,\gamma}^{\text{RRR}} - \hat{G}_{r,\gamma}^{\text{RRR}})\|$ and

$$\begin{aligned} C_\gamma^{1/2}(G_{r,\gamma}^{\text{RRR}} - \hat{G}_{r,\gamma}^{\text{RRR}}) &= (C_\gamma^{1/2} \hat{C}_\gamma^{-1} C_\gamma^{1/2}) \cdot \\ &\quad \left(C_\gamma^{-1/2}(\hat{C} - C) G_{r,\gamma}^{\text{RRR}} + C_\gamma^{-1/2}(\hat{T} - T) \hat{P}_r + B(\hat{P}_r - P_r) \right). \end{aligned} \quad (57)$$

Using that the norm of orthogonal projector \hat{P} is bounded by one and that $\|G_{r,\gamma}^{\text{RRR}}\| \leq \|G_\gamma\|$, applying Propositions 13 and 14 together with Propositions 3 and 17 completes the proof. \square

D.3.5 Variance of PCR

Recall that the PCR population estimator is given by $G_{r,\gamma}^{\text{PCR}} = \llbracket C_\gamma^{-1} \rrbracket_r T$ while the empirical PCR estimator is $\hat{G}_{r,\gamma}^{\text{PCR}} = \llbracket \hat{C}_\gamma^{-1} \rrbracket_r \hat{T}$. So, in this case by P_r and \hat{P}_r we denote the orthogonal projectors onto the subspace of leading r right singular vectors of C and \hat{C} , respectively. Then we have $G_{r,\gamma}^{\text{PCR}} = \llbracket C_\gamma^{-1} \rrbracket_r T = C_\gamma^{-1} P_r T = P_r C_\gamma^{-1} T$ and $\hat{G}_{r,\gamma}^{\text{PCR}} = \llbracket \hat{C}_\gamma^{-1} \rrbracket_r \hat{T} = \hat{C}_\gamma^{-1} \hat{P}_r \hat{T} = \hat{P}_r \hat{C}_\gamma^{-1} \hat{T}$.

Proposition 19. *Let (RC), (SD) and (KE) hold for some $\alpha \in [1, 2]$, $\beta \in (0, 1]$ and $\tau \in [\beta, 1]$. Given $\delta > 0$ and $\gamma > 0$, if $\varepsilon_n^1(\gamma, \delta) < 1$, then for $G = \llbracket C_\gamma^{-1} \rrbracket_r T$ and $\hat{G} = \llbracket \hat{C}_\gamma^{-1} \rrbracket_r \hat{T}$ with probability at least $1 - \delta$ in the i.i.d. draw of $(x_i, y_i)_{i=1}^n$ from ρ ,*

$$\|S(G_{r,\gamma}^{\text{PCR}} - \hat{G}_{r,\gamma}^{\text{PCR}})\| \leq \frac{c \varepsilon_n^2(\gamma, \delta/4)}{1 - \varepsilon_n^1(\gamma, \delta/4)} \sqrt{\frac{1 + \varepsilon_n^1(\gamma, \delta/4)}{1 - \varepsilon_n^1(\gamma, \delta/4)}} + \frac{\sigma_1(S)}{\sigma_r(S) - \sigma_{r+1}(S)} \varepsilon_n(\delta/4), \quad (58)$$

where $c := 1 + a c_{\mathcal{H}}^{(\alpha-1)/2}$.

Proof. Start by observing that $\|S(G_{r,\gamma}^{\text{PCR}} - \hat{G}_{r,\gamma}^{\text{PCR}})\| = \|C^{1/2}(G_{r,\gamma}^{\text{PCR}} - \hat{G}_{r,\gamma}^{\text{PCR}})\|$ and

$$C^{1/2}(G_{r,\gamma}^{\text{PCR}} - \hat{G}_{r,\gamma}^{\text{PCR}}) = C^{1/2}(P_r - \hat{P}_r)G_\gamma + C^{1/2}\hat{P}_r(G_\gamma - \hat{G}_\gamma).$$

Therefore,

$$\begin{aligned} \|S(G_{r,\gamma}^{\text{PCR}} - \hat{G}_{r,\gamma}^{\text{PCR}})\| &\leq \|C^{1/2}(P_r - \hat{P}_r)G_\gamma\| + \|C_\gamma^{1/2}\hat{P}_r C_\gamma^{-1/2} C_\gamma^{1/2}(G_\gamma - \hat{G}_\gamma)\| \\ &\leq \sigma_1(S) \|P_r - \hat{P}_r\| + \|C_\gamma^{1/2}\hat{P}_r C_\gamma^{-1/2}\| \|C_\gamma^{1/2}(G_\gamma - \hat{G}_\gamma)\| \\ &\leq \sigma_1(S) \|P_r - \hat{P}_r\| + \|C_\gamma^{1/2}\hat{C}_\gamma^{-1/2}\hat{P}_r \hat{C}_\gamma^{1/2} C_\gamma^{-1/2}\| \|C_\gamma^{1/2}(G_\gamma - \hat{G}_\gamma)\|, \end{aligned}$$

Thus, using

$$\|S(G_{r,\gamma}^{\text{PCR}} - \hat{G}_{r,\gamma}^{\text{PCR}})\| \leq \sqrt{c_{\mathcal{H}}} \|P_r - \hat{P}_r\| + \|C_\gamma^{1/2}\hat{C}_\gamma^{-1/2}\| \|\hat{C}_\gamma^{1/2} C_\gamma^{-1/2}\| \|C_\gamma^{1/2}(G_\gamma - \hat{G}_\gamma)\|, \quad (59)$$

and applying Propositions 13 and 16 together with Propositions 3 and 12 we complete the proof. \square

D.4 Operator Norm Error Bounds

Summarising previous sections, in order to prove Theorem 2, we just need to analyse the bounds ε_n^1 , ε_n^2 and ε_n^3 . To that end, we use the following result, see Lemma 11 in [18].

Lemma 1. *Let (SD) hold for some $\beta \in (0, 1]$. Then, if $\beta < 1$, for all $\gamma > 0$*

$$\text{tr}(C_\gamma^{-1}C) \leq \begin{cases} \frac{b^\beta}{1-\beta} \gamma^{-\beta} & , \beta < 1, \\ c_{\mathcal{H}} \gamma^{-1} & , \beta = 1. \end{cases}$$

Now, since $\alpha \geq 1$ and $\beta \leq \tau$, we clearly obtain that for large enough n one has

$$\varepsilon_n^1(\gamma, \delta) \lesssim \frac{n^{-1/2}}{\gamma^{\tau/2}} \ln \delta^{-1}, \quad (60)$$

as well as

$$\varepsilon_n^2(\gamma, \delta) \lesssim \left(\frac{n^{-1/2}}{\gamma^{\beta/2}} \vee \frac{n^{-1}}{\gamma^{\tau/2}} \right) \ln \delta^{-1} \quad \text{and} \quad \varepsilon_n^3(\gamma, \delta) \lesssim \left(\frac{n^{-1/2}}{\gamma^{(\beta+1)/2}} \vee \frac{n^{-1}}{\gamma^{(\tau+1)/2}} \right) \ln \delta^{-1}. \quad (61)$$

Therefore, as a consequence we have two following results on the estimation of singular values of Z , and on the operator norm error.

Proposition 20. *Let (RC) and (SD) hold for some $\alpha \in [1, 2]$ and $\beta \in (0, 1]$. Then, there exists a constant $c > 0$ such that for every given $\delta \in (0, 1)$, large enough $n > r$ and small enough $\gamma > 0$ with probability at least $1 - \delta$ in the i.i.d. draw of $(x_i, y_i)_{i=1}^n$ from ρ*

$$|\sigma_j(\widehat{B}) - \sigma_j(A_\pi S)| \lesssim n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}. \quad (62)$$

Proof. The proof is direct consequence of Propositions 6 and 17 using (60)-(61). \square

Theorem 2. *Assume the operator A_π satisfies $\sigma_r(A_\pi S) > \sigma_{r+1}(A_\pi S) \geq 0$ for some $r \in \mathbb{N}$. Let (SD) and (RC) hold for some $\beta \in (0, 1]$ and $\alpha \in [1, 2]$, respectively, and let $\text{cl}(\text{Im}(S)) = L_\pi^2(\mathcal{X})$. Let*

$$\gamma \asymp n^{-\frac{1}{\alpha+\beta}} \quad \text{and} \quad \varepsilon_n^* := n^{-\frac{\alpha}{2(\alpha+\beta)}}. \quad (15)$$

Let $\delta \in (0, 1)$. Then, there exists a constant $c > 0$, depending only on \mathcal{H} , such that for large enough $n \geq r$, with probability at least $1 - \delta$ in the i.i.d. draw of \mathcal{D}_n from ρ

$$\mathcal{E}(\widehat{G}) \leq \begin{cases} \sigma_{r+1}(A_\pi S) + c \varepsilon_n^* \ln \delta^{-1} & \text{if } \widehat{G} = \widehat{G}_{r,\gamma}^{\text{RRR}}, \\ \sigma_{r+1}(S) + c \varepsilon_n^* \ln \delta^{-1} & \text{if } \widehat{G} = \widehat{G}_{r,\gamma}^{\text{PCR}} \text{ and } \sigma_r(S) > \sigma_{r+1}(S). \end{cases} \quad (16a)$$

$$(16b)$$

Proof. Since $\sigma_r(Z) > \sigma_{r+1}(Z)$ implies due to Proposition 6 that for small enough $\gamma > 0$ we have that $\sigma_r(B) > \sigma_{r+1}(B)$, Propositions 5, 7 and 18, ensure that for large enough n (16a) holds.

Similarly, Propositions 5, 8 and 19, assure that for large enough n (16b) holds if $\sigma_r(S) > \sigma_{r+1}(S)$. \square

Setting $\gamma_n = n^{-\frac{1}{\alpha+\beta}}$ we have that

$$\gamma_n^{\alpha/2} = \frac{n^{-1/2}}{\gamma_n^{\beta/2}} = n^{-\frac{\alpha}{2(\alpha+\beta)}}, \quad \frac{n^{-1/2}}{\gamma_n^{\tau/2}} = n^{-\frac{\alpha+\beta-\tau}{2(\alpha+\beta)}} \quad \text{and} \quad \frac{n^{-1/2}}{\gamma_n^{(1+\beta)/2}} = n^{-\frac{\alpha-1}{2(\alpha+\beta)}},$$

and, hence, recalling that $\alpha \geq 1$,

$$\lim_{n \rightarrow \infty} \varepsilon_n^2(\gamma_n, \delta/5) = \lim_{n \rightarrow \infty} \varepsilon_n^1(\gamma_n, \delta/5) = \lim_{n \rightarrow \infty} \varepsilon_n(\delta/5) = 0,$$

while $\varepsilon_n^3(\gamma_n, \delta/5)$ is, in the worst case, bounded.

Therefore, for this choice of regularization parameter Theorem 2 assures that

$$\mathcal{E}(\widehat{G}_{r,\gamma}^{\text{RRR}}) - \sigma_{r+1}(B) \leq \mathcal{E}(\widehat{G}_{r,\gamma}^{\text{RRR}}) - \sigma_{r+1}(Z) \lesssim n^{-\frac{\alpha}{2(\alpha+\beta)}}. \quad (63)$$

Using the same arguments as the above, the following lower bounds follow.

Theorem 6. *Under the assumptions of Theorem 2, recalling (15), there exists a constant $c > 0$ such that for large enough $n \geq r$ and every $i \in [r]$ with probability at least $1 - \delta$ in the i.i.d. draw of \mathcal{D}_n from ρ*

$$\mathcal{E}(\hat{G}_{r,\gamma}^{\text{RRR}}) \geq \sigma_{r+1}(A_\pi S) - c \varepsilon_n^* \ln \delta^{-1},$$

and

$$\mathcal{E}(\hat{G}_{r,\gamma}^{\text{PCR}}) \geq \sigma_{r+1}(S) - \sqrt{\mathcal{R}_0^{r+1}} - c \varepsilon_n^* \ln \delta^{-1},$$

where $\mathcal{R}_0^{r+1} := \langle (SS^* - A_\pi S S^* A_\pi^*) \ell_{r+1}, \ell_{r+1} \rangle \geq 0$ is the irreducible risk restricted to the $(r+1)$ -st left singular function ℓ_{r+1} of S .

The Example 3 in the main body, that we now elaborate, shows that lower bounds of the previous theorem are tight.

Example 4 (Example3). *Consider the 1D equidistant sampling of the Ornstein–Uhlenbeck process, obtained by integrating Langevin equation of Example 1 with $\beta = 1$ and $U(x) = x^2/2$*

$$X_t = e^{-1} X_{t-1} + \sqrt{1 - e^{-2}} \epsilon_t$$

where $x \in \mathbb{R} =: \mathcal{X}$ and $\{\epsilon_t\}_{t \geq 1}$ are independent standard Gaussians. For this process it is well-known that π is $\mathcal{N}(0, 1)$ and that A_π admits a spectral decomposition $(\mu_i, f_i)_{i \in \mathbb{N}}$ in terms of Hermite polynomials. Namely, $\{\mu_j = e^{-(j-1)}\}_{j \geq 1}$ with corresponding orthonormal eigenbasis $f_j(x) = \frac{1}{\sqrt{(j-1)!}} H_{j-1}(x)$ where $H_j(x) = \partial_1^j G(0, x)$, for $G(s, x) = e^{sx - \frac{1}{2}s^2}$, denotes the j -th probabilists' Hermite polynomial. See e.g. [23] for more details.

We now make use of the spectral decomposition of A_π to design the following class of kernel functions

$$k_{\Pi, \nu}(x, x') := \sum_{i \in N} \mu_{\Pi(i)}^{2\nu} f_i(x) f_i(x'),$$

where Π is a permutation of the indices of the eigenvalues and ν is a scaling factor. The rationale behind these kernels is to morph the original metric structure of A_π in a way which is harder and harder to revert when learning from finite sets of data. In particular, for any target rank r , we set $\nu := 1/r^2$ and Π to the permutation such that $i \mapsto 2r - i + 1$ ($i \geq r$), $i \mapsto i - r$ ($r + 1 \leq i \leq 2r$) and $i \mapsto i$ elsewhere. Then, we immediately have $\mathcal{R}_0^{r+1} := \langle (SS^* - ZZ^*) \ell_{r+1}, \ell_{r+1} \rangle = \mu_{r+1}^{2\nu} - \mu_{r+1}^{2\nu} \mu_1^2 = 0$, while $\sigma_{r+1}(S) = \sqrt{\mu_{r+1}^{2\nu}} = e^{-\nu r} = e^{-1/r}$. On the other hand, $\sigma_{r+1}(Z) = \sqrt{\mu_{r+1}^2 \mu_1^{2\nu}} = e^{-r}$. In view of Theorem 6, observing that in this example we can set $\beta \rightarrow 0$,

$$|\mathcal{E}(\hat{G}_{r,\gamma}^{\text{PCR}}) - e^{-1/r}| \lesssim n^{-1/2} \ln \delta^{-1},$$

and

$$|\mathcal{E}(\hat{G}_{r,\gamma}^{\text{RRR}}) - e^{-r}| \lesssim n^{-1/2} \ln \delta^{-1}.$$

D.5 Minimax optimal operator norm bounds

In this section we first derive a minimax lower bound for the operator norm convergence rate of the Koopman operator. This approach follows a standard framework already used for instance in [18, 29] to derive a lower bound on the Hilbert-Schmidt norm convergence rate of the Conditional Mean Embedding. We typically assume (as in [18, 29]) that condition (SD) is sharp, i.e. there exists $\beta \in (0, 1]$ and a constant $b > 1$ such that for every $j \in J$

$$b^{-1} j^{-1/\beta} \leq \lambda_j(C) \leq b j^{-1/\beta}. \quad (64)$$

Then we prove that the learning rate of the RRR estimator is sharp (up to a log factor).

Theorem 7. *Let $0 < \beta \leq \tau \leq 1$ and $\alpha \in (0, 2]$ be such that conditions (SD), (RC) and (KE) hold. Then for any $r \geq 2$, there exist absolute constants $c, q > 0$ such that for all learning methods $\mathcal{D}_n \rightarrow \hat{G} \in \mathcal{B}_r(\mathcal{H})$, $\delta \in (0, 1)$, and all sufficiently large $n \geq 1$, there is a distribution ρ on $\mathcal{X} \times \mathcal{X}$ used to sample \mathcal{D}_n , with marginal distribution π on \mathcal{X} , such that with probability at least $1 - \delta$,*

$$\mathcal{E}(\hat{G}) \geq c \delta^q n^{-\frac{\alpha \vee \tau}{2((\alpha \vee \tau) + \beta)}}.$$

Proof. We remind first that condition **(SRC)** is more restrictive than condition **(RC)** (see Remark 6). This means that a minimax lower under Condition **(SRC)** also holds under less restrictive Condition **(RC)**. We remark next that a Koopman operator is fully characterized by the conditional distribution of $X_{t+1}|X_t$. Therefore we will consider the set of well-separated conditional distributions on \mathcal{X} introduced in [20, 29]: $\{p_j(\cdot|x), j \in [M]\}$. Hence the corresponding set of Koopman operators $\Omega = \{A_j, j \in [M]\}$ admits rank at most 2 and when restricted to the RKHS \mathcal{H} satisfies $g_p(x) = \mathbb{E}[\phi(X_{t+1})|X_t = x] = \int_{\mathcal{X}} p_j(dy|x)\phi(y)$.

Fix now $a \in \mathcal{X}$ such that $\|\phi(a)\| \geq c_{\mathcal{H}}/2$ (such a always exists under condition **(KE)**). The boundedness of the kernel also guarantees that

$$\int \langle p_j(dy|x) - p_{j'}(dy|x), \phi(a) \rangle^2 \pi(dx) \leq \|\phi(a)\|^2 \|(A_j - A_{j'})S\|^2 \leq c_{\mathcal{H}} \|(A_j - A_{j'})S\|^2.$$

Hence it is sufficient to lower bound the left-hand-side to guarantee the existence of well-separated hypothesis. It follows from Lemma 7 in [29] that $M \geq 2^c n^{\frac{\beta}{(\alpha \vee \tau) + \beta}}$ and

$$\frac{1}{c c_{\mathcal{H}}} n^{-\frac{\alpha \vee \tau}{(\alpha \vee \tau) + \beta}} \leq \|(A_j - A_{j'})S\|^2 \leq c n^{-\frac{\alpha \vee \tau}{(\alpha \vee \tau) + \beta}},$$

for some absolute constant $c > 1$. Then, we can apply Lemma 8 in [29] to get the result. \square

This result states that, under conditions **(SD)**, **(RC)** and **(KE)**, for $\alpha \geq 1$ no estimator can achieve a learning rate faster than $n^{-\frac{\alpha}{2(\alpha+\beta)}}$ in the operator norm. Note that this lower bound is matching our upper bound on the variance term (up to $\ln \delta^{-1}$) in (16a) and (16b) of the RRR and PCR estimators respectively. We note however that a significant difference between these bounds concerns the "bias" term which is equal to $\sigma_{r+1}(A_{\pi}S)$ for RRR and $\sigma_{r+1}(S)$ for PCR.

Hence, to have an optimal learning rate for PCR, due to **((SD))**, we need to choose $r \geq n^{\frac{\alpha\beta}{\alpha+\beta}} - 1$. On the other hand, assuming for instance that $A_{\pi}S$ admits finite rank r , then the bias disappears completely in (16a) and consequently the RRR estimator is minimax optimal (up to a log factor). Moreover, even if the Koopman operator is infinite-dimensional with e.g. eigenvalues decaying exponentially fast, then the bias term is negligible in front of the variance term. Thus RRR is still minimax (up to a log).

D.6 Extension to misspecified setting

Next, recalling results from [29] on HS-norm error of the KRR estimator, we extend our results for RRR and PCR beyond the case **(RC)** for $\alpha \in [1, 2]$. Namely, in the reminder of this section we will assume **(SRC)** instead of **(RC)**, i.e. we will assume that $G_{\mathcal{H}}^{\alpha}$ is not only bounded by also a HS-operator when $\alpha < 1$.

Theorem 8. *Let **(SRC)**, **(SD)** and **(KE)** hold for some $\alpha \in (0, 2]$, $\beta \in (0, 1]$ and $\tau \in [\beta, 1]$, respectively, and let $\text{cl}(\text{Im}(S)) = L_{\pi}^2(\mathcal{X})$. Let*

$$\gamma \asymp n^{-\frac{1}{\alpha+\beta}} \text{ and } \varepsilon_n^* := n^{-\frac{\alpha}{2(\alpha+\beta)}} \text{ when } \alpha \geq \tau - \beta, \quad (65)$$

and for some $s > 1$

$$\gamma \asymp \left(\frac{n}{\ln^s n}\right)^{-\frac{1}{\tau}} \text{ and } \varepsilon_n^* := \left(\frac{n}{\ln^s n}\right)^{-\frac{\alpha}{2\tau}} \text{ when } \alpha < \tau - \beta. \quad (66)$$

Then, for every $\delta \in (0, 1)$ there exists a constant $c > 0$, depending only on \mathcal{H} , such that for large enough $n \geq r$, with probability at least $1 - \delta$ in the i.i.d. draw of \mathcal{D}_n from ρ

$$\mathcal{E}(\widehat{G}_{\gamma}) \leq c \varepsilon_n^* \ln \delta^{-1}.$$

Proof. Recalling that $\mathcal{E}(\widehat{G}_{\gamma}) = \|A_{\pi}S - S\widehat{G}_{\gamma}\| \leq \|A_{\pi}S - S\widehat{G}_{\gamma}\|_{\text{HS}}$, this result is a direct consequence of [29, Theorem 2].

In particular, recalling notation denoting $M = C_{\gamma}^{-1/2}(\widehat{T} - T)$ and $N = C_{\gamma}^{-1/2}(\widehat{C} - C)$ we have that

$$\|M - NG_{\gamma}\|_{\text{HS}} \leq \varepsilon_n^2(\gamma, \delta) \quad \text{w.p. at least } 1 - \delta \quad (67)$$

where

$$\tilde{\varepsilon}_n^2(\gamma, \delta) := 4\sqrt{2} \ln \frac{2}{\delta} \sqrt{\frac{c_{\mathcal{H}} \operatorname{tr}(C_{\gamma}^{-1}C)}{n} + \frac{4c_{\mathcal{H}}c_{\tau}}{n^2\gamma^{\tau}} + \frac{a^2c_{\tau}}{n}\gamma^{\alpha-\tau}}, \quad (68)$$

and, since $R = C_{\gamma}^{1/2}(\hat{G}_{\gamma} - G_{\gamma}) = C_{\gamma}^{1/2}\hat{C}_{\gamma}^{-1}C_{\gamma}^{1/2}(M - NG_{\gamma})$, consequently,

$$\mathbb{P} \left\{ \|C_{\gamma}^{1/2}(\hat{G}_{\gamma} - G_{\gamma})\|_{\text{HS}} \leq \frac{\tilde{\varepsilon}_n^2(\gamma, \delta/2)}{1 - \varepsilon_n^1(\gamma, \delta/2)} \right\} \geq 1 - \delta.$$

Combining this with the approximation bound of Proposition 5, we obtain the result. \square

Therefore, when $\alpha > \tau - \beta$ previous theorem shows that the operator norm error of KRR estimator is optimal. It is important to note that regime $\alpha \leq \tau - \beta$ the optimal rates are an open problem in operator learning as well as in the classical regression setting. We conclude this section by extending Theorem 2 and showing that PCR and RRR also achieve optimal rates when $\alpha > \tau - \beta$.

Theorem 9. Assume the operator A_{π} satisfies $\sigma_r(A_{\pi}S) > \sigma_{r+1}(A_{\pi}S) \geq 0$ for some $r \in \mathbb{N}$. Let (SRC), (SD) and (KE) hold for some $\alpha \in (0, 2]$, $\beta \in (0, 1]$ and $\tau \in [\beta, 1]$, respectively, and let $\operatorname{cl}(\operatorname{Im}(S)) = L_{\pi}^2(\mathcal{X})$. Let γ and ε_n^* be given by (65)-(66) and let $\delta \in (0, 1)$. Then, there exists a constant $c > 0$, depending only on \mathcal{H} , such that for large enough $n \geq r$, with probability at least $1 - \delta$ in the i.i.d. draw of \mathcal{D}_n from ρ

$$\mathcal{E}(\hat{G}) \leq \begin{cases} \sigma_{r+1}(A_{\pi}S) + c\varepsilon_n^* \ln \delta^{-1} & \text{if } \hat{G} = \hat{G}_{r,\gamma}^{\text{RRR}}, \\ \sigma_{r+1}(S) + c\varepsilon_n^* \ln \delta^{-1} & \text{if } \hat{G} = \hat{G}_{r,\gamma}^{\text{PCR}} \text{ and } \sigma_r(S) > \sigma_{r+1}(S). \end{cases} \quad (69a)$$

$$\quad (69b)$$

Proof. Proof of (69b) readily follows from (59) by applying (67) instead of Proposition 16, i.e. using $\tilde{\varepsilon}_n^2$ given in (68) instead of ε_n^2 given in (48). To show (69a), recalling (55), note that

$$\hat{B}^* \hat{B} - B^* B = B^*(M - NG_{\gamma}) + M^* B + (M - NG_{\gamma})^* C_{\gamma}^{1/2} \hat{C}_{\gamma}^{-1} C_{\gamma}^{1/2} (M - NG_{\gamma})$$

we have that w.p. $1 - \delta$

$$\|\hat{B}^* \hat{B} - B^* B\| \leq \sqrt{c_{\mathcal{H}}} [\tilde{\varepsilon}_n^2(\gamma, \delta/3) + \varepsilon_n^2(\gamma, \delta/3)] + \frac{[\tilde{\varepsilon}_n^2(\gamma, \delta/3)]^2}{1 - \varepsilon_n^1(\gamma, \delta/3)} \lesssim \tilde{\varepsilon}_n^2(\gamma, \delta/3).$$

Therefore, since

$$C_{\gamma}^{1/2}(G_{r,\gamma}^{\text{RRR}} - \hat{G}_{r,\gamma}^{\text{RRR}}) = C_{\gamma}^{1/2}[\hat{G}_{\gamma} - G_{\gamma}]\hat{P}_r + B[\hat{P}_r - P_r],$$

we conclude that w.p. $1 - \delta$

$$\|S(G_{r,\gamma}^{\text{RRR}} - \hat{G}_{r,\gamma}^{\text{RRR}})\| \lesssim \tilde{\varepsilon}_n^2(\gamma, \delta/3) \left(1 + \frac{\sigma_1(B)}{\sigma_r^2(B) - \sigma_{r+1}^2(B)} \right),$$

which after balancing with the approximation bound of Proposition 5 yields the final result. \square

E Spectral Learning Rates

We now prove the statements of Theorems 3 and 4 in slightly more general form.

Theorem 10 (RRR). Let A_{π} be a compact self-adjoint operator. Under the assumptions of Theorem 3, if $\hat{G} = \hat{G}_{r,\gamma}^{\text{RRR}}$, then there exists a constant $c > 0$ (depending only on the RKHS) such that for every $\delta \in (0, 1)$, every large enough $n \geq r$ and every $i \in [r]$ with probability at least $1 - \delta$ in the i.i.d. draw of \mathcal{D}_n from ρ

$$|\hat{\lambda}_i - \mu_{j(i)}| \leq \frac{\sigma_{r+1}(Z)}{\sigma_r(Z)} \left(1 + c n^{-\frac{\alpha-1}{2(\alpha+\beta)}} \ln \delta^{-1} \right) + c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}. \quad (70)$$

Moreover,

$$|\hat{\lambda}_i - \mu_{j(i)}| \leq \hat{\eta}_i \sigma_{r+1}(\hat{B}) + c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}, \quad (71)$$

and

$$\|\hat{f}_i - f_{j(i)}\|^2 \leq \frac{2(\hat{\eta}_i \sigma_{r+1}(\hat{B}) + c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1})}{[\operatorname{gap}_{j(i)}(A_{\pi}) - \hat{\eta}_i \sigma_{r+1}(\hat{B}) - c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}]_+}. \quad (72)$$

Proof. First, observe that Proposition 1 and Weyl's inequality imply that for every $i \in [r]$

$$\eta(\widehat{\psi}_i) \leq \frac{(|\widehat{\lambda}_i| \text{cond}(\widehat{\lambda}_i)) \wedge \|\widehat{G}_{r,\gamma}^{\text{RRR}}\|}{[\sigma_r(SG_{r,\gamma}^{\text{RRR}}) - \|S(\widehat{G}_{r,\gamma}^{\text{RRR}} - G_{r,\gamma}^{\text{RRR}})\|]_+} \leq \frac{(|\widehat{\lambda}_i| \text{cond}(\widehat{\lambda}_i)) \wedge \|\widehat{G}_\gamma\|}{[\sigma_r(SG_{r,\gamma}^{\text{RRR}}) - \|S(\widehat{G}_{r,\gamma}^{\text{RRR}} - G_{r,\gamma}^{\text{RRR}})\|]_+}.$$

But,

$$\sigma_r(SG_{r,\gamma}^{\text{RRR}}) \geq \sigma_r(C_\gamma^{1/2} G_{r,\gamma}^{\text{RRR}}) - \sqrt{\gamma} \|G_{r,\gamma}^{\text{RRR}}\| \geq \sigma_r(B) - \sqrt{\gamma} \|\widehat{G}_\gamma\|,$$

and, using Proposition 6 and $\|G_\gamma\| \leq 1$ we conclude

$$\eta(\widehat{\psi}_i) \leq \frac{(|\widehat{\lambda}_i| \text{cond}(\widehat{\lambda}_i)) \wedge \|\widehat{G}_\gamma\|}{[\sigma_r(Z) - a\|C\|^{\alpha/4} \gamma^{\alpha/4} - \sqrt{\gamma} - \|S(\widehat{G}_{r,\gamma}^{\text{RRR}} - G_{r,\gamma}^{\text{RRR}})\|]_+}.$$

Hence, in view of Propositions 16 and 18, we conclude that for some constant $c' > 0$

$$\eta(\widehat{\psi}_i) \leq \frac{(|\widehat{\lambda}_i| \text{cond}(\widehat{\lambda}_i)) \wedge (1 + c' n^{-\frac{\alpha-1}{2(\alpha+\beta)}} \ln \delta^{-1})}{[\sigma_r(Z) - c' n^{-\frac{\alpha}{4(\alpha+\beta)}} \ln \delta^{-1}]_+}$$

and, consequently, for some $c > 0$,

$$\eta(\widehat{\psi}_i) \leq 2 \frac{|\widehat{\lambda}_i| \text{cond}(\widehat{\lambda}_i)}{\sigma_r(Z)} \wedge \left(\frac{2}{\sigma_r(Z)} + c n^{-\frac{\alpha-1}{2(\alpha+\beta)}} \ln \delta^{-1} \right).$$

Therefore, using (63), (70) directly follows.

Next, since we have obtain that $\eta(\widehat{\psi}_i)$ is bounded, due to Proposition 4, empirical metric distortions $\widehat{\eta}_i$ are bounded, too.

So, assuming that $\sigma_{r+1}(Z) > 0$, we have that

$$\begin{aligned} \eta(\widehat{\psi}_i) \mathcal{E}(\widehat{G}_{r,\gamma}^{\text{RRR}}) &\leq \eta(\widehat{\psi}_i) \sigma_{r+1}(B) + \eta(\widehat{\psi}_i) \left(\mathcal{E}(\widehat{G}_{r,\gamma}^{\text{RRR}}) - \sigma_{r+1}(B) \right) \\ &\leq \widehat{\eta}_i \sigma_{r+1}(B) + \widehat{\eta}_i \eta^2(\widehat{\psi}_i) \varepsilon_n(\delta/3) + \eta(\widehat{\psi}_i) \left(\mathcal{E}(\widehat{G}_{r,\gamma}^{\text{RRR}}) - \sigma_{r+1}(B) \right) \\ &\leq \widehat{\eta}_i \sigma_{r+1}(\widehat{B}) + \widehat{\eta}_i \frac{|\sigma_{r+1}^2(\widehat{B}) - \sigma_{r+1}^2(B)|}{\sigma_{r+1}(B)} + \widehat{\eta}_i \eta^2(\widehat{\psi}_i) \varepsilon_n(\delta/3) \\ &\quad + \eta(\widehat{\psi}_i) \left(\mathcal{E}(\widehat{G}_{r,\gamma}^{\text{RRR}}) - \sigma_{r+1}(B) \right). \end{aligned}$$

Thus, recalling Proposition 17, and using (63) we obtain that there exists a constant $c > 0$ depending only on the RKHS so that

$$\eta(\widehat{\psi}_i) \mathcal{E}(\widehat{G}_{r,\gamma}^{\text{RRR}}) \leq \widehat{\eta}_i \sigma_{r+1}(\widehat{B}) + c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}.$$

Finally, applying Theorem 1 concludes the proof. \square

Specifying the previous result to finite rank Koopman operators we obtain the following.

Corollary 1 (RRR). *If A_π is of finite rank $r \in \mathbb{N}$, under the assumptions of Theorem 10, with probability at least $1 - \delta$ in the i.i.d. draw of $(x_i, y_i)_{i=1}^n$ from ρ*

$$|\widehat{\lambda}_i - \mu_{j(i)}| \leq c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}, \quad (73)$$

and

$$\|\widehat{f}_i - f_{j(i)}\|^2 \leq \frac{2 c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}}{[\text{gap}_i(\widehat{G}_{r,\gamma}^{\text{RRR}}) - 3 c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}]_+}. \quad (74)$$

Proof. First, observe that in the previous proof we could conclude the same having $\sigma_{r+1}(Z)$ instead of $\sigma_{r+1}(\widehat{B})$. But, since rank of Z is at most r , we conclude (73). Next, since the ranks of $\widehat{G}_{r,\gamma}^{\text{RRR}}$ and A_π are equal, we have the same number of nonzero eigenvalues. Hence, in view of the eigenvalue ordering $\widehat{\lambda}_1 \geq \widehat{\lambda}_2 \geq \dots \widehat{\lambda}_r$ and $\mu_{j_1} \geq \mu_{j_2} \geq \dots \mu_{j_r}$, we obtain

$$\text{gap}_i(\widehat{G}_{r,\gamma}^{\text{RRR}}) = |\widehat{\lambda}_i - \widehat{\lambda}_{i-1}| \wedge |\widehat{\lambda}_i - \widehat{\lambda}_{i+1}| \leq \text{gap}_i(A_\pi) + |\widehat{\lambda}_i - \mu_{j(i)}| + (|\widehat{\lambda}_{i-1} - \mu_{j_{i-1}}| \vee |\widehat{\lambda}_{i+1} - \mu_{j_{i+1}}|).$$

So, applying (73), we conclude (74). \square

In analogous way we have the following theorem for PCR estimator.

Theorem 11 (PCR). *Let A_π be a compact self-adjoint operator. Under the assumptions of Theorem 3, if $\hat{G} = \hat{G}_{r,\gamma}^{\text{PCR}}$, then there exists a constant $c > 0$ (depending only on the RKHS) such that for every $\delta \in (0, 1)$, every large enough $n \geq r$ and every $i \in [r]$ with probability at least $1 - \delta$ in the i.i.d. draw of \mathcal{D}_n from ρ*

$$|\hat{\lambda}_i - \mu_{j(i)}| \leq \frac{\sigma_{r+1}(S)}{[\sigma_r(Z) - \sigma_{r+1}^\alpha(S)]_+} \left(2 + c n^{-\frac{\alpha-1}{2(\alpha+\beta)}} \ln \delta^{-1} \right) + c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}. \quad (75)$$

Moreover,

$$|\hat{\lambda}_i - \mu_{j(i)}| \leq \hat{\eta}_i \sigma_{r+1}(\hat{S}) + c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}, \quad (76)$$

and

$$\|\hat{f}_i - f_{j(i)}\|^2 \leq \frac{2(\hat{\eta}_i \sigma_{r+1}(\hat{S}) + c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1})}{[\text{gap}_{j(i)}(A_\pi) - \hat{\eta}_i \sigma_{r+1}(\hat{S}) - c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}]_+}. \quad (77)$$

Proof. The proof follows along the same lines of reasoning as for the RRR estimator, with the exception that we use

$$\sigma_r(SG_{r,\gamma}^{\text{PCR}}) = \sigma_r(P_r C^{1/2} G_\gamma) \geq \sigma_r(C^{1/2} G_\gamma) - \|(I - P_r) C^{1/2} C_\gamma^{-1} T\| \geq \sigma_r(B) - \sqrt{\gamma} - \sigma_{r+1}^\alpha(S),$$

where P_r denotes the projector onto a subspace of r leading eigenfunctions of C . \square

We finally remark that from the proofs of Theorems 11 and 11, recalling notation of the uniform bias in Eq. (17), we also have

$$\frac{|\hat{\lambda}_i - \mu_{j(i)}|}{|\hat{\lambda}_i| \text{cond}(\hat{\lambda}_i)} \leq s + c n^{-\frac{\alpha}{2(\alpha+\beta)}} \ln \delta^{-1}, \quad (78)$$

which may also be of interest.

Remark 8 (Limitation to (RC) for $\alpha \geq 1$). *Note that, according to Proposition 5, G_γ may become unbounded as $\gamma \rightarrow 0$ for $\alpha < 1$. This, recalling Proposition 1, implies that the metric distortion of estimated eigenfunctions may deteriorate. So, proving tight spectral rates in the misspecified regime remains an interesting open problem.*

F Experiments

A Realistic Example: Langevin Dynamics. We here append some additional data on the Langevin dynamics experiment not fitting in the main body. We recall that the eigenpairs of A_π in this experiment have a tangible physical interpretation [37]. Indeed μ_i is related to the typical time scale needed for a particle to cross one of the potential barriers, while f_i is approximately constant in the regions of the phase space where the particle spends a lot of time (metastable states) and have sharp variations near the unstable points of the dynamics (transition states)

In Figure 4 we show the eigenvalue and eigenfunction errors, and fitted the lines appearing in the log-log plots to get the decay rates for the eigenvalue and eigenfunctions errors reported in Table 2 and 3.

	RRR	PCR
$\hat{\lambda}_0$	-0.801699	-0.009406
$\hat{\lambda}_1$	-0.617175	-0.410670
$\hat{\lambda}_2$	-0.553720	-0.631005
$\hat{\lambda}_3$	-0.699849	-0.697323

Table 2: Fitted decay rates for the estimated Koopman eigenvalues for the Langevin dynamics experiment




Figure 4: Eigenvalue errors (left pair) and Eigenfunction errors (right pair) vs. sample size. Solid, dashed, dash-dotted, and dotted lines denote the error on the four leading eigenpairs, in order. Both axes are in log scale, so straight lines correspond to a polynomial decay of the error with rate the line slope. RRR estimator is unbiased and attains smaller errors than the biased PCR one.

	RRR	PCR
\hat{f}_0	-0.635680	-0.139309
\hat{f}_1	-0.577666	-0.326457
\hat{f}_2	-0.573375	-0.505743
\hat{f}_3	-0.476403	-0.450845

Table 3: Fitted decay rates for the estimated Koopman eigenfunctions for the Langevin dynamics experiment.

Model selection and the Alanine dipeptide dataset We use a simulation of the small molecule Alanine dipeptide reported in Ref. [44]. For each RRR estimator we set the rank $r = 5$ and the Tikhonov regularization $\gamma = 10^{-6}$. The 19 different kernels are the following: 7 RBF kernels with length scales $\sigma \in \{0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35\}$, and 12 Matérn kernels corresponding to each possible combination of $\nu \in \{1.5, 2.5\}$ and length scale $\sigma \in \{0.05, 0.1, 0.15, 0.2, 0.25, 0.3\}$. The optimal kernel turned out to be RBF with $\sigma = 0.35$. We report forecasting RMSE of $30 = 3 * 10$ positions of the 10 atoms of the system.