Finite-data Error Bounds for Approximating the Koopman Operator

Daniel Fassler

A Thesis in
The Department of
Mathematics and Statistics

Presented in Partial Fulfillment of the Requirements
for the Degree of Master of Science (Mathematics) at
Concordia University
Montréal, Québec, Canada

July, 2025

Bornes d'erreur pour approximation de l'opérateur de Koopman avec données finies

Daniel Fassler

Mémoire présenté
au Département de
Mathématiques et Statistiques

comme exigence partielle au grade de
Maître ès Sciences (Mathématiques)
Université Concordia
Montréal, Québec, Canada

Juillet, 2025

CONCORDIA UNIVERSITY
School of Graduate Studies

This is to certify that the thesis prepared
By: **Daniel Fassler**
Entitled: **Finite-Data Error Bounds for Approximating the Koopman Operator**
And submitted in partial fulfillment of the requirements for the degree of

Master of Science (Mathematics)

complies with the regulations of the University and meets the accepted standards with respect to originality and quality.
Signed by the final Examining Committee:

_____ Chair
Dr. Marco Bertola

_____ Thesis Supervisor
Dr. Jason J. Bramburger

_____ Thesis Supervisor
Dr. Simone Brugiappaglia

_____ Examiner
Dr. Maria Ntekoume

Approved by

_____ Dr. Lea Popovic, Graduate Program Director

_____ Dr. Pascale Sicotte, Dean of Faculty of Arts and Science

Date: June 2025

Abstract

Finite-data Error Bounds for Approximating the Koopman Operator

Daniel Fassler

The Koopman operator is a powerful tool for the study of dynamical systems that allows the study of nonlinear systems through the lens of observables functions forming an equivalent linear formulation of the dynamics in an infinite-dimensional space. Recent developments in data-driven approximation techniques allow the approximation of the Koopman operator without any a priori knowledge of the underlying system. However, despite the existence of asymptotic results on the capabilities of such techniques, only a few guarantees exist in the more realistic finite-data setting.

Approximation of the Koopman operator through Extended Dynamic Mode Decomposition (EDMD) has been observed to converge at the Monte Carlo rate of the inverse, i.e., proportionally to the square root of number of samples. In this thesis, we bridge the gap between EDMD and the theory of least squares to provide a proof of this statement with minimal assumptions. Moreover, leveraging known results from function approximation via least squares, we investigate the effect of the sampling routine and the choice and size of dictionary on the convergence of the EDMD method. Additionally, we develop a similar approach in the context of compressed sensing, where we provide recovery guarantees when the Koopman operator is sparse. Finally, we validate theoretical findings through extensive numerical illustrations.

Acknowledgements

I would like to thank my supervisors Dr. Jason J. Bramburger and Dr. Simone Brugiapaglia for their continued support and guidance throughout this project. Without them, this work and this leg of my academic journey would not have been possible.

I would also like to thank Dr. Ben Adcock whose work and insight have been of tremendous help for this project and Dr. Maria Ntekoume for taking part in the Examining Committee.

I would like to thank my family for their continued support and ecouragement without whom none of this would have been possible in the first place.

Finally, I would like to thank Dr. Geoffrey McGregor and Dr. Jean-Christophe Nave for their guidance during my undergraduate studies, allowing me to find a place where my passion and skills could be put to good use.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Dynamical systems are a ubiquitous part of modern day science, from the study of movements of bodies arising from mechanics in physics or engineers trying to predict the motion of a system, to biologists trying to understand the behaviour of a population. Dynamical systems analysis aims to study time-evolving processes from a geometric perspective. It focuses on understanding salient features of systems, such as equilibria, periodic orbits and invariant manifolds. As such, it can be seen as analysis, geometry, topology and applied mathematics all working together to study real life systems. Because of its importance, it is crucial to have tools to study, understand and control such systems.

In his original work in 1931 [24], Bernard O. Koopman introduced a new concept that would become fundamental for the study of dynamical systems. Using the theory of Hilbert spaces, Koopman demonstrated how one could study a dynamical system through an equivalent linear dynamical system, albeit an infinite-dimensional one. To reach this goal, he used an operator now known as the Koopman operator which acts on *observables* functions of the system, scalar-valued functions on the phase space. It turns out that by studying how those observable functions evolve in time, one can infer the behaviour of the original system. Those functions are often already quantities of interest such as momentum, energy, or concentration of chemicals. This is the basis of the Koopman theory.

Koopman's work saw significant re-discovery in the late 1990s and early 2000s, and the recent success of the Koopman Theory has been dubbed "Koopmania" by some researchers. This moniker is not without merit, as Koopman theory has a very wide range of applications such as fluid dynamics [28], control theory [12] or even quantum computing [23] through Koopman-Von Neumann mechanics.

The rise of data science over the past decades was of tremendous importance for scientific experimentations, a context where the true underlying system is often unknown or too complicated to work with directly. The abundance of data and increase of computational power has also made it much more affordable to approximate the Koopman operator using only data gathered from the system. To this day, approximation of the Koopman operator is a very active field of research. While many theoretical results exist on the asymptotic convergence of such methods, the theory sometimes lacks practical results with easily verifiable assumptions. This is precisely the problem that this thesis aims to address.

We explore the approximation of the Koopman operator from data, both in the over sampled regime with an Extended Dynamic Mode Decomposition (EDMD) approach and

in the under sampled regime with a compressed sensing approach. We will provide error bounds for both approaches and a convergence rate for the EDMD method.

For the EDMD approach, we will provide a proof relying on minimal assumptions for the convergence rate of the method. While this result is not new, as it as been extensively observed and proved in different settings in the literature [30, 39, 40]. Our approach generalizes those proofs to both stochastic and deterministic systems, as well as leveraging the use of two different dictionaries, which allows to apply our work to the identification of Lyapunov functions [10] and the discovery of invariant measures [11]. Our proof unites results from the theory of least squares and the theory of EDMD to obtain new convergence guarantees. Finally, we demonstrate the versatility of our result through extensive testing on both continuous and discrete time dynamics, as well as stochastic dynamical systems.

For the compressed sensing approach, we provide new approximation schemes based on the EDMD approach in the compressed sensing setting. We present theoretical recovery guarantees on the sparsity of the recovered Koopman operator using methods such as LASSO and Quadratically constrained basis pursuit. Finally, we demonstrate the usefulness of these techniques on a high-dimensional dynamical system in continuous-time.

We now provide a roadmap of what will be discussed in this thesis. In Chapter 2, we will present a brief overview of the Koopman operator and its properties. We will also introduce EDMD as a method to approximate the Koopman operator. Finally, we will then review least-squares theory and compressed sensing theory, the tools we will use to derive our results. In Chapter 3, we will present our main results for approximating the Koopman operator in the Monte Carlo setting using EDMD and Least Square Theory, we give a proof of the convergence rate of $\frac{1}{\sqrt{n}}$, where n is the number of data points. We will also talk about exponential convergence by letting the dictionary size grow with the number of data points. In Chapter 4, we talk about approximation of the Koopman Operator in a different setting, compressed sensing. We will show that in the context of sparsity, one can derive sparsity recovery guarantees for methods like the Least Absolute Shrinkage and Selection Operator (LASSO) and Quadratically Constrained Basis Pursuit (QCBP). In Chapter 5, we conclude the thesis by discussing our results and providing some possible future directions for this work.

Chapter 2

Background

In this chapter, we will present the basic notions, notations, tools and definitions used throughout this thesis. We will start by presenting the analytical framework that we work in. We then introduce the Koopman operator, its properties and how to approximate it, and finally we explain the theory of function approximations through least squares.

2.1 The basics

In this section, we present the basic notation and tools we used throughout our work. We start with the idea of orthogonal bases of function spaces. Then we move on to lower sets, which are of particular importance in the context of high dimensional approximation.

2.1.1 Orthogonal bases of L^2_ρ

Throughout this work, we are interested in the approximation of functions in the space $L^2_\rho(\mathbb{X})$, defines as follows.

Definition 2.1 ($L^2_\rho(\mathbb{X})$ space)**.** *The set of square integrable functions with respect to the measure ρ on the domain $\mathbb{X} \subseteq \mathbb{R}^d$ is defined as*

$$L^2_\rho(\mathbb{X}) = \left\{ \varphi : \mathbb{X} \to \mathbb{R} : \int_{\mathbb{X}} |\varphi(x)|^2 d\rho(x) < +\infty \right\}.$$

It is equipped with the inner product

$$\langle f, g \rangle_{L^2_\rho(\mathbb{X})} = \int_{\mathbb{X}} f(x) g(x) d\rho(x), f, g, \in L^2_\rho(\mathbb{X}),$$

and the L^2-norm is defined as

$$\|f\|_{L^2_\rho(\mathbb{X})} = \sqrt{\langle f, f \rangle_{L^2_\rho(\mathbb{X})}} = \left(\int_{\mathbb{X}} |f(x)|^2 d\rho(x) \right)^{\frac{1}{2}}.$$

The space $L^2_\rho(\mathbb{X})$ is a Hilbert space with respect to the inner product $\langle \cdot, \cdot \rangle_{L^2_\rho(\mathbb{X})}$ and the induced norm $\| \cdot \|_{L^2_\rho(\mathbb{X})}$.

Additionally, we need to define the notion of a *basis* for such a space. We note that throughout this thesis, we will use multi-index notation to denote basis functions in multiple dimension. We denote by $\boldsymbol{\nu} = (\nu_k)_{k=1}^d$ a multi-index in $\mathbb{N}_0^d$.

Definition 2.2 (Basis)**.** *A set of functions $\{\psi_{\boldsymbol{\nu}} : \nu \in \mathbb{N}_0^d\}$ is a basis for a subspace $\mathcal{P}$ of $L^2_\rho(\mathbb{X})$ if the following holds:*

$$\forall \varphi \in \mathcal{P}, \exists c_{\boldsymbol{\nu}} \in \mathbb{R} : \varphi = \sum_{\boldsymbol{\nu} \in \Lambda} c_{\boldsymbol{\nu}} \psi_{\boldsymbol{\nu}}.$$

for $\Lambda \subseteq \mathbb{N}_0^d$. Moreover, the basis must satisfy span$\{\psi_{\boldsymbol{\nu}}\}_{\boldsymbol{\nu} \in \mathbb{N}_0^d} = \mathcal{P}$.

We recall and consider useful bases for $L^2_\rho(\mathbb{X})$, for $\mathbb{X} \subset \mathbb{R}^d$. To this end, we need to define the idea of orthogonality.

Definition 2.3 (Orthogonality)**.** *Given a basis $\{\psi_{\boldsymbol{\nu}} : \nu \in \mathbb{N}_0^d\}$ of L^2_ρ, we say that the basis is orthogonal if for all $\boldsymbol{\nu}, \boldsymbol{\mu} \in \mathbb{N}_0^d$,*

$$\langle \psi_{\boldsymbol{\nu}}, \psi_{\boldsymbol{\mu}} \rangle := \int_{\mathbb{X}} \psi_{\boldsymbol{\nu}} \psi_{\boldsymbol{\mu}} \, d\rho = 0, \forall \boldsymbol{\nu} \neq \boldsymbol{\mu}. \tag{2.1}$$

Moreover, if additionally we have that $\|\psi_{\boldsymbol{\nu}}\|_{L^2_\rho(\mathbb{X})} := \sqrt{\langle \psi_{\boldsymbol{\nu}}, \psi_{\boldsymbol{\nu}} \rangle} = 1, \forall \boldsymbol{\mu} \in \mathbb{N}_0^d$, we say that the basis is orthonormal.

Throughout this work, we consider different choices of ρ and their corresponding orthogonal bases. For example, we consider the uniform measure on $\mathbb{X} = [-1, 1]^d, d\rho(\mathbf{x}) = 2^{-d} d\mathbf{x}$. In this case, we consider the following orthogonality bases: The Legendre polynomials and the Fourier basis. The Legendre polynomials are defined through Rodrigues' formula

$$P_n(x) = \frac{1}{2^n n!} \frac{d^n}{dx^n} (x^2 - 1)^n.$$

The first few are given in Table 2.1, it can be shown that with the uniform measure, the Legendre polynomials satisfy

$$\int_{-1}^{1} P_n(x) P_m(x) dx = \frac{2}{2n+1} \delta_{nm}.$$

For $d \geq 2$, we consider the following multi-variate basis of Legendre polynomials:

$$\psi_{\boldsymbol{\nu}} = P_{\nu_k} \otimes \cdots \otimes P_{\nu_d}. \tag{2.2}$$

The Fourier basis on $[-1, 1]$ is comprised of functions of the form $\sin(n\pi x), \cos(n\pi x)$ for $n \in \mathbb{N}_0$. Similarly to form a basis on $[-1, 1]^d$ we consider the tensor product of these functions, as in (2.2). It can be shown that the Fourier basis is orthonormal with respect to the uniform measure on $[-1, 1]^d$.

Another family of orthogonal polynomials that we consider are the Chebyshev polynomials of the first kind. They are defined on the interval $[-1, 1]$ and are orthogonal with respect to the measure $\rho(x) = \frac{1}{\pi\sqrt{1-x^2}}$. The Chebyshev polynomials are defined as:

$$T_n(x) = \cos(n \arccos(x)).$$

Degree n	Legendre Polynomial $P_n(x)$	Chebyshev Polynomial $T_n(x)$
0	1	1
1	x	x
2	$\frac{1}{2}(3x^2-1)$	$2x^2-1$
3	$\frac{1}{2}(5x^3-3x)$	$4x^3-3x$
4	$\frac{1}{8}(35x^4-30x^2+3)$	$8x^4-8x^2+1$
5	$\frac{1}{8}(63x^5-70x^3+15x)$	$16x^5-20x^3+5x$

Table 2.1: First few Legendre polynomials $P_n(x)$ and Chebyshev polynomials of the first kind $T_n(x)$.

The first few Chebyshev polynomials are given in Table 2.1. When $d \geq 2$, one can define the multivariate Chebyshev polynomials as

$$\psi_{\boldsymbol{\nu}}(\mathbf{x}) = 2^{\|\boldsymbol{\nu}\|_0/2} \prod_{k=1}^{d} \cos(\nu_k \arccos(x_k)),$$

where $\|\cdot\|_0$ counts the number of non-zero entries of a vector. They are orthogonal with respect to the measure $d\rho(\mathbf{x}) = \prod_{k=1}^{d} \frac{1}{\pi\sqrt{1-x_k}} d\mathbf{x}$.

We note that both the Legendre polynomials and Chebyshev polynomials are part of the family of the Jacobi polynomials which are orthogonal with respect to the measure $d\rho(\mathbf{x}) = \prod_{k=1}^{d}(1-x_k)^{\alpha}(1+x_k)^{\beta} d\mathbf{x}$. The Jacobi polynomials are defined as

$$P_n^{(\alpha,\beta)}(x) = \frac{(-1)^n}{2^n n!}(1-x)^{-\alpha}(1+x)^{-\beta}\frac{d^n}{dx^n}\left((1-x)^{\alpha}(1+x)^{\beta}(1-x^2)^n\right). \tag{2.3}$$

The Legendre polynomials are obtained when $\alpha = \beta = 0$ and the Chebyshev polynomials are obtained when $\alpha = \beta = -1/2$. We will sometimes use the Jacobi Polynomials with $\alpha = 1, \beta = 0$ to study the effect of having an orthogonality measure and polynomials that are non-symmetric (skewed). Finally, we note that most of these bases are orthogonal but not orthonormal, and so we will normalize them to have unit norm $\|\psi_{\boldsymbol{\nu}}\|_{L^2_\rho(\mathbb{X})} = 1$.

2.1.2 Lower sets

In the high dimensional setting, the size of our sample sets grows exponentially if we consider a tensor product basis. This is a problem as it makes numerical experiments costly and impractical. To alleviate this, we can consider using different multi-index sets that are smaller but encompasses the most important multi indices. Below are a few examples compared to the tensor product index set.

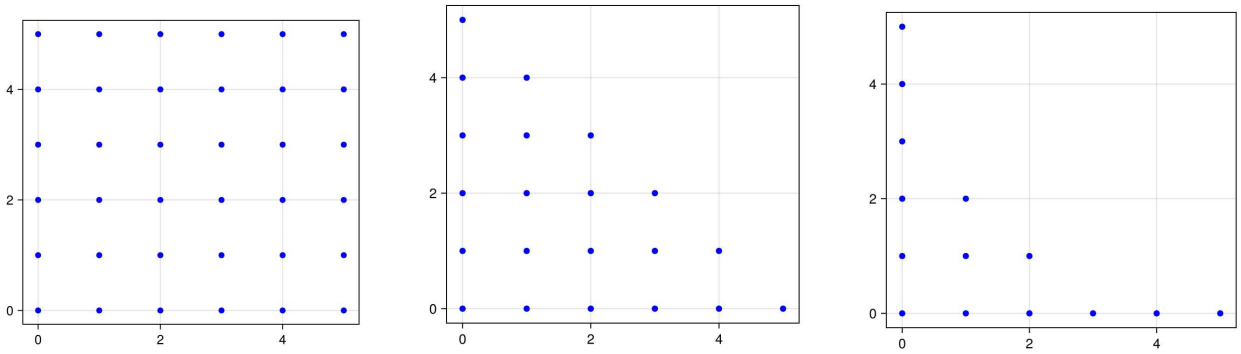

(a) Λ_5^{TP} (b) Λ_5^{TD} (c) Λ_5^{HC}

Figure 2.1: Comparison of the different multi-index sets. The left figure shows the tensor product set of degree 5 Λ_5^{TP}, the middle figure shows the total degree set of degree 5 Λ_5^{TD} and the right figure shows the hyperbolic cross set of degree 5 Λ_5^{HC}.

Definition 2.4 (Multi-index sets.)**.** *We define the following multi-index sets:*

$$\Lambda_n^{TP} = \left\{ \boldsymbol{\nu} = (\nu_k)_{k=1}^d \in \mathbb{N}_0^d : \max_{k\in[d]} \nu_k \leq n \right\}, \qquad \textit{(Tensor product set)}$$

$$\Lambda_n^{TD} = \left\{ \boldsymbol{\nu} = (\nu_k)_{k=1}^d \in \mathbb{N}_0^d : \sum_{k=1}^d \nu_k \leq n \right\}, \qquad \textit{(Total degree set)}$$

$$\Lambda_n^{HC} = \left\{ \boldsymbol{\nu} = (\nu_k)_{k=1}^d \in \mathbb{N}_0^d : \prod_{k=1}^d (\nu_k + 1) \leq n+1 \right\}. \qquad \textit{(Hyperbolic cross set)}$$

All of the above are examples of *lower sets*, they are displayed in Figure 2.1.

Definition 2.5 (Lower sets)**.** *A multi-index set $\Lambda \subseteq \mathbb{N}_0^d$ is lower if the following holds for all $\nu, \mu \in \mathbb{N}_0^d$:*

$$\nu \in \Lambda, \mu \leq \nu \Rightarrow \mu \in \Lambda. \tag{2.4}$$

Lower sets are especially useful, as they allow reducing the number of indices we need to consider without losing too much information. Informally, we can think of lower sets as sets without any holes connected to the origin. The hyperbolic cross set is especially useful as it can be seen as the union of all lower sets up to a given size.

Proposition 2.1 (Union of lower set is hyperbolic cross set, from [3, Chapter 2])**.** *Let $d, n \in \mathbb{N}$. Then,*

$$\bigcup \left\{ S \subseteq \mathbb{N}_0^d : |S| \leq n, S \textit{ lower} \right\} = \Lambda_{n-1}^{HC}.$$

Finally, we conclude this section by mentioning the cardinality of the hyperbolic cross set, as it is the one we will consider in our experiments.

Proposition 2.2 (Cardinality of the hyperbolic cross, from [3, Chapter 2])**.** *The hyperbolic cross set Λ_n^{HC} satisfies*

$$|\Lambda_{n-1}^{HC}| \sim \frac{n \log(n)^{d-1}}{(d-1)!}, \quad n \to \infty.$$

With all these basic notions introduced, we can start presenting the theory of the Koopman operator and its approximation.

2.2 The Koopman operator

In this section, we introduce the main object of the thesis, the Koopman operator, and how to approximate it from data.

2.2.1 The Lie derivative and the Koopman operator

We begin by introducing the Koopman operator. We will work on stochastic processes $\{\mathbf{x}_t\}_{t\geq 0}$ with $t \in \mathbb{T}$ where $\mathbb{T} = \mathbb{R}$ for continuous time dynamics and $\mathbb{T} = \mathbb{N}$ for discrete time dynamics.

Definition 2.6. *Let $\tau \in \mathbb{T}$. The linear operator $\mathcal{K}^\tau : C_b(\mathbb{X}) \to C_b(\mathbb{X})$ defined as*

$$\mathcal{K}^\tau \varphi(\mathbf{x}) := \mathbb{E}\left[\varphi(\mathbf{x}_{t+\tau})|\mathbf{x}_t = \mathbf{x}\right] \tag{2.5}$$

is called the (stochastic) Koopman operator.

The Koopman operator is closely related to another important quantity from dynamical system analysis–the Lie derivative–through the following:

Remark 2.1 (From [10])**.** *The family of operators $\mathcal{K}^\tau$ is a one-parameter semigroup of linear operators on $C_b(\mathbb{X})$ generated by the Lie derivative $\mathcal{L}$.*

This gives the following definition of the Lie derivative for stochastic processes.

Definition 2.7. *Let $\mathbf{x}_t$ denote the state at time t of a stochastic process on a probability space $(\Omega, \mathcal{F}, \pi)$ which evolves in a subset $\mathbb{X}$ of a Banach space over $\mathbb{T} = \mathbb{R}$ or $\mathbb{T} = \mathbb{N}$. We write $\mathbb{E}\left[\varphi(\mathbf{x}_s)|\mathbf{x}_t = \mathbf{x}\right]$ for the expected value of $\varphi(\mathbf{x}_s)$ at time $s \geq t$ given that $\mathbf{x}_t = \mathbf{x}$, with the understanding that $\mathbb{E}\left[\varphi(\mathbf{x}_s)|\mathbf{x}_t = \mathbf{x}\right] = \varphi(\mathbf{x}_s)$ for deterministic dynamics. We define the Lie derivative of φ the linear operator $\mathcal{L}$ on the space $C_b(\mathbb{X})$ as*

$$\mathcal{L}\varphi(\mathbf{x}) := \mathbb{E}\left[\varphi(\mathbf{x}_{t+1})|\mathbf{x}_t = \mathbf{x}\right] - \varphi(t\mathbf{x}), \tag{2.6}$$

for discrete time, and

$$\mathcal{L}\varphi(\mathbf{x}) := \lim_{\tau \to 0^+} \frac{\mathbb{E}\left[\varphi(\mathbf{x}_{t+\tau})|\mathbf{x}_t = \mathbf{x}\right] - \varphi(\mathbf{x})}{\tau} \tag{2.7}$$

for continuous time, provided the limit exists uniformly on $\mathbb{T} \times \mathbb{X}$. We call $\mathcal{L}\varphi$ the Lie Derivative of φ.

In the context of dynamics governed by an Ordinary Differential Equation (ODE), the Lie derivative simplifies to the following:

Definition 2.8. *Given an ODE $\dot{\mathbf{x}}_t = f(\mathbf{x}_t)$, the Lie derivative of a function $\varphi : \mathbb{R}^d \to \mathbb{R}$ is defined as the derivative alongside trajectories of the dynamical system:*

$$\mathcal{L}\varphi(\mathbf{x}) := f(\mathbf{x}) \cdot \nabla\varphi(\mathbf{x}). \tag{2.8}$$

A classical example of an application to the Lie derivative is the characterization of global and local stability. For $\mathbf{x}_t \in \mathbb{R}^d$ be the solution of the nonlinear ODE $\dot{\mathbf{x}}_t = f(\mathbf{x}_t)$ with $f(0) = 0$. The equilibrium $\mathbf{x} = 0$ is globally stable if there exists a continuously differentiable function $V : \mathbb{R}^d \to \mathbb{R}$ satisfying

$$\begin{aligned} V(\mathbf{x}) &\geq 0, \ \forall \mathbf{x} \in \mathbb{R}^d, \\ \mathcal{L}V(\mathbf{x}) &< 0, \ \forall \mathbf{x} \neq 0, \\ V(\mathbf{x}) &\to \infty, \text{ as } \|\mathbf{x}\| \to \infty. \end{aligned}$$

Such a function is called a Lyapunov function.

We now present a few examples on how to compute Koopman operators and Lie derivatives for specific systems and observables.

Example 2.1 (The stochastic logistic map)**.** *Consider the stochastic logistic map on $[-1, 1]$*

$$X_{n+1} = 2\mu_n X_n^2 - 1 \tag{2.9}$$

for μ_n drawn from the uniform distribution on $[0, 1]$, for each $n \leq 1$. The Koopman operator applied to the observable $\varphi(x) = x^2$ is given by

$$\begin{aligned} \mathcal{K}\varphi &= \mathbb{E}\left[\varphi(X_{n+1})|X_n = x\right] \\ &= \int_0^1 \varphi(2\mu x^2 - 1) d\mu \\ &= \int_0^1 (2\mu x^2 - 1)^2 d\mu \\ &= \int_0^1 4\mu^2 x^4 - 4\mu x^2 + 1 d\mu \\ &= \left[\frac{4}{3} x^4 \mu^3 - 2\mu^2 x^2 + \mu\right]_0^1 \\ &= \frac{4}{3} x^4 - 2x^2 + 1. \end{aligned}$$

Example 2.2 (The Thomas model)**.** *The Thomas model is a 3-dimensional continuous dynamical system governed by the ODE system*

$$\dot{x} = \sin(y) - bx \tag{2.10}$$
$$\dot{y} = \sin(z) - by \tag{2.11}$$
$$\dot{z} = \sin(x) - bz \tag{2.12}$$

for some parameter $b > 0$. The Lie derivative of the observable $\varphi(x, y, z) = e^{x+y+z}$ is given by

$$\begin{aligned}\mathcal{L}\varphi(x,y,z) &= f(x,y,z)\cdot\nabla\varphi(x,y,z)\\ &= (\sin(y)-bx)\,\frac{\partial}{\partial x}e^{x+y+z} + (\sin(z)-by)\,\frac{\partial}{\partial y}e^{x+y+z} + (\sin(x)-bz)\,\frac{\partial}{\partial z}e^{x+y+z}\\ &= e^{x+y+z}\left(\sin(y)-bx+\sin(z)-by+\sin(x)-bz\right).\end{aligned}$$

2.2.2 Approximating the Koopman operator

Here, we present a technique to approximate the Koopman operator from data called Extended Dynamic Mode Decomposition (EDMD). We start by introducing the concept of *observables*. Let $\{\phi_, \ldots, \phi_l\}$ and $\{\psi_1, \ldots, \psi_m\}$ be two finite dictionaries in $C_b(\mathbb{X})$ whose elements are referred to as *observables*. We set

$$\boldsymbol{\phi} = \begin{bmatrix}\phi_1\\ \vdots\\ \phi_l\end{bmatrix}, \qquad \psi = \begin{bmatrix}\psi_1\\ \vdots\\ \psi_m\end{bmatrix}. \tag{2.13}$$

We denote by span $\boldsymbol{\phi}$, and span $\boldsymbol{\psi}$ respectively, the linear spans of elements belonging to $\boldsymbol{\phi}$ and $\boldsymbol{\psi}$. Consider n 'data snapshots' pairs $(\mathbf{x}_i, \mathbf{y}_i)$, where $\mathbf{y}_i$ is the state of the system exactly τ units in the future from $\mathbf{x}_i$ for all $i = 1, \ldots n$. EDMD seeks to generate an approximate Koopman operator $\mathcal{K}^\tau_{mn}$: span $\boldsymbol{\phi} \to$ span $\boldsymbol{\psi}$. It approximates the action of the exact Koopman operator $\mathcal{K}^\tau$ on span $\boldsymbol{\phi}$ using the basis $\boldsymbol{\psi}$. To this end, we build the matrices

$$\boldsymbol{\Psi_n} := \begin{bmatrix}| & & |\\ \psi(\mathbf{x}_1(t)) & \cdots & \psi(\mathbf{x}_n(t))\\ | & & |\end{bmatrix} \in \mathbb{R}^{m\times n}, \tag{2.14}$$

$$\boldsymbol{\Phi^\tau_n} := \begin{bmatrix}| & & |\\ \boldsymbol{\phi}(\mathbf{x}_1(t+\tau)) & \cdots & \boldsymbol{\phi}(\mathbf{x}_n(t+\tau))\\ | & & |\end{bmatrix} \in \mathbb{R}^{l\times n}. \tag{2.15}$$

Then, we solve the following minimization problem:

$$\min_{A\in\mathbb{R}^{l\times m}} \|\boldsymbol{\Phi^\tau_n} - A\boldsymbol{\Psi_n}\|_F^2, \tag{2.16}$$

whose solution is given by

$$K^\tau_{mn} = \boldsymbol{\Phi^\tau_n}\boldsymbol{\Psi_n}^\dagger = \boldsymbol{\Phi^\tau_n}\boldsymbol{\Psi_n}^T(\boldsymbol{\Psi_n}\boldsymbol{\Psi_n}^T)^\dagger, \tag{2.17}$$

where $A^\dagger$ denotes the Moore-Penrose pseudoinverse. Then one can reproduce the action of the Koopman operator on some $\varphi = \mathbf{c}\cdot\boldsymbol{\phi}$ with $\mathbf{c} \in \mathbb{R}^l$ by

$$\mathcal{K}^\tau_{mn}\varphi := \mathbf{c}\cdot K^\tau_{mn}\boldsymbol{\psi}. \tag{2.18}$$

Remark 2.2. *The EDMD method using two dictionaries is an extremely versatile technique. For example, picking* $\boldsymbol{\phi} = \psi = (f_i(\mathbf{x}))_{i=1}^d$ *where* $f_i(\mathbf{x}) = \mathbf{x}_i$ *simplifies to the classical Dynamic Mode Decomposition (DMD) method [37]. Another possibility is to pick* $\psi = \{f_i(\mathbf{x})\}_{i=1}^d$ *as before and* $\boldsymbol{\phi}$ *any choice of dictionary. With this choice, EDMD becomes model identification, a technique related to the Sparse Identification of Nonlinear Dynamics (SINDy) method [13] without the sparsity constraint.*

With an approximate Koopman operator, it is now possible to obtain an approximate Lie derivative. We will need the following assumption.

Assumption 2.1. *The dictionaries* $\boldsymbol{\phi}$ *and* $\boldsymbol{\psi}$ *satisfy span* $\boldsymbol{\phi} \subseteq \mathcal{D}(\mathcal{L})$ *and span* $\boldsymbol{\phi} \subseteq$ *span* $\boldsymbol{\psi}$.

The goal of this assumption is to enforce that elements in span $\boldsymbol{\phi}$ have well-defined Lie derivatives and the second part of the assumption guarantees that there exists an $l \times m$ matrix Θ_m such that $\boldsymbol{\phi} = \Theta_m \boldsymbol{\psi}$. One can then build an approximate Lie derivative operator $\mathcal{L}_{mn}^\tau : \text{span } \boldsymbol{\phi} \to \text{span } \boldsymbol{\psi}$, where, for every $\varphi = \mathbf{c} \cdot \boldsymbol{\phi}$,

$$\mathcal{L}_{mn}^\tau \varphi = \mathbf{c} \cdot L_{mn}^\tau \psi \quad \text{where} \quad L_{mn}^\tau := \frac{K_{mn}^\tau - \Theta_m}{\tau} \tag{2.19}$$

We are now in a position to present the main known results on the convergence of EDMD to the Koopman operator in the limit of infinite data and infinite dictionary. While our main results come from [10], many similar results are present in the literature such as [30, 39, 40].

Theorem 2.1 (From [10])**.** *Suppose the assumptions stated above hold. Then*

$$\lim_{m\to\infty} \lim_{n\to\infty} \|\mathcal{K}_{mn}^\tau[\varphi] - \mathcal{K}^\tau[\varphi]\|_{L_\rho^2} = 0, \tag{2.20}$$

In particular, $\mathcal{K}_{mn}^\tau[\varphi](\mathbf{x}) \to \mathcal{K}[\varphi](\mathbf{x})$ *almost everywhere on* $\mathbb{X}$.

It should be noted that one could also look at the limit of infinite sampling rate ($\tau \to 0$) and show that the convergence still holds as long as that limit is taken after the limit of infinite data. It is imperative in this result that $n \to \infty$ (the limit of data) be the innermost limit, or else the analysis done in [10] would not hold. The case of m and n growing together has been handled in [25].

Remark 2.3 (On the infinite data limit)**.** *Since we will derive convergence results for the approximation of the Koopman operator, we will work on the limit of infinite data, i.e.,* $n \to \infty$. *It will be useful for us to define*

$$K_m^\tau := \lim_{n\to\infty} K_{mn}^\tau \quad \textit{and} \quad \mathcal{K}_m^\tau := \lim_{n\to\infty} \mathcal{K}_{mn}^\tau, \tag{2.21}$$

as is done in [10].

2.3 Least-squares theory

We now present the theory of function approximation using the least-squares method. We will introduce the method and very important results on the approximation error and convergence of the method.

2.3.1 Definition of the least-squares problem

We consider the problem of approximating scalar-valued functions of finitely many variables. Given a function $f : \mathbb{X} \to \mathbb{C}$, $f \in L^2_\rho(\mathbb{X}) \cap C(\mathbb{X})$, where $C(\mathbb{X})$ is the space of continuous functions on $\mathbb{X}$. We study the problem of approximating f in an arbitrary finite dimensional subspace $\mathcal{P} \subset L^2_\rho(\mathbb{X}) \cap C(\mathbb{X})$ of dimension $\dim(\mathcal{P}) = s$. Let $\{\mathbf{y}_i\}_{i=1}^n, \mathbf{y}_i \in \mathbb{X}$ for $i = 1, \dots, n$ be n *sample points*. We will try to approximate f in the subspace $\mathcal{P}$ from its possibly noisy sample values

$$f(\mathbf{y}_i) + n_i, \qquad i = 1, \dots, n, \tag{2.22}$$

where each n_i represent a measurement error. In this section, we will present various results to obtain approximations that are quasi-optimal in $\mathcal{P}$. That is the corresponding error is, up to noise, proportional to the best approximation error of f in $\mathcal{P}$. We will also quantify the number of sample points required to achieve it. We note that we will focus on the case where the n_i's are bounded.

Assuming $n \geq \dim(\mathcal{P})$, we define the *least-squares approximation* $\hat{f}$ *of* f as the solution of the following minimization problem:

$$\hat{f} = \underset{p \in \mathcal{P}}{\arg\min} \frac{1}{n} \sum_{i=1}^{n} |f(\mathbf{y}_i) - p(\mathbf{y}_i) + n_i|^2. \tag{2.23}$$

Remark 2.4. *While in general, the least-squares approximation $\hat{f}$ is not unique, for the purpose of this work, it will always be the case that $\hat{f}$ is unique as we always oversample the space $\mathcal{P}$ with $n \geq s$, and the samples $\{\mathbf{y}_i\}_{i=1}^n$ are distinct with probability 1 in our context.*

Let $\{\psi_i\}_{i=1}^s$ be an orthonormal basis of $\mathcal{P}$. We define the normalized *measurement matrix* and the (noisy) normalized *measurement vector* as

$$A = \left(\frac{1}{\sqrt{n}} \psi_j(\mathbf{y}_i) \right)_{i,j=1}^{n,s} \in \mathbb{C}^{n \times s}, \qquad \mathbf{b} = \frac{1}{\sqrt{n}} \left(f(\mathbf{y}_i) + n_i \right)_{i=1}^n \in \mathbb{C}^n. \tag{2.24}$$

We can provide an equivalent formulation of the least-squares problem (2.23) in terms of the measurement matrix A and the measurement vector $\mathbf{b}$. This fact is very useful as it allows us to reduce our analysis to a classical linear algebra least-squares problem. Moreover, many tools that our analysis will rely on are based on this alternate formulation.

Lemma 2.1. *Suppose that $\hat{f} = \sum_{i=1}^s \hat{c}_i \psi_i$, we can then rewrite the minimization problem* (2.23) *as*

$$\hat{\mathbf{c}} \in \underset{\mathbf{c} \in \mathbb{C}^s}{\arg\min} \|\mathbf{b} - A\mathbf{c}\|_2^2. \tag{2.25}$$

Proof. We have that

$$\begin{aligned}\|\mathbf{b}-A\mathbf{c}\|_2^2 &= \sum_{i=1}^{n}\left|\frac{1}{\sqrt{n}}\left(f(\mathbf{y_i})+n_i-\sum_{j=1}^{s}c_j\psi_j(\mathbf{y}_i)\right)\right|^2\\ &= \frac{1}{n}\sum_{i=1}^{n}\left|f(\mathbf{y}_i)+n_i-\sum_{j=1}^{s}c_j\psi_j(\mathbf{y}_i)\right|^2\\ &= \frac{1}{n}\sum_{i=1}^{n}|f(\mathbf{y}_i)+n_i-p(\mathbf{y}_i)|,\end{aligned}$$

where $p=\sum_{j=1}^{s}c_j\psi_j$. □

2.3.2 The discrete stability constant and Christoffel function

Our first goal will be to analyze the error of the least-squares approximation (2.23). We introduce the following useful tool:

Definition 2.9 (Discrete semi-inner product and seminorm)**.** *Given $f,g\in L^2_\rho(\mathbb{X})\cap C(\mathbb{X})$, we define the discrete semi-inner product of f and g as*

$$\langle f,g\rangle_{\text{disc}}=\frac{1}{n}\sum_{i=1}^{n}f(\mathbf{y}_i)\overline{g(\mathbf{y}_i)}, \tag{2.26}$$

and the corresponding discrete semi-norm as

$$\|f\|_{\text{disc}}^2=\langle f,f\rangle_{\text{disc}}=\frac{1}{n}\sum_{i=1}^{n}|f(\mathbf{y}_i)|^2. \tag{2.27}$$

To study the behaviour of the semi-inner product and seminorm, we introduce the *discrete stability constant*.

Definition 2.10 (Discrete stability constant)**.** *The discrete stability constant is defined as*

$$\alpha=\alpha\left(\mathcal{P},\{\mathbf{y}_i\}_{i=1}^n\right):=\inf\left\{\|p\|_{\text{disc}}:p\in\mathcal{P},\|p\|_{L^2_\rho(\mathbb{X})}=1\right\}. \tag{2.28}$$

This constant gives a one-sided estimate between the discrete seminorm and the $L^2_\rho-$norm. Specifically,

$$\|p\|_{\text{disc}}\geq\alpha\|p\|_{L^2_\rho(\mathbb{X})},\qquad\forall p\in\mathcal{P}. \tag{2.29}$$

The discrete semi-inner product can be reformulated using the measurement matrix A. This formulation will be generally more convenient to work with. Letting $f=\sum_{i=1}^{s}c_i\psi_i$ and $g=\sum_{i=1}^{s}d_i\psi_i$, we have that

$$\begin{aligned}\langle f,g\rangle_{\text{disc}} &= \frac{1}{n}\sum_{i=1}^{n}f(\mathbf{y}_i)\overline{g(\mathbf{y}_i)}\\ &= \sum_{i=1}^{n}\left(\frac{1}{\sqrt{n}}\sum_{j=1}^{n}c_j\psi_j(\mathbf{y}_i)\right)\overline{\left(\frac{1}{\sqrt{n}}\sum_{j=1}^{s}d_j\psi_j(\mathbf{y}_i)\right)}\\ &= \mathbf{d}^*A^*A\mathbf{c}.\end{aligned}$$

Similarly,

$$\|f\|_{\text{disc}}^2 = \langle f, f\rangle_{\text{disc}} = \mathbf{c}^* A^* A\mathbf{c} = \langle A\mathbf{c}, A\mathbf{c}\rangle = \|A\mathbf{c}\|_2^2,$$

so $\|f\|_{\text{disc}} = \|A\mathbf{c}\|_2$ Letting $p = \sum_{i=1}^s c_i\psi_i$, we can furthermore express the discrete stability constant as

$$\alpha = \inf\left\{\|A\mathbf{c}\|_2 : \mathbf{c} \in \mathbb{C}^s, \|\mathbf{c}\|_2 = 1\right\}. \tag{2.30}$$

Hence, thanks to the min-max theorem [33, Theorem 1.3.2], we obtain the following:

$$\begin{aligned}\alpha &= \inf\left\{\|A\mathbf{c}\|_2 : \mathbf{c} \in \mathbb{C}^s, \|\mathbf{c}\|_2 = 1\right\} \\ &= \sup_{\dim(V)=s} \inf\left\{\|A\mathbf{c}\|_2 : \mathbf{c} \in V, \|\mathbf{c}\|_2 = 1\right\},\end{aligned}$$

as there is only one s-dimensional subspace of $\mathbb{C}^s$ with dimension s, that is necessarily $V = \mathbb{C}^s$. Therefore,

$$\begin{aligned}\alpha &= \sup_{\dim(V)=s} \inf_{\mathbf{c}\in V, \|\mathbf{c}\|_2=1} \|A\mathbf{c}\|_2 \\ &= \sup_{\dim(V)=s} \inf_{\mathbf{c}\in V, \|\mathbf{c}\|_2=1} \mathbf{c}^* A^* A\mathbf{c} = \lambda_{\min}(A^*A) = \sigma_{\min}(A).\end{aligned}$$

The first insight that we can gather from the discrete-semi inner product is in regard to the Monte Carlo approximation of the integral

$$\int_{\mathbb{X}} |p(\mathbf{y})| d\rho(\mathbf{y}).$$

If the $\mathbf{y}_i$'s are sampled i.i.d, then due to the strong law of large numbers, we have almost surely that

$$\|p\|_{\text{disc}}^2 = \frac{1}{n}\sum_{i=1}^n |p(\mathbf{y}_i)|^2 \to \int_{\mathbb{X}} |p(\mathbf{y})|^2 d\rho(\mathbf{y}) = \|p\|_{L^2_\rho(\mathbb{X})}^2, \quad \text{as} \quad n \to \infty.$$

Since $\mathcal{P}$ is finite dimensional, this in turn almost surely implies that $\alpha \to 1$ as $n \to \infty$ since $\|p\|_{\text{disc}} \to \|p\|_{L^2_\rho(\mathbb{X})}$ and, therefore

$$\alpha = \inf_{\|p\|_{L^2_\rho(\mathbb{X})}=1} \|p\|_{\text{disc}} \to \inf_{\|p\|_{L^2_\rho(\mathbb{X})}=1} \|p\|_{L^2_\rho(\mathbb{X})} = 1.$$

Moreover, as seen in [3, Chapter 5], the sequence of least-square approximations $\hat{f}$ for $n = 1, 2, \ldots$ converges almost surely to the best approximation of f in $\mathcal{P}$

$$f^* := \arg\min_{p\in\mathcal{P}} \|f - p\|_{L^2_\rho(\mathbb{X})}.$$

We conclude this subsection by providing a proof of the following useful result.

Proposition 2.3. *Let $f \in L^2_\rho(\mathbb{X}) \cap C(\mathbb{X})$, $f = \sum_{i=1}^s c_i\psi_i$ and $\hat{f}$ be the least-squares approximation* (2.23)*. Then, for any $p \in \mathcal{P}$, we have that*

$$\left\langle \hat{f}, p\right\rangle_{\text{disc}} = \langle f, p\rangle_{\text{disc}} + \frac{1}{n}\sum_{i=1}^n n_i g(\mathbf{y}_i).$$

Proof. Let $p = \sum_{i=1}^{s} d_i \psi_i$. Then,

$$\begin{aligned}\left\langle \hat{f}, p \right\rangle_{\text{disc}} &= \hat{\mathbf{c}}^* A^* A \mathbf{d} \\ &= \mathbf{b}^* A \mathbf{d},\end{aligned}$$

as $A^* A \hat{\mathbf{c}} = A^* \mathbf{b}$. However, $\mathbf{b} = \left(\frac{1}{\sqrt{n}}(f(\mathbf{y}_i + n_i))\right)_{i=1}^{n}$ and $A\mathbf{d} = \frac{1}{\sqrt{n}} \left(\sum_{j=1}^{s} d_i (\psi_j(\mathbf{y}_i))_{i=1}^{n}\right)_{i=1}^{n} = \frac{1}{\sqrt{n}} \left(g(\mathbf{y}_i)\right)_{i=1}^{n}$. Hence, we have that

$$\begin{aligned}\left\langle \hat{f}, p \right\rangle_{\text{disc}} &= \mathbf{b}^* A \mathbf{d} \\ &= \frac{1}{n} \sum_{i=1}^{n} f(\mathbf{y}_i) \overline{g(\mathbf{y}_i)} + \frac{1}{n} \sum_{i=1}^{n} n_i g(\mathbf{y}_i) \\ &= \langle f, g \rangle_{\text{disc}} + \frac{1}{n} \sum_{i=1}^{n} n_i g(\mathbf{y}_i).\end{aligned}$$

This concludes the proof. □

Our next goal is to bound α in the case where the $\mathbf{y}_i$'s are drawn independently from the orthogonality measure ρ. A key tool for this analysis is the *Christoffel function*.

Definition 2.11 (Christoffel function)**.** *Let $\mathcal{P}$ be an s-dimensional subspace of $L^2_\rho(\mathbb{X}) \cap C(\mathbb{X})$. The Christoffel function is defined as $\mathbf{y} \mapsto 1/K(\mathbf{y})$, where*

$$K(\mathbf{y}) = K(\mathcal{P})(\mathbf{y}) := \sum_{i=1}^{s} |\psi_i(\mathbf{y})|^2, \quad \forall \mathbf{y} \in \mathbb{X}, \tag{2.31}$$

and $\{\psi_i\}_{i=1}^{s}$ is any orthonormal basis of $\mathcal{P}$. In this case, we also write

$$\kappa = \kappa(\mathcal{P}) := \|K\|_{L^\infty(\mathbb{X})} := \sup_{\mathbf{y} \in \mathbb{X}} |K(\mathbf{y})| \tag{2.32}$$

for its L^∞-norm.

We note that even though the definition requires a choice of basis, the Christoffel function is independent of the choice of basis. This is due to the fact that the function $K(\mathbf{y})$ admits a different representation.

Lemma 2.2.

$$K(\mathbf{y}) = \sup\{|p(\mathbf{y})|^2 : p \in \mathcal{P}, \|p\|_{L^2_\rho} = 1\}, \qquad \forall \mathbf{y} \in \mathbb{X}. \tag{2.33}$$

Proof. Since $\{\psi_j\}_{j=1}^{s}$ forms an orthonormal basis of $\mathcal{P}$, from Parseval's identity, we have that

$$\begin{aligned}\sup_{\|p\|_{L^2_\rho(\mathbb{X})}=1} |p(\mathbf{y})|^2 &= \sup_{\|c\|_2=1} \left| \sum_{j=1}^{s} c_j \psi_j(\mathbf{y}) \right|^2 \\ &\leq \sup_{\|c\|_2=1} \|c\|_2^2 \left(\sum_{j=1}^{s} |\psi_j(\mathbf{y})|^2 \right) \\ &= K(\mathbf{y}),\end{aligned}$$

where the inequality follows from the Cauchy-Schwarz inequality. Now pick $c_j = \psi_j(\mathbf{y})/\sqrt{K(\mathbf{y})}$, to get

$$\|\mathbf{c}\|_2^2 = \sum_{i=1}^{s} \frac{\psi_j(\mathbf{y})^2}{K(\mathbf{y})} = \frac{K(\mathbf{y})}{K(\mathbf{y})} = 1,$$

and, moreover,

$$\begin{aligned}\sup_{\|p\|_{L^2_\rho(\mathbb{X})}=1} |p(\mathbf{y})|^2 &= \left|\sum_{j=1}^{s} c_j \psi_j(\mathbf{y})\right|^2 \\ &= \left|\sum_{j=1}^{s} \frac{\psi_j(\mathbf{y})}{\sqrt{K(\mathbf{y})}} \psi_j(\mathbf{y})\right|^2 \\ &= \frac{1}{K(\mathbf{y})}\left(\sum_{j=1}^{s} |\psi_j(\mathbf{y})|^2\right)^2 = \frac{K(\mathbf{y})^2}{K(\mathbf{y})} = K(\mathbf{y}).\end{aligned}$$

This concludes the proof. □

Note that the Christoffel function is undefined for some $\mathbf{y}$ if $K(\mathbf{y}) = 0$, or equivalently if $p(\mathbf{y}) = 0$ for all $p \in \mathcal{P}$. This is a very mild assumption to satisfy in practice, as we would simply require that $\mathcal{P}$ admits an orthonormal basis such that for a given $\mathbf{y}$, there exists at least one index i such that $\psi_i(\mathbf{y}) \neq 0$. We would like to use this Christoffel function to obtain an estimate of α.

Theorem 2.2 (Stability constant lower bound in probability (from [3, Chapter 5]))**.** *Let $0 < \delta, \epsilon < 1$, $\mathcal{P} \subset L^2_\rho(\mathbb{X}) \cap C(\mathbb{X})$ be a finite dimensional subspace with $\dim(\mathcal{P}) = s$ and suppose that $\mathbf{y}_1, \ldots, \mathbf{y}_n$ are drawn independantly from ρ. If*

$$n \geq c_\delta \cdot \kappa \cdot \log(s/\epsilon), \qquad c_\delta = ((1-\delta)\log(1-\delta) + \delta)^{-1}, \tag{2.34}$$

where $\kappa = \kappa(\mathcal{P})$ is as in (2.32), then, with probability at least $1 - \epsilon$, the discrete stability constant $\alpha = \alpha(\mathcal{P}, \{\mathbf{y}_i\}_{i=1}^n)$ defined in (2.28) satisfies

$$\alpha > \sqrt{1-\delta}. \tag{2.35}$$

Remark 2.5. *A proof is available in [3, Theorem 5.7]. It relies on defining appropriate independent, self-adjoint random matrices and applying the matrix Chernoff bound [36, Theorem 1.1]. An important observation of this proof is that it is necessary that $\mathbb{E}[A^*A] = I_d$.*

2.3.3 Preconditioning and Christofell sampling

In this subsection, we focus on the case of polynomial approximation. We will study what happens when the samples $\mathbf{y}_i$ are sampled from a probability measure μ_i that is different from ρ, and without imposing that $\mu_i = \mu_j$ for $i \neq j$, more specifically, we will study what happens to κ and present a strategy called *Christofell sampling*.

Let ρ be a measure on $\mathbb{X}$ that satisfies

$$\frac{1}{n}\sum_{i=1}^{n} d\mu_i(\mathbf{y}) = \frac{d\rho(\mathbf{y})}{w(\mathbf{y})}, \quad \mathbf{y} \in \mathbb{X}, \tag{2.36}$$

for some weight function $w : \mathbb{X} \to \mathbb{R}$ strictly positive and finite almost everywhere. Notice since the μ_i's are probability measures, $\int_{\mathbb{X}} 1/w(\mathbf{y}) d\rho(\mathbf{y}) = \sum_{i=1}^{n} \int_{\mathbb{X}} d\mu_i(\mathbf{x}) = 1$. We then assume that each $\mathbf{y}_i$ is drawn independently from μ_i. In order to apply the same analysis as the previous subsection, more specifically to apply an argument analogous to that of Theorem 2.2, we need to guarantee that $\mathbb{E}[A^*A] = I_d$ (See Remark 2.5). For this to hold, we need to solve instead a *weighted least-squares problem*:

$$\hat{f} \in \arg\min_{p \in \mathcal{P}} \frac{1}{n}\sum_{i=1}^{n} w(\mathbf{y}_i) \left| f(\mathbf{y}_i) - p(\mathbf{y}_i) + n_i \right|^2. \tag{2.37}$$

We can define similarly to the unweighted problem the measurement matrix, and measurement vector for this problem:

$$A_w = \left(\frac{1}{\sqrt{n}} \sqrt{w(\mathbf{y}_i)} \psi_j(\mathbf{y}_i) \right)_{i,j=1}^{n,s} \in \mathbb{C}^{n \times s}, \qquad \mathbf{b}_w = \frac{1}{\sqrt{n}} \sqrt{w(\mathbf{y}_i)} \left(f(\mathbf{y}_i) + n_i \right)_{i=1}^{n} \in \mathbb{C}^n. \tag{2.38}$$

Proposition 2.4. *Let $\hat{f} = \sum_{i=1}^{s} \hat{c}_i \psi_i$ be the solution of the weighted least-squares problem (2.37) with $\eta_i = 0$, $i = 1, \ldots, n$. Then,*

$$\hat{\mathbf{c}} = A_w^{\dagger} \mathbf{b}_w = (WA)^{\dagger} W \mathbf{b}, \tag{2.39}$$

where $W = diag(\sqrt{w(\mathbf{y}_i)})_{i=1}^{n} \in \mathbb{C}^{n \times n}$, and $A^{\dagger} = (A^T A)^{-1} A$ is the Moore-Penrose pseudoinverse of A.

Proof. We recall the identity between the functional least squares problem and the algebraic least squares problem

$$\frac{1}{n}\sum_{i=1}^{n} |f(\mathbf{y}_i) - p(\mathbf{y}_i)|^2 = \|A\mathbf{c} - \mathbf{b}\|_2^2,$$

where $A, \mathbf{b}$ are defined in (2.24). Now notice that

$$\begin{aligned} \frac{1}{n}\sum_{i=1}^{n} w(\mathbf{y}_i) |f(\mathbf{y}_i) - p(\mathbf{y}_i)|^2 &= \frac{1}{n}\sum_{i=1}^{n} |\sqrt{w(\mathbf{y}_i)} f(\mathbf{y}_i) - \sqrt{w(\mathbf{y}_i)} p(\mathbf{y}_i)|^2 \\ &= \|WA\mathbf{c} - W\mathbf{b}\|_2^2 \\ &= \|A_w\mathbf{c} - \mathbf{b}_w\|_2^2. \end{aligned}$$

A solution of the weighted least-squares problem is thus given by the minimizer of the algebraic least-squares problem

$$\hat{\mathbf{c}} \in \arg\min_{\mathbf{c} \in \mathbb{C}^s} \|A_w\mathbf{c} - \mathbf{b}_w\|_2^2.$$

A minimizer of this problem must satisfy $\nabla\|A_w\hat{\mathbf{c}} - \mathbf{b}_w\|_2^2 = 0$. Moreover,

$$\begin{aligned}
&\nabla\|A_w\hat{\mathbf{c}} - \mathbf{b}_w\|_2^2 = 0 \\
&\Leftrightarrow 2A_w^*(A_w\hat{\mathbf{c}} - \mathbf{b}_w) = 0 \\
&\Leftrightarrow A_w^*A_w\hat{\mathbf{c}} = A_w^*\mathbf{b}_w \\
&\Leftrightarrow \hat{\mathbf{c}} = \widehat{(A_w^*A_w)}A_w^*\mathbf{b}_w = (A_w)^\dagger\mathbf{b}_w.
\end{aligned}$$

Finally, we notice that $A_w = WA$, and $\mathbf{b}_w = W\mathbf{b}$ to get that

$$\hat{\mathbf{c}} = (WA)^\dagger W\mathbf{b}.$$

This concludes the proof. □

We can also define weighted discrete semi-inner product and semi-norm as

$$\langle f, g\rangle_{\mathrm{disc},w} = \frac{1}{n}\sum_{i=1}^n w(\mathbf{y}_i)f(\mathbf{y}_i)\overline{g(\mathbf{y}_i)}, \qquad \|f\|_{\mathrm{disc},w}^2 = \langle f, f\rangle_{\mathrm{disc}} = \frac{1}{n}\sum_{i=1}^n w(\mathbf{y}_i)|f(\mathbf{y}_i)|^2. \quad (2.40)$$

The discrete stability constant and the L^∞-norm of the Christoffel function are now given by

$$\alpha_w = \alpha_w\left(\mathcal{P}, \{\mathbf{y}_i\}_{i=1}^n\right) := \inf\left\{\|p\|_{\mathrm{disc},w} : p \in \mathcal{P}, \|p\|_{L^2_\rho(\mathbb{X})} = 1\right\}, \quad (2.41)$$

$$\kappa_w = \kappa_w(\mathcal{P}) := \|wK\|_{L^\infty(\mathbb{X})} = \sup_{\mathbf{y}\in\mathbb{X}} w(\mathbf{y})K(\mathbf{y}). \quad (2.42)$$

Note that this setup encompasses the unweighted case by setting $\mu_1 = \cdots = \mu_n = \rho$ and $w(\mathbf{y}) = 1$. Moreover, Theorem 2.2 can be extended to the weighted case as we have guaranteed that $\mathbb{E}[A^*A] = I_d$, and the $\mathbf{y}_i$'s are still sampled independently.

Let us now focus on a particular case. Let $\{\psi_\nu\}_{\nu\in\mathbb{N}_0^d}$ be the d-variate Legendre polynomial basis, $S \subset \mathbb{N}_0^d$, $\mathcal{P}_S = \mathrm{span}\{\psi_n\}_{\nu\in S}$ and define $\mu_1 = \ldots = \mu_n = \mu$ to be the tensor Chebyshev measure

$$d\mu(\mathbf{y}) = \prod_{i=1}^d \frac{1}{\pi\sqrt{1-y_i^2}}dy_i. \quad (2.43)$$

This can be motivated from the theory of polynomial approximation in low dimensions, where it is well known that uniformly distributed samples are a poor choice, whereas the Chebyshev measure concentrates the samples near the endpoints, avoiding issues such as Runge's phenomenon (see [6, Example 4.2]). Notice that the Legendre polynomials are orthogonal with respect to the uniform measure $d\rho(\mathbf{y}) = 2^{-d}d\mathbf{y}$, so the weight function is

$$w(\mathbf{y}) = \prod_{i=1}^d \frac{\pi}{2}\sqrt{1-y_i^2}, \quad \forall \mathbf{y} \in \mathbb{X}. \quad (2.44)$$

Another motivation for that choice is to notice that, in this setting,

$$\kappa_w(\mathcal{P}_S) = \sup_{\mathbf{y}\in\mathbb{X}}\sum_{\nu\in S}|\sqrt{w(\mathbf{y})}\psi_\nu(\mathbf{y})|^2 = \sup_{\mathbf{y}\in\mathbb{X}}\sum_{\nu\in S}|\boldsymbol{\phi}_\nu(\mathbf{y})|^2 = \kappa(\mathcal{P}_{S,w}), \quad (2.45)$$

where $\mathcal{P}_{S,w} = \text{span}\{\boldsymbol{\phi}_\nu\}_{\nu\in S}$ with

$$\boldsymbol{\phi}_\nu(\mathbf{y}) = \prod_{i=1}^{d} \sqrt{\frac{\pi}{2}} \left(1 - y_i^2\right)^{1/4} \psi_\nu(\mathbf{y}), \quad \forall \mathbf{y} \in \mathbb{X}. \tag{2.46}$$

This is the *preconditionned* basis. Its elements are orthonormal with respect to μ, hence this is equivalent to the unweighted case with the preconditioned basis. For this reason, this procedure is called *preconditionning* the Legendre polynomials. The reason why this gives better performance, is that one can see that $\|\boldsymbol{\phi}_\nu\|_{L^\infty(\mathbb{X})} < (\pi/2)^{d/2}(2/\sqrt{\pi})^{\|\nu\|_0} \quad \forall \nu \in \mathbb{N}_0^d$. These functions are bounded uniformly in $\boldsymbol{\nu}$ unlike the d-variate Legendre polynomials, so in general $\kappa(\mathcal{P}_{S,w})$ scales better than $\kappa(\mathcal{P}_S)$ as the dimension increases.

With this example in mind, one can ask the following question: *How can we reproduce the preconditionning effect to for an arbitrary $\mathcal{P}$ and ρ?*. The technique we are going to introduce now is called *Christofell sampling*. We can derive the optimal weight function w and sampling measures $\mu_1 = \cdots = \mu_n$ for a given measure ρ by utilizing the bound $\kappa_w \geq w(\mathbf{y})K(\mathbf{y})$ and integrating with respect to $\frac{1}{w(\mathbf{y})}d\rho(\mathbf{y})$:

$$\int_{\mathbb{X}} w(\mathbf{y})K(\mathbf{y})\frac{1}{w(\mathbf{y})}d\rho(\mathbf{y}) \leq \int_{\mathbb{X}} \kappa_w \frac{1}{w(\mathbf{y})}d\rho(\mathbf{y}). \tag{2.47}$$

Since $\int_{\mathbb{X}} K(\mathbf{y})d\rho(\mathbf{y}) = s$ and $\int_{\mathbb{X}} 1/w(\mathbf{y})d\rho(\mathbf{y})$, we deduce that $\kappa_w \geq s$. On the other hand, we minimize κ_w by taking $w(\mathbf{y})$ proportional to $1/K(\mathbf{y})$. Given our assumption that $K(\mathbf{y}) > 0$ for all $\mathbf{y} \in \mathbb{X}$, and $\int_{\mathbb{X}} 1/w(\mathbf{y})d\rho(\mathbf{y}) = 1$, this leads to the following choice for $w(\mathbf{y})$:

$$w(\mathbf{y}) = \left(\frac{K(\mathbf{y})}{s}\right)^{-1}, \quad \forall \mathbf{y} \in \mathbb{X}, \tag{2.48}$$

yielding the optimal value $\kappa_w = s$. Now our sampling measures need to satisfy

$$\frac{1}{n}\sum_{i=1}^{n} d\mu_i(\mathbf{y}) = \frac{1}{s}K(\mathbf{y})d\rho(\mathbf{y}). \tag{2.49}$$

There are multiple ways to achieve such a result, arguably the simplest and most straightforward solution is to set

$$d\mu_1(\mathbf{y}) = \cdots = d\mu_n(\mathbf{y}) = \frac{1}{s}K(\mathbf{y})d\rho(\mathbf{y}). \tag{2.50}$$

It should be noted that in the case where the assumption that $K(\mathbf{y}) > 0$ does not hold, one could replace the w we chose by any of the following [1]:

$$w(\mathbf{y}) = \left(\theta + (1-\theta)\frac{K(\mathbf{y})}{s}\right)^{-1}, \quad \theta \in (0,1). \tag{2.51}$$

The cost to pay is a slight increase in κ_w, as now we can only guarantee $\kappa_w \leq (1-\theta)^{-1}s$. Additionaly, in the case where $K(\mathbf{y})$ is small, our first choice of w would become very large,

and this could lead to numerical instability due to the noise. So this choice of w could also be more suitable in a high noise, low $K(\mathbf{y})$ setting.

For this choice of w, one would need the measures μ_i to satify

$$\frac{1}{n}\sum_{i=1}^{n} d\mu_i(\mathbf{y}) = \left(\theta + \frac{1}{s}(1-\theta)K(\mathbf{y})\right) d\rho(\mathbf{y}). \tag{2.52}$$

Once again, this can be solved in a straighforward manner by setting

$$d\mu_1(\mathbf{y}) = \ldots = d\mu_n(\mathbf{y}) = \left(\theta + \frac{1}{s}(1-\theta)K(\mathbf{y})\right) d\rho(\mathbf{y}). \tag{2.53}$$

2.3.4 Convergence rates

In this subsection, we present recent work done in [2] on the convergence rates of the least-squares approximation. First, we need to introduce an ubiquitous object in the study of least squares, the Bernstein polyellipse.

Definition 2.12 (Bernstein polyellipse)**.** *The Bernestein ellipse of parameter $\rho > 1$ is defined by*

$$\mathcal{E}_\rho = \left\{(z + z^{-1})/2 : z \in \mathbb{C}, 1 \leq |z| \leq \rho\right\}. \tag{2.54}$$

Given a parameter $\boldsymbol{\rho} = (\rho_1, \rho_2, \ldots, \rho_N)$, we define the Bernstein polyellipse of parameter $\boldsymbol{\rho}$ as

$$\mathcal{E}_{\boldsymbol{\rho}} = \mathcal{E}_{\rho_1} \times \mathcal{E}_{\rho_2} \times \ldots \times \mathcal{E}_{\rho_N} \subset \mathbb{C}^N. \tag{2.55}$$

We will now introduce the main result of this subsection about the convergence rate of the least-squares approximation in the case where the sample complexity scales *log-linearly* with the dimension of the underlying problem.

Theorem 2.3 ([15, Theorem B.1])**.** *Let $0 < \epsilon < 1$, $p \in (0,1)$, ρ be either the uniform or Chebyshev measure on $\mathbb{X} = [-1,1]^d, n \geq 3$ and $\mathbf{y}_1, \ldots \mathbf{y}_n$ sampled i.i.d. from ρ. Then there exists a set $S \subseteq \mathbb{N}_0^d$ of cardinality $|S| \leq \lceil n/\log(n/\epsilon)\rceil$ such that the following holds with probability at least $1-\epsilon$ for each fixed $f : \mathbb{X} \to \mathbb{C}$ that admits a holomorphic extension to a Bernstein polyellipse $\mathcal{E}_{\boldsymbol{\rho}} \subset \mathbb{C}^N$. For any $\mathbf{e} \in \mathbb{C}^d$, the least squares approximation $\hat{f}$ to the problem* (2.23) *is unique and satisfies*

$$\|f - \hat{f}\|_{L^2_\rho(\mathbb{X})} \leq C(\boldsymbol{\rho}, p)\left(\frac{n}{\log(n/\epsilon)}\right)^{\frac{1}{2}-\frac{1}{p}} + 2\|\mathbf{e}\|_\infty, \tag{2.56}$$

where $C(\boldsymbol{\rho}, p)$ is a constant that depends on $\boldsymbol{\rho}$ and p only.

Remark 2.6. *This theorem only applies to the case where the measure ρ is either the uniform or Chebyshev measure. For a recent extension of this result to the more general case of Jacobi measures, see [4].*

This theorem is an important result, as it shows that using the Chebyshev or uniform measure, it is possible to achieve a convergence rate faster than the well observed Monte Carlo rate $O(n^{-\frac{1}{2}})$. The value of p in the bound comes from the theory of functions of infinitely many variables, where we require the function to be approximated is $(\mathbf{b}, \epsilon)$-holomorphic in a Bernstein polyellipse. In this setting, this theorem requires $\mathbf{b} \in \ell^p(\mathbb{N})$. Because in the finite dimensional case $\mathbf{b}$ is necessarily a finite sequence, this theorem shows that we should expect a faster than algebraic convergence. See [2, Theorem 6.1 and Remark 5.2] for more details.

2.4 Compressed sensing theory

In this section, we introduce the main theory that will be used to understand the approximation of the Koopman operator via compressed sensing techniques. We will consider the Quadratically constrained basis pursuit problem, and the LASSO-type decoders, as well as recovery guarantee results. Our main references for this section are [3, 18–20].

2.4.1 Polynomial approximation through compressed sensing

Compressed sensing is a novel approximation techniques that arose in the mid 2000's by two papers of Candès, Romberg and Tao [14] and Donoho [16]. That theory has seen successful applications in signal processing, machine learning, and, the subject of this section, polynomial approximation. As in the previous section, our goal here is to present the necessary theory that we will use to approximate the Koopman operator. This section is based on multiple books, but mainly [3, 18, 20]. The theory that we outline in this section is the application of compressed sensing to function approximation. As before, we aim to recover a function f expanded in the basis as

$$f = \sum_{\nu \in \mathbb{N}_0^d} c_\nu \psi_\nu, \qquad c_\nu = \langle f, \psi_\nu \rangle_{L^2_\rho(\mathbb{X})}, \forall \nu \in \mathbb{N}_0^d. \tag{2.57}$$

We make an additional assumption for now that f has exactly s nonzero terms in the expansion, i.e.,

$$f = \sum_{\nu \in S} c_\nu \psi_\nu, S \subset \mathbb{N}_0^d, |S| = s. \tag{2.58}$$

The assumption is not always satisfied in practice, but it is a good starting point to understand the theory. In practice, we might not know S, but we will assume to we know a set $\Lambda \subset \mathbb{N}_0^d, S \subset \Lambda$. Λ may be a tensor-product, total degree or hyperbolic cross set.

As in the previous section, we aim to recover f using n sample points $\mathbf{y}_1, \ldots, \mathbf{y}_n$ randomly drawn i.i.d. from a measure ρ. If we assume $|\Lambda| = N$, then in this setting we have that $n < N$ (typically $n \ll N$), and $m \geq s$. We will need once again the measurement vector and matrix as defined in (2.38). Moreover, it will be useful to work in the case where the measurement matrix comes from sampling a *Bounded Orthonormal System* (BOS).

Definition 2.13. *A measurement matrix* $A = \left(\frac{1}{\sqrt{n}}\psi_j(\mathbf{y_i})\right)_{i,j=1}^{n,s}$ *is said to come from a Bounded Orthonormal System (BOS) if there exists a constant* $B > 0$ *such that* $\int_{\mathbb{X}} \psi_j \overline{\psi_i} d\rho = \delta_{j,k}$ *where* $\delta_{j,k}$ *is the Kronecker delta, and* $\|\psi_j\|_{L^\infty(\mathbb{X})} \leq B, \forall j = 1, \ldots, s$.

Before formulating a compressed sensing problem, we will introduce the concept of sparsity.

Definition 2.14 (Sparsity)**.** *Let* $\mathbf{c} = (c_j)_{j=1}^N \in \mathbb{C}^N$*. The support of* $\mathbf{c}$ *is defined as*

$$supp(\mathbf{c}) = \{j \in [N] : c_j \neq 0\}. \tag{2.59}$$

We say $\mathbf{c}$ *is* s*-sparse if*

$$\|\mathbf{c}\|_0 := |supp(\mathbf{c})| \leq s. \tag{2.60}$$

The set of all s*-sparse vectors in* $\mathbb{C}^N$ *is denoted by* Σ_s^N*.*

A concept of paramount importance in compressed sensing is the idea of the *best* s*-term approximation error*. This is the minimum error we can achieve by approximating a vector $\mathbf{c} \in \mathbb{C}^N$ with an s-sparse vector.

Definition 2.15 (Best s-term approximation error)**.** *Let* $0 < p < \infty, \mathbf{c} \in \mathbb{C}^N$*, and* $s \in \mathbb{N}_0$ *with* $s \leq N$*. The* ℓ^p*-norm best* s*-term approximation error of* $\mathbf{c}$ *is defined as*

$$\sigma_s(\mathbf{c})_p = \min\left\{\|\mathbf{c} - \mathbf{z}\|_p : \mathbf{z} \in \Sigma_s\right\}. \tag{2.61}$$

Ideally, we would like to solve the following optimization problem:

$$\min_{\mathbf{z} \in \mathbb{C}^N} \|\mathbf{z}\|_0 \text{ subject to } A\mathbf{z} = \mathbf{b}. \tag{2.62}$$

Unfortunately, (2.62) is non-convex and typically NP-hard [29]. To overcome this problem, we consider different relaxations of this problem. The first one we will consider is the *Quadratically Constrained Basis Pursuit* (QCBP) problem defined as

$$\min_{\mathbf{z} \in \mathbb{C}^N} \|\mathbf{z}\|_1 \text{ subject to } \|A\mathbf{z} - \mathbf{b}\|_2 \leq \sigma. \tag{2.63}$$

It is a convex optimization problem and can be solved efficiently. The QCBP problem is a well-known problem in compressed sensing, as it turns out the ℓ_1-norm promotes sparsity. To make this notion precise, we need to introduce the *nullspace property*.

Definition 2.16 (Nullspace property)**.** *The matrix* $A \in \mathbb{C}^{n \times N}$ *is said to have the robust nullspace property of order* s *with constant* $0 < \xi < 1$ *and* $\gamma > 1$ *if for all* $S \subseteq \{1, 2, \ldots, N\}$ *with* $|S| \leq s$*, the following holds:*

$$\|\mathbf{v}_S\|_1 < \frac{\rho}{\sqrt{s}}\|\mathbf{z}_{S^C}\|_1 + \gamma\|A\mathbf{v}\|_2, \qquad \forall \mathbf{v} \in \mathbb{C}^N, \mathbf{v} \neq 0, A\mathbf{v} = 0. \tag{2.64}$$

where $[N] = \{1, 2, \ldots, N\}$ *and*

$$(\mathbf{v}_S)_j = \begin{cases} \mathbf{v}_j & \text{if } j \in S, \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 2.5 (ℓ_1-norm promotes sparsity, from [18, Theorem 14.2])**.** *For all matrices* $A \in \mathbb{C}^{n \times N}$*, the following are equivalent:*

- *every* $\mathbf{x} \in \Sigma_s^N$ *is the unique minimizer of* (2.63) *subject to* $A\mathbf{x} = \mathbf{y}$.
- A *has the nullspace property of order* s.

This proposition tells us that if A has the nullspace property, the solution of its associated QCBP problem belongs in the set of s-sparse vectors.

We will use the following result from [3, Chapter 6] about recovery guarantees from the noisy QCBP problem.

Proposition 2.6 ([3, Theorem 6.9])**.** *Let* $x \in \mathbb{C}^d$ *and let* $A \in \mathbb{C}^{n\times d}$ *have the robust nullspace property with constants* ξ *and* γ. *For* $\mathbf{y} = A\mathbf{x} + \mathbf{e}$ *with* $\|\mathbf{e}\|_2 \leq \eta\sqrt{n}$ *for some* $\eta \geq 0$. *Let* $\mathbf{x}^\#$ *be a solution to*

$$\min_{\mathbf{z}\in\mathbb{C}^n} \|\mathbf{z}\|_1 \text{ subject to } \|A\mathbf{z} - \mathbf{y}\|_2 \leq \eta\sqrt{m}.$$

then with probability at least $1 - \epsilon$, *the reconstruction error satisfies*

$$\|\mathbf{x} - \mathbf{x}^\#\|_2 \leq C_1 \frac{\sigma_s(x)_1}{\sqrt{s}} + C_2\eta,$$

where $\sigma_s(x)_1$ *is the best* s*-term approximation error of* x *in the* ℓ_1 *norm, and*

$$C_1 = \frac{2(1+\xi)^2}{1-\xi}, \qquad C_2 = \frac{2(3+\xi)\gamma}{1-\xi}.$$

This proposition tells us that the robust nullspace property ensures that the solution of the QCBP problem can get very close to the solution of the original problem. The reconstruction error is controlled by the best s-term approximation error of the original problem, which should be close to zero if the original signal is close to being s-sparse.

Proving directly that a class of matrices has the nullspace property can be difficult. Fortunately, we can use some sufficient condition as a criterion for the nullspace property. One of which is the *Restricted Isometry Property* (RIP).

Definition 2.17 (Restricted Isometry Property)**.** *A matrix* $A \in \mathbb{C}^{n\times N}$ *has the Restricted Isometry Property of order* s *and constant* $\delta \in (0, 1)$ *if*

$$(1-\delta)\|\mathbf{z}\|_2 \leq \|A\mathbf{z}\|_2 \leq (1+\delta)\|\mathbf{z}\|_2, \qquad \forall \mathbf{z} \in \Sigma_s^N. \tag{2.65}$$

One can think of the RIP as the matrix A being almost an isometry when restricted to the s-sparse vectors, up to a small distortion δ. With the RIP, we can now present the following criterion for the NSP.

Theorem 2.4 (RIP $\Rightarrow$ NSP, from [3, Chapter 6])**.** *Let* $A \in \mathbb{C}^{m\times N}$ *be a matrix with the RIP of order* $2s$ *and constant* $\delta_{2s} < \sqrt{2} - 1$. *Then,* A *has the robust nullspace property of order* s *with constants* $\xi = \frac{\sqrt{2}\delta_{2s}}{1-\delta_{2s}}$ *and* $\gamma = \frac{\sqrt{1+\delta_{2s}}}{1-\delta_{2s}}$.

Finally, it remains to provide a criterion to guarantee that our measurement matrix A has the RIP property. This can be done using the following proposition.

Proposition 2.7 (Bounded orthonormal systems provide the RIP. From [5]). *Let $0 < \delta, \epsilon < 1, A \in \mathbb{C}^{m\times N}$ a measurement matrix obtained from a bounded orthonormal system with constant B and $s \in \{1, 2, \dots, N\}$. Suppose that*

$$m \geq Bc\delta^{-2}sL(s),$$

where $c > 0$ is a universal constant and

$$L(s) = \log\left(\frac{2B^2s}{\delta^2}\right) \cdot \left[\frac{1}{\delta^4}\log\left(\frac{2B^2s}{\delta^2}\right) \cdot \log(2N) + \frac{1}{\delta}\log\left(\frac{2}{\delta\epsilon}\log\left(\frac{2B^2s}{\delta^2}\right)\right)\right]. \tag{2.66}$$

Then, with probability at least $1-\epsilon$, the matrix A satisfies the RIP of order s with constant $\delta_s \leq \delta$.

We can now provide a recovery guarantee results for the QCBP problem combining all that we have seen so far.

Theorem 2.5 (Recovery guarantee for QCBP). *Let $0, \delta, \epsilon < 1$. $A \in \mathbb{C}^{m\times N}$ a measurement matrix obtained from a BOS with constant B. Suppose that $m \geq Bc\delta^{-2}(2s)L(2s)$ with $c > 0$ a universal constant and $L(s)$ be as in (2.66). Then, with probability at least $1-\epsilon$, for $\mathbf{y} = A\mathbf{x} + \mathbf{e}$ with $\|\mathbf{e}\|_2 \leq \eta\sqrt{m}$ for some $\eta \geq 0$. If $\mathbf{x}^\#$ is a solution of the QCBP problem (2.63), then the reconstruction error satisfies*

$$\|\mathbf{x} - \mathbf{x}^\#\|_2 \leq C_1\frac{\sigma_s(\mathbf{x})_1}{\sqrt{s}} + C_2\eta, \tag{2.67}$$

for C_1, C_2 constants depending on δ.

Proof. From Proposition 2.7, we know that A satisfies the RIP of order $2s$ with constant $\delta_{2s} \leq \delta$ with probability at least $1-\epsilon$. Hence, we can apply Theorem 2.4 to obtain that A satisfies the robust nullspace property with constants ξ, γ depending on δ_{2s}. We can then apply Proposition 2.6 to obtain the desired result. □

2.4.2 LASSO-type decoders

Another alternative to the ℓ_0-minimization problem is to introduce a sparse regularization term to the Least Square problem. This leads to the following class of optimization problems called *Least Absolute Shrinkage and Selection Operator* (LASSO) problems:

$$\min_{\mathbf{z}\in\mathbb{C}^N} \|A\mathbf{z} - \mathbf{b}\|_p^q + \lambda\frac{q}{r}\|\mathbf{z}\|_1^r. \tag{2.68}$$

We will focus on the cases where $q = p = 2, r = 1$ (the standard LASSO problem) and $q = r = 2$ (the square LASSO problem). These problems enjoy the following recovery guarantees hinging on the RIP. The RIP allows us to obtain sparsity recovery guarantees from the LASSO-type decoders. This result has been proved in [19].

Proposition 2.8 ([19, Theorem 1])**.** *Let $p \in [1,2], q \geq 1, r \geq 1$. Consider a vector $x \in \mathbb{C}^N$ that is $s-$sparse. Suppose that the vector x is measured via $\mathbf{y} = A\mathbf{x} \in \mathbb{C}^m$ for some $A \in \mathbb{R}^{n \times N}$ satisfying the RIP of order $t = 4s+1$. Then for any $\lambda > 0$, the solution $\mathbf{x}^\lambda$ of the LASSO-type procedure*

$$\min_{\mathbf{z} \in \mathbb{R}^N} \|\mathbf{y} - A\mathbf{z}\|_p^q + \lambda \frac{q}{r} \|\mathbf{z}\|_1^r$$

has sparsity at most proportional to s, namely

$$\|\mathbf{x}^\lambda\|_0 \leq 4s.$$

At this point, we have all the necessary tools to obtain the recovery guarantees for the LASSO problems, as we have seen that sampling from a BOS guarantees in probability the nullspace property.

Chapter 3

Monte Carlo approximation of the Koopman operator

In this chapter, we will present the main results of this thesis for the EDMD problem. We will present a convergence rate for the least-squares problem, and then show how to apply this result to the EDMD problem. Then we will confirm our results through numerical experiments.

3.1 Convergence results

In this section, we provide the main convergence results for the approximation of the Koopman operator using EDMD. We will first provide a convergence results for the least-squares problem, and then show how to apply this result to the EDMD problem. Finally, we show how to apply Theorem 2.3 to the EDMD problem.

3.1.1 Convergence for the least-squares problem

In this subsection, we briefly go back to the least-squares setting of Section 2.3. We will first provide a convergence rate in this setting, and then explain how to apply this result to the EDMD problem. First, we introduce the following definitions that will be helpful in deriving the desired bound.

Definition 3.1 (Truncation operator)**.** *The truncation operator* $\mathcal{F}_L : L^2_\rho(\mathbb{X}) \to L^2_\rho(\mathbb{X})$ *is defined as*

$$\mathcal{F}_L(g) := \begin{cases} g & \text{if } \|g\|_{L^2_\rho(\mathbb{X})} \leq L \\ Lg/\|g\|_{L^2_\rho(\mathbb{X})} & \text{otherwise.} \end{cases}, \quad \forall g \in L^2_\rho(\mathbb{X}) \tag{3.1}$$

Definition 3.2 (Best approximation of f in $\mathcal{P}$)**.** *Let* f *be a function in* $L^2_\rho(\mathbb{X})$ *and* $\mathcal{P}$ *a subspace of* $L^2_\rho(\mathbb{X})$*. The best approximation of* f *in* $\mathcal{P}$ *is defined as*

$$f^* \in \arg\min_{g \in \mathcal{P}} \|f - g\|_{L^2_\rho(\mathbb{X})}. \tag{3.2}$$

Lemma 3.1. *The truncation operator is a contraction on $L^2_\rho(\mathbb{X})$, that is*

$$\|\mathcal{F}_L(g) - \mathcal{F}_L(f)\|_{L^2_\rho(\mathbb{X})} \leq \|g - f\|_{L^2_\rho(\mathbb{X})}. \tag{3.3}$$

Proof. We observe that the truncation operator $\mathcal{F}_L$ is the projection operator on the ball of radius L in $L^2_\rho(\mathbb{X})$. Since the truncation operator is a projection onto a closed convex set, using the result in [7, Corollary 8.5], it is Lipschitz with constant 1, hence a contraction. □

With this helpful operator established, we will use the following lemmas to establish the desired convergence rate.

Lemma 3.2. *Let $0 < \delta, \epsilon < 1$. Let f^* be as in (3.2), and $\hat{f}$ the noiseless least-squares approximation (2.23) (i.e., with noise vector $\mathbf{e} = (e_i)_{i=1}^n = 0$) of f in $\mathcal{P}$ with $n \geq c_\delta \cdot \kappa \cdot \log(s/\epsilon)$ where c_δ is as in Theorem 2.2. Then for an orthonormal basis $\{\psi_i\}_{i=1}^s$, we have that*

$$\|\hat{f} - f^*\|^2_{L^2_\rho(\mathbb{X})} \leq \frac{1}{(1-\delta)^2} \sum_{i=1}^{s} |\langle f - f^*, \psi_i \rangle_{\text{disc}}|^2 . \tag{3.4}$$

Proof. Let $\mathbf{d} \in \mathbb{C}^s$ such that $\hat{f} - f^* = \sum_{i=1}^{s} d_i \psi_i$. Then

$$\begin{aligned}
\|\hat{f} - f^*\|^2_{\text{disc}} &= \left\langle \hat{f} - f^*, \hat{f} - f^* \right\rangle_{\text{disc}} \\
&= \left\langle f - f^*, \hat{f} - f^* \right\rangle_{\text{disc}} \\
&= \left\langle f - f^*, \sum_{i=1}^{s} d_i \psi_i \right\rangle_{\text{disc}} \\
&= \sum_{i=1}^{s} d_i \langle f - f^*, \psi_i \rangle_{\text{disc}} \\
&\leq \|\mathbf{d}\|_2 \sqrt{\sum_{i=1}^{s} |\langle f - f^*, \psi_i \rangle_{\text{disc}}|^2} \\
&\leq \|\mathbf{d}\|_2 \sqrt{\sum_{i=1}^{s} |\langle f - f^*, \psi_i \rangle_{\text{disc}}|^2} \\
&= \|\hat{f} - f^*\|_{L^2_\rho(\mathbb{X})} \sqrt{\sum_{i=1}^{s} |\langle f - f^*, \psi_i \rangle_{\text{disc}}|^2},
\end{aligned}$$

where the last equality comes from Parseval's Identity. Now we recall that $\frac{\|p\|_{\text{disc}}}{\|p\|_{L^2_\rho}} \geq \alpha, \forall p \in \mathcal{P}$.

Hence,

$$\|\hat{f} - f^*\|^2_{L^2_\rho(\mathbb{X})} \leq \alpha^{-2}\|\hat{f} - f^*\|^2_{\text{disc}}$$

$$\|\hat{f} - f^*\|^2_{L^2_\rho(\mathbb{X})} \leq \alpha^{-2}\|\hat{f} - f^*\|_{L^2_\rho(\mathbb{X})}\sqrt{\sum_{i=1}^{s} |\langle f - f^*, \psi_i\rangle_{\text{disc}}|^2}$$

$$\|\hat{f} - f^*\|_{L^2_\rho(\mathbb{X})} \leq \alpha^{-2}\sqrt{\sum_{i=1}^{s} |\langle f - f^*, \psi_i\rangle_{\text{disc}}|^2}$$

$$\|\hat{f} - f^*\|^2_{L^2_\rho(\mathbb{X})} \leq \frac{1}{(1-\delta)^2}\sum_{i=1}^{s} |\langle f - f^*, \psi_i\rangle_{\text{disc}}|^2,$$

where the last inequality comes from Theorem 2.2. □

Lemma 3.3. *For $\{\psi_i\}_{i=1}^\infty$ orthonormal basis of $L^2_\rho(\mathbb{X})$ and $\mathcal{P} = \text{span}\{\psi_i\}_{i=1}^s$, then we have that for $i = 1, \ldots, s$,*

$$\mathbb{E}\left(\langle f - f^*, \psi_i\rangle^2_{\text{disc}}\right) = \frac{\|(f - f^*)\psi_i\|^2_{L^2_\rho(\mathbb{X})}}{n}.$$

Proof. Let $f = \sum_{i=1}^{\infty} c_i\psi_i$, then since $\{\psi_i\}_{i=1}^\infty$ is an orthonormal basis of $L^2_\rho(\mathbb{X})$, we have that $f^* = \sum_{i=1}^{s} c_i\psi_i$. Therefore,

$$\mathbb{E}\left[(f - f^*)\psi_i\right] = \int_{\mathbb{X}} (f - f^*)\psi_i d\rho(\mathbf{y}) = \int_{\mathbb{X}} \left(\sum_{j=1}^{\infty} c_j\psi_j - \sum_{j=1}^{s} c_j\psi_j\right)\psi_i d\rho(\mathbf{y}) \tag{3.5}$$

$$= \int_{\mathbb{X}} \left(\sum_{j=s+1}^{\infty} c_j\psi_j\right)\psi_i d\rho(\mathbf{y}) \tag{3.6}$$

$$= 0, \tag{3.7}$$

where the last equality comes from the fact that $\{\psi_i\}$ is an orthonormal basis, and $i \leq s$ while $j \geq s+1$. Now, consider the random variable $Y = (f(\mathbf{y}) - f^*(\mathbf{y}))\psi_i(\mathbf{y}), \mathbf{y}$ sampled from ρ. Then from the above computation, $\mathbb{E}[Y] = 0$. Letting $\overline{Y} = \langle f - f^*, \psi_i\rangle_{\text{disc}}$, we see that $\mathbb{E}[\overline{Y}] = \mathbb{E}[Y] = 0$ since samples are i.i.d. with $\mathbf{y}_1, \ldots, \mathbf{y}_n$ sampled from ρ. Then we have

$$\text{Var}(\overline{Y}) = \mathbb{E}\left[\left(\overline{Y} - \underbrace{\mathbb{E}(\overline{Y})}_{=0}\right)^2\right] = \mathbb{E}\left[\langle f - f^*, \psi_i\rangle^2_{\text{disc}}\right] = \frac{\sigma^2}{n},$$

where $\sigma^2 = \text{Var}(Y) = \mathbb{E}[(Y - \mathbb{E}[Y])^2] = \mathbb{E}[Y^2] = \|(f - f^*)\psi_i\|^2_{L^2_\rho(\mathbb{X})}$, as $\mathbb{E}[Y] = 0$. This gives the desired result. □

With these two previous results, we arrive at the following theorem.

Theorem 3.1. *Let $0 < \delta, \epsilon < 1$. Let $\mathcal{P} \subset L^2_\rho(\mathbb{X})$ be a finite dimensional subspace of dimension $dim(\mathcal{P}) = s$ with $\{\psi_i\}_{i=1}^s$ an orthonormal basis of $L^2_\rho(\mathbb{X})$ and $\{\psi_i : i = 1, \dots, s\}$ an orthnormal basis of $\mathcal{P}$. Let $\mathbf{y}_1, \dots \mathbf{y}_n$ be i.i.d. $n \geq c_\delta \cdot \kappa \cdot \log(s/\epsilon)$ sample points in $\mathbb{X}$ sampled from ρ, where*

$$c_\delta = ((1-\delta)\log(1-\delta) + \delta)^{-1}.$$

Suppose that $\alpha = \alpha(\mathcal{P}, \{\mathbf{y}_i\}_{i=1}^n)$. Moreover, let L be such that $\|f^*\|_{L^2_\rho(\mathbb{X})} \leq L$, where f^* is defined in (3.2). Then, for all $f \in L^2_\rho(\mathbb{X}) \cap C(\mathbb{X})$, with probability $1 - \epsilon$,

$$\mathbb{E}\|\mathcal{F}_L(\hat{f}) - f^*\|^2_{L^2_\rho(\mathbb{X})} \leq \frac{1}{n} \cdot \frac{1}{(1-\delta)^2} \sum_{i=1}^s \|(f - f^*)\psi_i\|^2_{L^2_\rho(\mathbb{X})} + 4L^2\epsilon.$$

Proof. Let $\mathbf{d} \in \mathbb{C}^s$ such that $\mathcal{F}_L(\hat{f}) - f^* = \sum_{i=1}^s d_i \psi_i$. Moreover, let $E = \{(\mathbf{y}_1, \dots, \mathbf{y}_n) \in \mathbb{X}^n : \alpha > \sqrt{1-\delta}\}$. Then

$$\begin{aligned}
\mathbb{E}\|\mathcal{F}_L(\hat{f}) - f^*\|^2_{L^2_\rho(\mathbb{X})} &= \int_E \|\mathcal{F}_L(\hat{f}) - f^*\|^2_{L^2_\rho(\mathbb{X})} d\rho^n + \int_{E^c} \|\mathcal{F}_L(\hat{f}) - f^*\|^2_{L^2_\rho(\mathbb{X})} d\rho^n \\
&\leq \int_E \|\hat{f} - f^*\|^2_{L^2_\rho(\mathbb{X})} d\rho^n + \underbrace{\int_{E^c} (\|\mathcal{F}_L(\hat{f})\|_{L^2_\rho} + \|f^*\|_{L^2_\rho(\mathbb{X})})^2 d\rho^n}_{\leq 4L^2\epsilon}.
\end{aligned}$$

Since $\mathbb{P}(\alpha \leq \sqrt{1-\delta}) \leq \epsilon$ from Theorem 2.2. Let us now focus on the first integral above:

$$\begin{aligned}
\int_E \|\hat{f} - f^*\|^2_{L^2_\rho(\mathbb{X})} d\rho^n &\leq \int_E \frac{1}{(1-\delta)^2} \sum_{i=1}^s \langle f - f^*, \psi_i \rangle^2_{\text{disc}} \, d\rho^n \quad \text{(using Lemma 3.2)} \\
&\leq \int_{\mathbb{X}} \frac{1}{(1-\delta)^2} \sum_{i=1}^s \langle f - f^*, \psi_i \rangle^2_{\text{disc}} \, d\rho^n \text{ as } \langle f - f^*, \psi_i \rangle^2_{\text{disc}} \geq 0 \\
&= \frac{1}{(1-\delta^2)} \sum_{i=1}^s \mathbb{E}\left[\langle f - f^*, \psi_i \rangle^2_{\text{disc}}\right] \\
&= \frac{1}{n} \cdot \frac{1}{(1-\delta)^2} \sum_{i=1}^s \|(f - f^*)\psi_i\|^2_{L^2_\rho(\mathbb{X})} \quad \text{(using Lemma 3.3).}
\end{aligned}$$

This proves the desired result. □

Remark 3.1. *The truncation operator is necessary to bound the expected value of the error over E^c. By construction, we have that $\|f^*\|_{L^2_\rho(\mathbb{X})} \leq L$, but we do not have any control over $\|\hat{f}\|_{L^2_\rho(\mathbb{X})}$.*

3.1.2 Convergence for the EDMD problem

In this subsection, we show how to extend Theorem 3.1 to the EDMD problem. We start by making the following observation. We can rewrite the EDMD problem (2.16) as follows:

$$\min_{A \in \mathbb{R}^{l \times m}} \|\mathbf{\Phi}_n^\tau - A\mathbf{\Psi}_n\|_F^2 = \min_{A \in \mathbb{R}^{l \times m}} \|(\mathbf{\Phi}_n^\tau)^* - \mathbf{\Psi}_n^* A^*\|_F^2$$

Denote by $\mathbf{c}_i, \mathbf{b}_i$ the i-th columns of $\mathbf{\Phi}_n^{\tau^*}, A^*$ respectively. Then using the separability of the objective function in the columns of A^* we have that

$$\begin{aligned}
\min_{A\in\mathbb{R}^{l\times m}} \|(\mathbf{\Phi}_n^\tau)^* - \mathbf{\Psi}_n^* A^*\|_F^2 &= \min_{\mathbf{c}_1,\dots,\mathbf{c}_l\in\mathbb{R}^m} \sum_{i=1}^{l} \|\mathbf{b}_i - \mathbf{\Psi}_n^* \mathbf{c}_i\|_2^2 \\
&= \sum_{i=1}^{l} \min_{\mathbf{c}_i\in\mathbb{R}^m} \|\mathbf{b}_i - \mathbf{\Psi}_n^* \mathbf{c}_i\|_2^2 \\
&= \sum_{i=1}^{l} \min_{\mathbf{c}_i\in\mathbb{R}^m} \|\mathbf{b}_i^* - \mathbf{c}_i^* \mathbf{\Psi}_n\|_2^2 \\
&= \sum_{i=1}^{l} \min_{p\in\operatorname{span}\{\boldsymbol{\psi}\}} \|p - \mathcal{K}_{mn}^\tau[\phi_i]\|_{L_\rho^2(\mathbb{X})}^2
\end{aligned}$$

Hence the EDMD problem is equivalent to solving l independent least-squares problems, one for each observable in the dictionary $\boldsymbol{\phi}$. Thus, we can apply the theory what we developed in Section 2.3 to the EDMD problem. We are now ready to state our convergence result for EDMD.

Theorem 3.2. *Let $0 < \delta, \epsilon < 1$. Let $\mathbf{\Phi}_n^\tau, \mathbf{\Psi}_n$, be the observables matrices as in (2.14) from n sample points $\mathbf{x}_1, \dots, \mathbf{x}_n$ sampled i.i.d from ρ, with $n \geq c_\delta \cdot \kappa \cdot \log(\frac{s}{\epsilon})$. Moreover, suppose the dictionary $\boldsymbol{\psi}$ is orthonormal with respect to ρ and assume that $\|\mathcal{K}_m^\tau[\phi_j]\|_{L_\rho^2(\mathbb{X})} \leq L_j$, where $\mathcal{K}_m^\tau$ is defined in eq2.21. Let A be the diagonal matrix such that $A_{jj} = \min\left\{1, \frac{L_j}{\|\mathcal{K}_m^\tau[\phi_j]\|}\right\}$. Define $\tilde{K}_{mn}^\tau = AK_{mn}^\tau$. Then, with probability at least $1 - l\epsilon$,*

$$\mathbb{E}\|\tilde{K}_{mn}^\tau - K_m^\tau\|_F^2 \leq \frac{1}{n} \cdot \frac{1}{(1-\delta)^2} \sum_{j=1}^{l}\sum_{i=1}^{m} \|(\mathcal{K}^\tau[\phi_j] - \mathcal{K}_m^\tau[\phi_j])\psi_i\|_{L_\rho^2(\mathbb{X})} + 4\left(\sum_{j=1}^{l} L_j^2\right)\epsilon. \quad (3.8)$$

Proof.

$$\begin{aligned}
\mathbb{E}\|\tilde{K}_{mn}^\tau - K_m^\tau\|_F^2 &= \mathbb{E}\left(\sum_{j=1}^{l} \|A_{jj}\,(K_{mn}^\tau)_j - (K_m^\tau)_j\,\|_2^2\right) \\
&= \sum_{j=1}^{l} \mathbb{E}\left(\|A_{jj}\,(K_{mn}^\tau)_j - \left(K_m^\tau\|_2^2\right)_j\right)\|_2^2 \\
&= \sum_{j=1}^{l} \mathbb{E}\left(\|\mathcal{F}_L(\mathcal{K}_{mn}^\tau[\phi_j]) - \mathcal{K}_m^\tau[\phi_j]\|_2^2\right) \\
&\leq \frac{1}{n} \cdot \frac{1}{(1-\delta)^2} \sum_{j=1}^{l}\sum_{i=1}^{m} \left(\|(\mathcal{K}^\tau[\phi_j] - \mathcal{K}_m^\tau[\phi_j])\psi_i\|_{L_\rho^2(\mathbb{X})}\right) + 4\sum_{j=1}^{l}\left(L_j^2\right)\epsilon,
\end{aligned}$$

where the last inequality comes from using Theorem 3.1 for each of the l least-squares problems. Each of these problems is independent, and the corresponding error bounds hold

each with probability at least $1-\epsilon$. Denote by P_i the event that the bound on the error of the i-th least-squares problem holds. Then,

$$\begin{aligned}\mathbb{P}\left(\bigcap_{i=1}^{l} P_i\right) &= 1-\mathbb{P}\left(\bigcup_{i=1}^{l} P_i^c\right)\\ &\geq 1-\sum_{i}^{l}\mathbb{P}(P_i^c)\\ &= 1-l\epsilon,\end{aligned}$$

where we used the fact that $\mathbb{P}\left(\bigcup_{i=1}^{l}\overline{P}_i\right) \leq \sum_i^n \mathbb{P}(\overline{P_i})$ through subadditivity of the probability measure (this is called the union bound). This concludes the proof. □

This theorem seems to contradict our claim that this method converges to the Monte Carlo rate of $\frac{1}{\sqrt{n}}$ because of the constant term in the bound. However, we note that we can replace this ϵ by solving for it in the sample complexity estimate $n \geq c_\delta \cdot \kappa \cdot \log(s/\epsilon)$. Solving for ϵ yields $\epsilon \leq s e^{-\frac{n}{c_\delta \cdot \kappa}}$. Hence, this constant terms vanishes at an exponential rate in n, and we can indeed achieve the desired Monte Carlo rate of $\frac{1}{\sqrt{n}}$.

We also note that this result requires orthonormality of the dictionary $\boldsymbol{\psi}$, this is used on (3.6). This is a strong assumption, necessary for this proof but from empirical observations it seems that this assumption could be relaxed.

3.1.3 Beating the Monte Carlo rate

In this subsection, we extend the results presented in Subsection 2.3.4 to the EDMD problem. We do this by once again noticing that the EDMD problem is equivalent to solving the l independent least-squares problems

$$\min_{p\in\operatorname{span}\{\boldsymbol{\psi}\}} \|p-\mathcal{K}_{mn}^{\tau}[\phi_i]\|_{L^2_\rho(\mathbb{X})}^2.$$

Hence, we can simply apply Theorem 2.3 to each of these problems to obtain the desired result.

Theorem 3.3. *Let $0<\epsilon<1, p\in(0,1), \rho$ be either the uniform or Chebyshev measure on $\mathbb{X}=[-1,1]^N, n\geq 3$ and $\mathbf{y}_1,\dots,\mathbf{y}_n$ sampled i.i.d. from ρ. Then there exist sets $S_i\subseteq\mathbb{N}_0^N$ of cardinality $|S_i|\leq\lceil n/\log(n/\epsilon)\rceil$ such that for all $i=1,\dots,l$, we have that with probability at least $1-l\epsilon$, if $\mathcal{K}[\phi_i]$ admits a holomorphic extension to a Bernstein polyellipse $\mathcal{E}_{\boldsymbol{\rho}}\subset\mathbb{C}^N$ for all $i=1,\dots,l$, then the EDMD approximation of the Koopman operator $\mathcal{K}_{mn}^{\tau}$ of $\mathcal{K}^{\tau}$ is unique and satisfies the following bound for all $i=1,\dots,m$:*

$$\|\mathcal{K}_{mn}^{\tau}[\phi_i]-\mathcal{K}^{\tau}[\phi_i]\|_{L^2_\rho(\mathbb{X})} \leq C(\boldsymbol{\rho},p)\left(\frac{n}{\log(n/\epsilon)}\right)^{\frac{1}{2}-\frac{1}{p}}, \tag{3.9}$$

where $C(\boldsymbol{\rho},p)$ is a constant that depends on $\boldsymbol{\rho}$ and p only.

Proof. We simply apply Theorem 2.3 to the l independent least-squares problems. We note that the only difference with Theorem 3.2 is that we have a different constant $C(\boldsymbol{\rho}, p)$, that is the maximum of each of the constants of the subproblems C_i. □

Remark 3.2. *This result tells us that it is possible to beat the Monte Carlo rate of $\frac{1}{\sqrt{n}}$ for the EDMD problem under suitable smoothness assumptions. In fact, just as Theorem 2.3, it predicts that it should be possible to beat any algebraic convergence rate.*

3.2 Numerical experiments

In this section, we provide numerical experiments to observe the convergence rates that were derived in Theorem 3.2 and Theorem 3.3. We study the Koopman Operator associated with (i) the logistic map, a deterministic and discrete-time dynamical system, (ii) the stochastic logistic map, a stochastic and disrete-time dynamical system, and finally (iii) the Thomas model, a nonpolynomial, 3-dimensional continuous-time dynamical system. Each experiment will be run for the unweighted EDMD problem, as the effect of adding weights in the EDMD problem were not considered in our results, the code is available for free at [17].

3.2.1 Error measurement

In this subsection, we will explain how the error is measured in each experiment. We distinguish between two different setups:

1. **Fixed Dictionary size:** In this case, we fix our dictionary size $|\boldsymbol{\psi}| = m, |\boldsymbol{\phi}| = l$, and we only vary the number of data points n.

2. **Variable Dictionary size:** In this case, we let the size of our dictionary to scale with the number of data points.

In the first case, we know from Theorem 2.1 that EDMD converges to the true Koopman operator as the number of data points goes to infinity. Hence, the error is simply measured by comparing the EDMD approximation of the Koopman operator to a high-fidelity approximation of the Koopman operator obtained through an EDMD simulation with 100 times more data points than the maximum number of data points used in the experiments. This approximation is denoted by $K_m^{\tau,HF}$ and the error is measured as

$$\text{error} = \frac{\|K_{mn}^{\tau} - K_m^{\tau,HF}\|_F}{\|K_m^{\tau,HF}\|_F}.$$

In the second case, since we are approaching the limit in a different order than Theorem 2.1, we will not be able to rely on a high-fidelity approximation to compute the error. Instead, we pick a smooth test function φ and measure the error as a Monte Carlo approximation $\frac{\|\mathcal{K}_{mn}^{\tau}[\varphi] - \mathcal{K}^{\tau}[\varphi]\|_{L^2_{\rho}(\mathbb{X})}}{\|\mathcal{K}^{\tau}[\varphi]\|_{L^2_{\rho}(\mathbb{X})}}$ for discrete time dynamics, with

$$\mathcal{K}_{mn}^{\tau}[\varphi] = \mathbf{c}^T \cdot K_{mn}^{\tau} \cdot \psi,$$

where $\mathbf{c}$ is the vector of coefficients of the projection of φ onto $\text{span}\{\boldsymbol{\psi}\}$. $\mathcal{K}^\tau[\varphi]$ is computed explicitly.

For continuous-time dynamics, we instead rely on the Lie derivative. The error is then measured as a Monte Carlo approximation of $\frac{\|\mathcal{L}_{mn}[\varphi]-\mathcal{L}[\varphi]\|_{L^2_\rho(\mathbb{X})}}{\|\mathcal{L}[\varphi]\|_{L^2_\rho(\mathbb{X})}}$. We use the fact than an approximate Lie derivative is immediately available once our EDMD approximation of the Koopman operator is computed, as we can set

$$L_{mn} = \frac{K^\tau_{mn} - \Theta_{mn}}{\tau}, \qquad \mathcal{L}_{mn}[\varphi] = \mathbf{c}^T \cdot L_{mn} \cdot \psi,$$

where Θ_{mn} is as in Assumption 2 of [10]. We use the Lie derivative instead of the Koopman operator as this eliminates the scaling induced by the sampling rate τ. The L^2_ρ-error is then approximated by sampling i.i.d. M samples points $\mathbf{z}_i$ from the uniform distribution on $\mathbb{X}$, and computed as follows:

$$\text{error} = \frac{\sqrt{\frac{1}{M}\sum_{i=1}^M |\mathcal{K}^\tau_{mn}[\varphi](\mathbf{z}_i) - \mathcal{K}^\tau[\varphi](\mathbf{z}_i)|^2}}{\sqrt{\frac{1}{M}\sum_{i=1}^M |\mathcal{K}^\tau[\varphi](\mathbf{z}_i)|^2}} \quad \text{for discrete time dynamics,}$$

$$\text{error} = \frac{\sqrt{\frac{1}{M}\sum_{i=1}^M |\mathcal{L}_{mn}[\varphi](\mathbf{z}_i) - \mathcal{L}[\varphi](\mathbf{z}_i)|^2}}{\sqrt{\frac{1}{M}\sum_{i=1}^M |\mathcal{L}[\varphi](\mathbf{z}_i)|^2}} \quad \text{for continuous time dynamics.}$$

3.2.2 The Logistic map

The logistic map is a one-dimensional, deterministic, and discrete dynamical system. It is defined by the following recurrence relation $X_{n+1} = 2X_n^2 - 1$, for $X_n \in [-1, 1]$. It has been studied extensively in the literature and is known to exhibit chaotic behaviour. First, we study the impact of the choice of dictionaries $\boldsymbol{\phi}, \boldsymbol{\psi}$. The summary of those experiments are outlined in Figure 3.1 and Figure 3.2 . Focusing first on the case when $\boldsymbol{\psi} = \boldsymbol{\phi}$, For polynomial bases, we consider observables functions going up to degree 5. For the Fourier basis, we consider the observables $\{\cos(n\pi x), \sin(n\pi x)\}_{n=0}^5$. In general, we will say that the Fourier basis $\{\cos(n\pi x), \sin(n\pi x)\}_{n=0}^N$ has order N. We can observe the Monte Carlo rate as predicted by Theorem 3.2. We should note that even though we do not always sample from the orthogonality measure (or in fact use orthogonal dictionaries for the case of the monomials), we always observe the same rate, as documented in the literature [30, 39, 40]. We observe that the best performance are obtained when we sample according to the Christofell sampling scheme, followed by the orthogonality measure in all cases but the Jacobi polynomial basis. This can be explained as the Jacobi polynomials are skewed to the left, but our dynamical system is symmetric. For the Chebyshev polynomial, we get that sampling from the orthogonality measure is slightly more advantageous than Christofell sampling. Finally, we remark that we did not investigate the effect of ergodic sampling, as for the logistic map, ergodic sampling is equivalent to sampling from the Chebyshev measure [27, Example 4.1.2].

For the case of $|\boldsymbol{\psi}| > |\boldsymbol{\phi}|$, we picked a larger dictionary for $|\boldsymbol{\psi}|$ in such a way that the nonlinearity $\mathcal{K}[\phi]$, for $\phi \in \text{span}\{\boldsymbol{\phi}\}$ is fully described in $\text{span}\{\boldsymbol{\psi}\}$ for the polynomial bases.

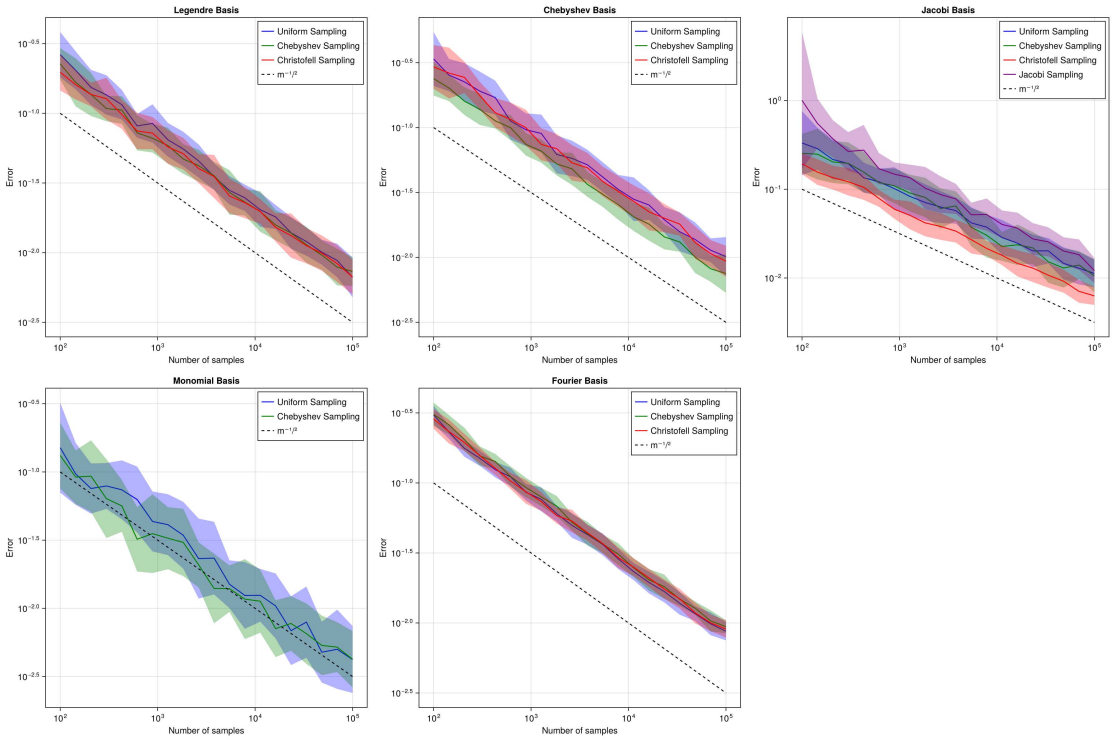

Figure 3.1: Convergence of the approximation of the Koopman operator for the logistic map with different dictionary size, $\boldsymbol{\psi} = \boldsymbol{\phi}$. Each figure represent different bases for the dictionaries, and each curve represents a different sampling measure for the data.

Since the logistic map is a polynomial of degree 2, we pick $\boldsymbol{\phi}$ to contain polynomials up to degree 5 and $|\boldsymbol{\psi}|$ to contain polynomials up to degree 10. This leads to an immediate recovery of the Koopman operator, with an error approaching machine precision almost instantly. The only case where this is not happening is for the Fourier basis, as it is impossible to fully resolve the nonlinearities for this basis.

Next we turn our attention to the question of recovery of non-smooth functions, for this experiment, we added the indicator function $f(x) = \mathbb{1}_{[-\frac{1}{2},\frac{1}{2}]}(x)$ to both dictionaries $\boldsymbol{\psi}$ and $\boldsymbol{\phi}$. The result is presented in Figure 3.3. We are in the setting when $|\boldsymbol{\psi}| > |\boldsymbol{\phi}|$. The addition of the step function is more akin to the result obtained for the Fourier basis, as it can never be fully recovered due to its non-smoothness. It should be noted that ergodic sampling —i.e. sampling from a single trajectory— seems to be performing worse than the other sampling methods, this is expected as ergodic sampling can converge arbitrarily slowly [26].

In the last test we will try is to see how the convergence rate can be improved by letting the dictionary size vary with the number of data points as in Theorem 3.3. We pick the relation $n = 3m^2 \log(m)$, where the constant of 3 has been determined through numerous tests to determine a scaling coefficient large enough to satisfy our theorem. Experiments to determine the correct scaling are presented in Figure 3.4. The results with all the different bases are shown in Figure 3.5. As we can see, the convergence is much faster than the Monte Carlo rate, the graph is presented on a semilog scale, so that exponential trends look linear, hence we can infer that we obtain exponential convergence in this case, with the

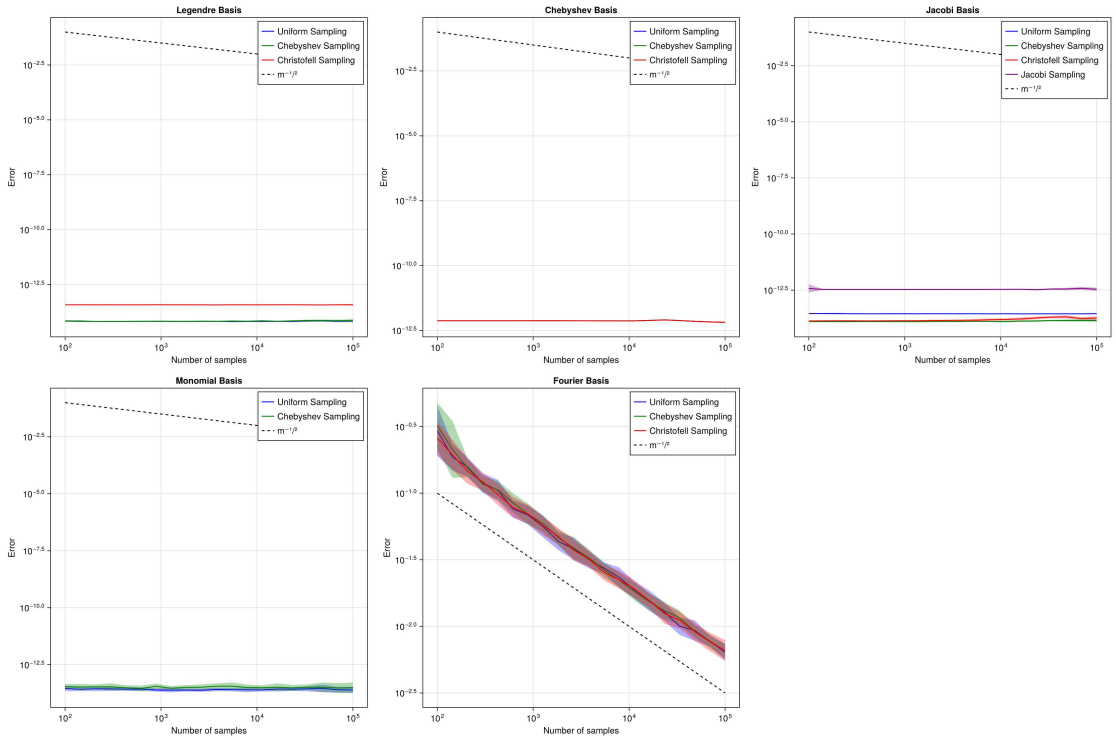

Figure 3.2: Convergence of the approximation of the Koopman operator for the logistic map with same dictionary size, $|\boldsymbol{\psi}| > |\boldsymbol{\phi}|$. Each figure represent different bases for the dictionaries, and each curve represents a different sampling measure for the data.

error approaching machine precision in the cases of Legendre and Chebyshev polynomials. It should be noted that the monomials seem to drift off the convergence rate, but this could be tied to the condition number of the measurement matrix in that case, leading to numerical instabilities [22].

3.2.3 The stochastic logistic map

The stochastic logistic map is a one-dimensional, stochastic, and discrete dynamical system. It is defined on (2.9). While this is very similar to the deterministic logistic map, we decided to study it to showcase that the theory presented does indeed hold in the stochastic regime. The results are presented in Figure 3.6 and Figure 3.7. We observe that the Monte Carlo rate is still achieved, even though the recovery is generally worse than in the deterministic case. Once again, Christofell sampling seems to be the best sampling scheme, followed by the orthogonality measure in most cases.

We observe that for the deterministic and stochastic logistic map, using Jacobi polynomials with $\alpha = 1, \beta = 0$ seem to be yielding relatively worse results, comparable to those of monomials, even when sampling from the orthogonality measure. This can be explained as with those parameters, the Jacobi polynomials are skewed to the left, while the dynamics of the system are symmetric, hence it is harder to describe the dynamics in that setting.

In the experiment where we let the dictionary size vary with the number of data points,

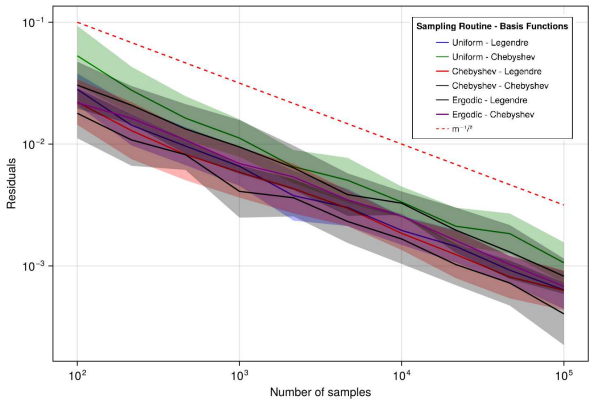

Figure 3.3: Convergence of the approximation of the logistic map with non-smooth function $f(x) = \mathbb{1}_{[-\frac{1}{2},\frac{1}{2}]}(x)$ added to both dictionaries. $|\boldsymbol{\psi}| > |\boldsymbol{\phi}|$ like in Figure 3.2.

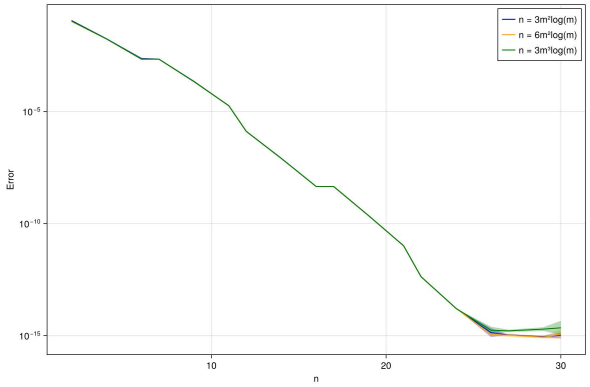

Figure 3.4: Convergence of the approximation of the Koopman operator with variable dictionary size for the logistic map. Multiple scaling laws are tested.

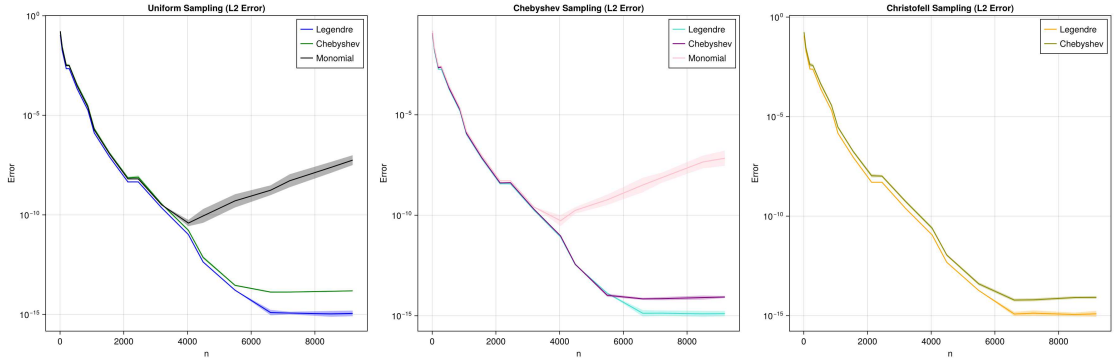

Figure 3.5: Convergence of the approximation of the Koopman operator with variable dictionary size for the logistic map. The relation is $n = 3m^2 \log(m)$.

the results are outlined in Figure 3.8. We notice that at first, the convergence is much faster than the Monte Carlo rate, but beyond 10^4 data points, it slows down drastically. For the Legendre polynomial, growing the size of the dictionary doesn't even seem to improve the recovery. It could be the case that not enough data was sampled for the dictionary size taken, as the inherent stochasticity of the map would require a larger number of data points to fully explore the parameter space. In practice this was not feasible as those simulations would take too long to run. Another reason explaining why this experiment is not conclusive is through numerical instability. Our choice of observable was $\phi(x) = x^2$. As soon as the dictionary size reaches 3, this observable belongs in our approximation space. As we increase the number of data points and dictionary size, there is no new information to be gained, leading to numerical artifacts populating the approximation. This effect is studied in Figure 3.9 where the Monte Carlo rate can be recovered in this setting, but not improved.

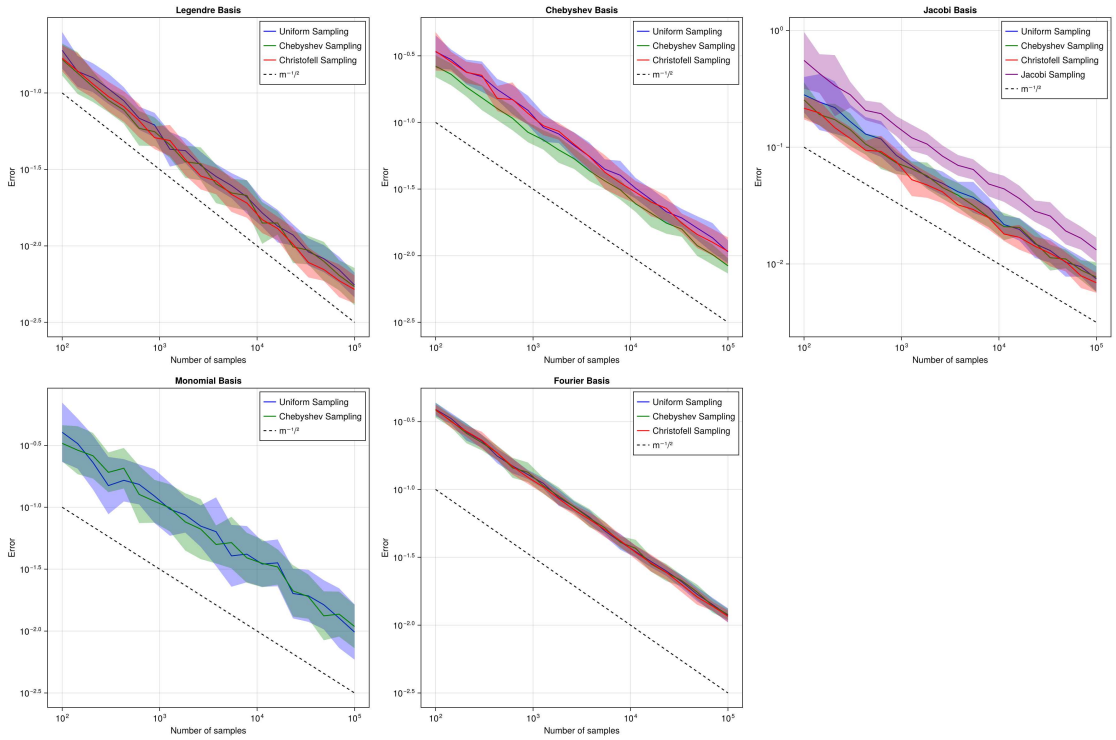

Figure 3.6: Convergence of the approximation of the Koopman operator for the stochastic logistic map with $\boldsymbol{\psi} = \boldsymbol{\phi}$.

3.2.4 The Thomas model

The Thomas model is a three-dimensional, continuous-time dynamical system introduced previously on (2.10). The behaviour of this system is chaotic, with a cyclically symmetric attractor originally identified by René Thomas in [34]. It is defined as follows:

$$\dot{\mathbf{x}} = f(\mathbf{x}) = \begin{pmatrix} \sin(x_2) - bx_1 \\ \sin(x_3) - bx_2 \\ \sin(x_1) - bx_3 \end{pmatrix}, \tag{3.10}$$

where b is a parameter taken to be 0.208186 for chaotic behaviour, the chaotic attractor can be seen on Figure 3.10. We perform a change of variable to ensure the solution stays within the cube $[-1, 1]^3$ so that we can once again use the Legendre and Chebyshev polynomial bases. By observing that the original solutions of the Thomas model with that choice of parameter are within the cube $[-5, 5]^3$, we perform the change of variable $\mathbf{y} = \frac{1}{5}\mathbf{x}$ leading to the re-scaled system $\dot{\mathbf{y}} = \frac{1}{5}f(5\mathbf{y})$. Figure 3.11 presents the experiments with both dictionaries being the same (functions of order 3), while Figure 3.12 has dictionaries of various size, with $\boldsymbol{\phi}$ containing bases functions of order 3 and $\boldsymbol{\psi}$ containing functions of order 6. We observe the predicted convergence rate of $\frac{1}{\sqrt{n}}$ in all cases, both when $\boldsymbol{\psi} = \boldsymbol{\phi}$ and when $|\boldsymbol{\psi}| > |\boldsymbol{\phi}|$. Once again, the best sampling scheme seem to be Christofell sampling followed by sampling from the orthogonality measure, except the Jacobi polynomials. We observe that the Jacobi polynomials with $\alpha = 1, \beta = 0$ seem to be performing rather poorly, even when sampling

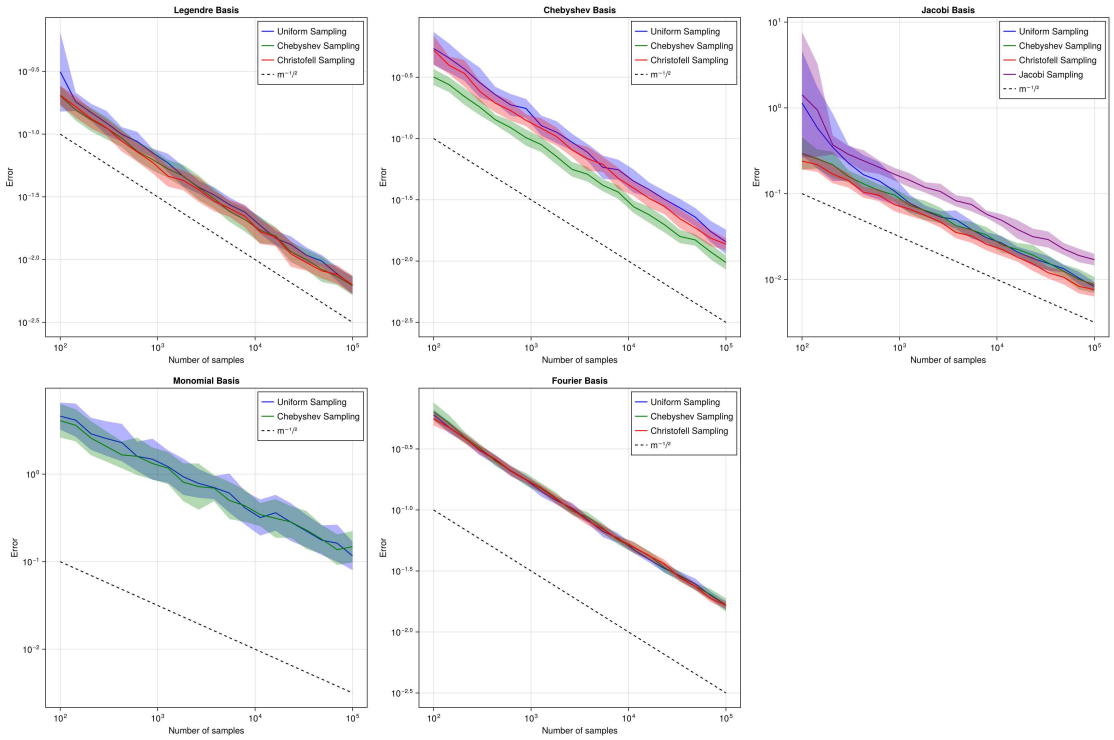

Figure 3.7: Convergence of the approximation of the Koopman operator for the stochastic logistic map with $|\boldsymbol{\psi}| > |\boldsymbol{\phi}|$.

from the orthogonality measure. This is similar to what we observed in the case of the previous examples (Figure 3.1 and Figure 3.6). We suppose that this is due to the fact that the dynamics of the system are symmetric, while the Jacobi polynomials are skewed to the left. The experiments outlined in this section show that the EDMD technique is able to

recover the Koopman operator for a variety of systems, both deterministic and stochastic. Our results are consistent with the results we presented in this thesis. We observe that the technique works best for discrete, deterministic systems, where we are able to observe a convergence rate faster than algebraic. However, this rate was difficult to recover in our experiments: we experienced numerical challenges when trying to recover the Koopman operator for the stochastic system or higher dimensional systems.

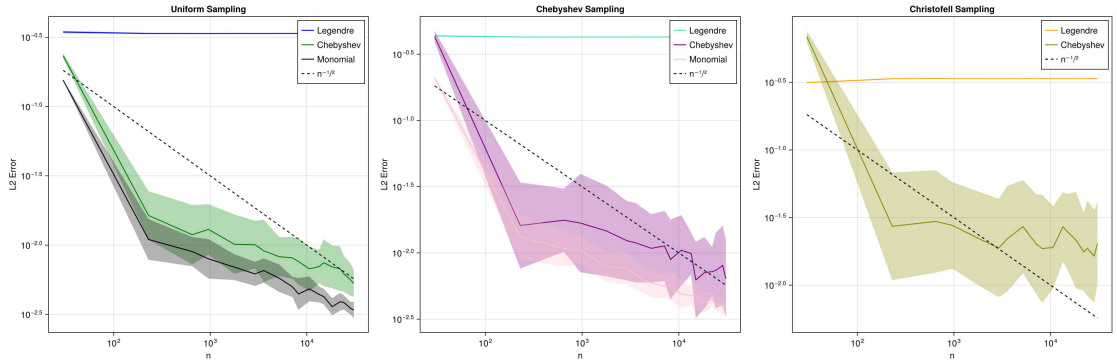

Figure 3.8: Convergence of the approximation of the Koopman operator with variable dictionary size for the stochastic logistic map. The relation is $n = 5m^2 \log(m)$

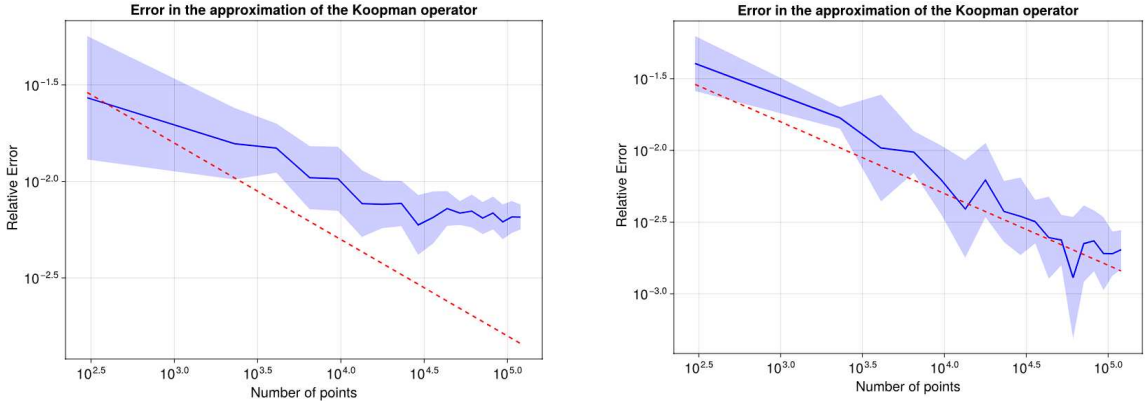

(a) No thresholding

(b) With thresholding

Figure 3.9: Exploring the impact of thresholding on the observable $\varphi(x) = x^2$ for the stochastic logistic map using Legendre polynomials and uniform sampling. The dashed line represents $\frac{1}{\sqrt{n}}$.

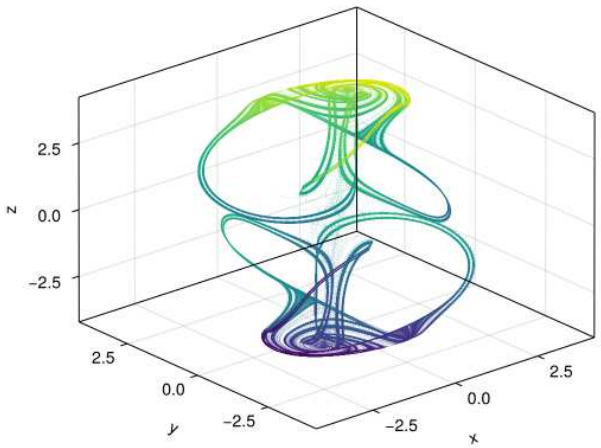

Figure 3.10: The chaotic attractor of the Thomas model.

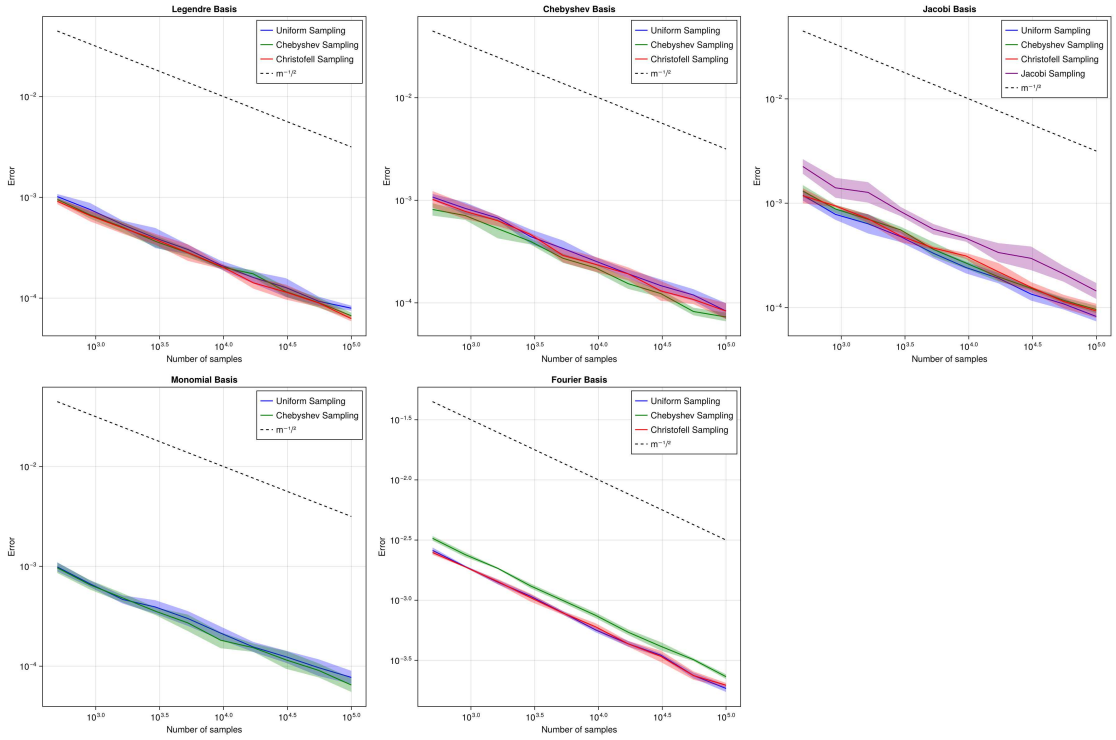

Figure 3.11: Convergence of the approximation of the Koopman operator for the Thomas model with $\psi = \boldsymbol{\phi}$.

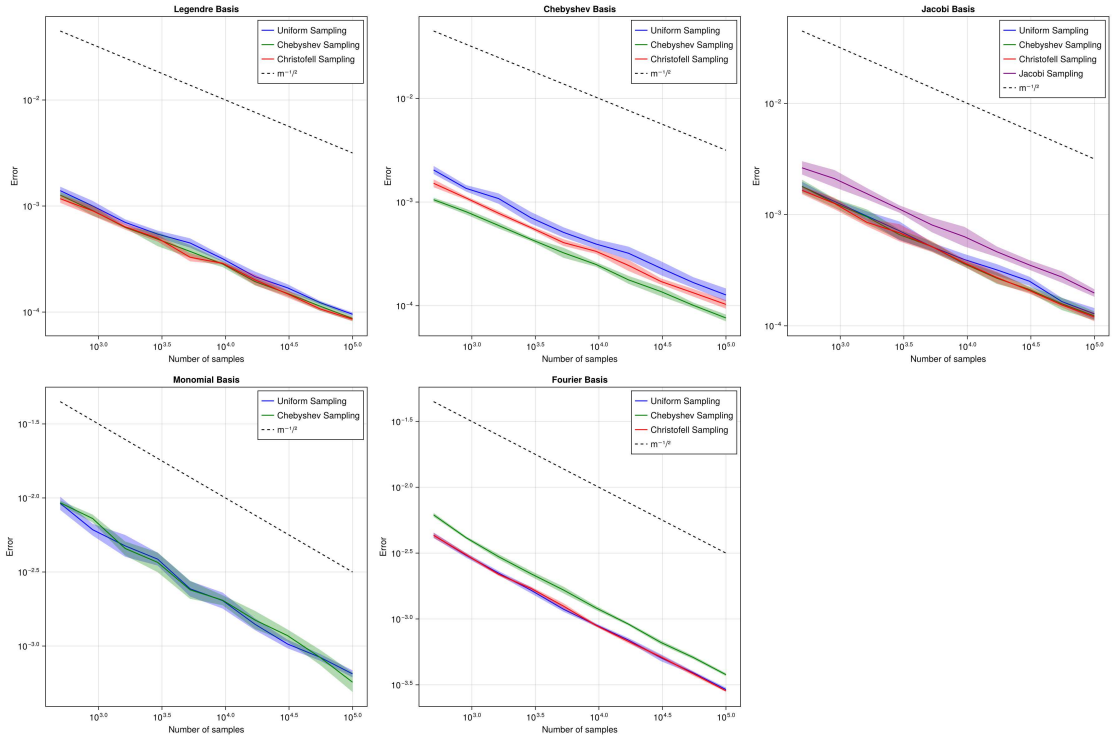

Figure 3.12: Convergence of the approximation of the Koopman operator for the Thomas model with $|\boldsymbol{\psi}| > |\boldsymbol{\phi}|$.

Chapter 4

The compressed sensing approach

In this chapter, we will explore the use of compressed sensing techniques to approximate the Koopman operator. Similarly as in the previous chapter, we will show how the task of approximating the Koopman operator can be recast as standard compressed sensing problems, and apply the known recovery results to the EDMD problem.

4.1 Sparse recovery guarantees

In this section, we explore the effectiveness of compressed sensing techniques to approximate the Koopman operator inspired by the EDMD formulation and the work presented in [32], aiming to generalize their work from model identification to the approximation of the Koopman operator, and experiment with different techniques. We will introduce the techniques utilized and derive some recovery guarantees.

4.1.1 The sparsity of the Koopman operator

Before delving into the compressed sensing techniques, there is one critical question one must ask: *Should we expect the Koopman operator to be sparse in the first place?* After all, compressed sensing hinges on the sparsity of the object we desire to reconstruct in the undersampled regime. In this subsection, we will explore the sparsity of the Koopman operator thanks to some key insights from dynamical systems.

Our motivating example is the Lorenz-96 model. This model is a d-dimensional continuous dynamical system for $d > 3$, defined as follows:

$$\frac{dx_i}{dt} = (x_{i+1} - x_{i-2})x_{i-1} - x_i + F, \quad i = 1, \dots, d, \tag{4.1}$$

where $x_{-1} = x_{d-1}, x_0 = x_d, x_{n+1} = x_1$ and $F = 8$ for chaotic behaviour. This system exhibits the two observations that lead to sparsity of the Koopman operator:

1. **Dynamical Systems are simple:** Most useful dynamical system encountered in science from physics or chemistry rely on a few simple laws governing the dynamics. For the case of the Lorenz System, the dynamics are governed by a 2nd order polynomial equation. We try to work with reduced order models that only capture the most

important dynamics of the system. Hence, in those cases, we can expect the systems to contain few nonlinear terms, leading to sparsity.

2. **High-Dimensional Dynamics are often sparser:** High-dimensional systems are often sparse. For examplem they might come from the recasting of a high-order ODE as a system of first-order ODEs. In this case, most of the state variables are only impacted by a single other state variable. A little more generally, many systems—like the Lorenz 96 system that we will be using— have the dynamics of a given state variable x_j rarely depend on the $x_i, i \in [j-k, j+l], k, l \in \mathbb{N}$. We can then expect the Koopman Operator to be sparser as dimension increases. This gives the compressed sensing approach a more competitive edge as it will scale better with dimensionality than the EDMD based techniques. The Lorenz-96 model can be arbitrarily high dimensional, so we can use it for our tests and verify this observation.

These observations are standard in the literature [32]. To make this link more explicit, in the case of the Lorenz-96 model we can see it via Taylor expansion. Consider the observable function $\phi(\mathbf{x}) = x_3$. Then

$$\begin{aligned}
\mathcal{K}^\tau[\phi](\mathbf{x}(t)) &= \phi(\mathbf{x}(t+\tau)) \\
&\approx \phi(\mathbf{x}(t)) + \tau\nabla(\mathbf{x}(t)) \\
&= x_3(t) + \tau\frac{d\mathbf{x}}{dx_3}(t) \\
&= (1-\tau)x_3(t) + \tau x_5 x_2 - \tau x_1 x_2 + \tau F.
\end{aligned}$$

Notice that in the space spanned by the monomials, the action of the Koopman operator to this observable can be described as a 4-sparse vector.

The analysis conducted above can only work if the underlying dynamics of the system are known. However, this is not always the case. Nonetheless, in practice, sparsity can still be checked through the use of EDMD since we know this method converges to the true Koopman operator via Theorem 2.1. To assess the sparsity of the Koopman operator from data, it is possible to use the EDMD method with dictionaries satisfying the conditions of Theorem 2.1 and then, after thresholding the smaller values to 0, compute the sparsity of the approximation. We can define a sparsity coefficient $k = s/(lm)$ of the approximation to get an idea of how sparse the operator is where s is the sparsity, and l, m are the number of observables functions in our dictionaries. For the Lorenz-96 model with $d = 5$, using the tensorized monomials up to degree 1 and 2 for $\boldsymbol{\phi}$ and $\boldsymbol{\psi}$ respectively, we obtain a sparsity of $s = 39$ for a total of 126 entries, yielding $k = 0.310$. For $d = 25$, we obtain $k \approx 0.0456$. The sparsity rate decreases as the dimension increases, just as we observed above. The choice of dictionaries is very important. If we instead switch to Legendre polynomials, we instead obtain $k = 0.372$ when $d = 5$ and $k = 0.103$ when $d = 25$. This is a clear indication that, in general, one can expect the Koopman operator to be at least compressible if not sparse when expanded in the appropriate basis.

4.1.2 Compressed sensing techniques applied to Koopman operator approximation

In this subsection, we present the compressed sensing problem we will study to approximate the Koopman operator, and provide a link between it and the theory of compressed sensing. We start by presenting the compressed sensing problems we will be studying.

Definition 4.1 (QCBP-EDMD and LASSO-EDMD)**.** *We define the Quadratically Constrained Basis Pursuit EDMD (QCBP-EDMD) problem as*

$$\min_{A\in\mathbb{C}^{l\times n}} \|A\|_F \quad \text{subject to} \quad \|\mathbf{\Phi}_n^\tau - A\mathbf{\Psi}_n\|_F \leq \sigma, \tag{4.2}$$

the LASSO-EDMD problem as

$$\min_{A\in\mathbb{C}^{l\times n}} \|\mathbf{\Phi}_n^\tau - A\mathbf{\Psi}_n\|_F^2 + 2\lambda\|\mathbf{A}\|_{F,1}, \tag{4.3}$$

and the Square LASSO-EDMD problem as

$$\min_{A\in\mathbb{C}^{l\times n}} \|\mathbf{\Phi}_n^\tau - A\mathbf{\Psi}_n\|_F^2 + \lambda\|\mathbf{A}\|_{F,1}^2, \tag{4.4}$$

where $\|A\|_{F,1} = \sum_{i=1}^{l}\sum_{j=1}^{m}|A_{ij}|$ *is the Frobenius 1-norm.*

For (4.3) and (4.4), there is a clear link with the theory of compressed sensing. Let's start with the LASSO problem. Note that we can rewrite it as

$$\begin{aligned}
\min_{A\in\mathbb{C}^{l\times n}} \|\mathbf{\Phi}_n^\tau - A\mathbf{\Psi}_n\|_F^2 + 2\lambda\|\mathbf{A}\|_{F,1} &= \min_{A\in\mathbb{C}^{l\times n}} \|(\mathbf{\Phi}_n^\tau)^* - (\mathbf{\Psi}_n)^* A^*\|_F^2 + 2\lambda\|\mathbf{A}\|_{F,1} \\
&= \min_{\mathbf{a_1},\dots,\mathbf{a}_l\in\mathbb{C}^n} \sum_{i=1}^{l} \|\mathbf{b}_i - \mathbf{\Psi}_n^*\mathbf{a}_i\|_2^2 + 2\lambda\sum_{i=1}^{l}\|\mathbf{a}_i\|_1 \\
&= \sum_{i=1}^{l} \min_{\mathbf{a}_i\in\mathbb{C}^n} \|\mathbf{b}_i - \mathbf{\Psi}_n^*\mathbf{a}_i\|_2^2 + 2\lambda\|\mathbf{a}_i\|_1,
\end{aligned}$$

where $\mathbf{a}_i, \mathbf{b}_i$ are the i-th columns of $A^*, (\mathbf{\Phi}_n^\tau)^*$ respectively. Hence, the LASSO-EDMD problem can be solved by working on l independent LASSO problems. For the Square LASSO problem (4.4), a similar technique can be used, but we will need the following equivalent formulation of the problem. While this specific formulation is found in [18, Chapter 6.2], a proof was not included, so we provide one for completeness.

Lemma 4.1 ([18, Chapter 6.2])**.** *The square LASSO problem*

$$\min_{\mathbf{c}\in\mathbb{C}^n} \|\mathbf{b} - A\mathbf{c}\|_2^2 + \mu^2\|\mathbf{c}\|_1^2$$

can be written as the non-negative least-squares problem

$$\min_{\mathbf{c}^+,\mathbf{c}^-\in\mathbb{C}^n} \left\| \begin{bmatrix} A & -A \\ \mu\mathbb{1}^T & \mu\mathbb{1}^T \end{bmatrix} \begin{bmatrix} \mathbf{c}^+ \\ \mathbf{c}^- \end{bmatrix} - \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix} \right\|_2^2 \text{ subject to } \mathbf{c}^+ \geq 0, \mathbf{c}^- \geq 0, \tag{4.5}$$

where $\mathbf{c} = \mathbf{c}^+ - \mathbf{c}^-$ *such that*

$$c_i^+ = \begin{cases} c_i & \text{if } c_i \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad c_i^- = \begin{cases} -c_i & \text{if } c_i < 0 \\ 0 & \text{otherwise} \end{cases}$$

and $\mathbb{1}$ *is a vector of ones of appropriate size.*

Proof. First, Let $I = \{i = 1, \ldots, n : \mathbf{c}_i \geq 0\}$ we note that $\|\mathbf{c}\|_1^2 = \sum_{i \in I} c_i^+ + \sum_{i \notin I} c_i^- = \begin{bmatrix} \mathbb{1}^T & \mathbb{1}^T \end{bmatrix} \begin{bmatrix} \mathbf{c}^+ \\ \mathbf{c}^- \end{bmatrix}$. Hence, $\mu^2 \|\mathbf{c}\|_1^2 = \left(\begin{bmatrix} \mu\mathbb{1}^T & \mu\mathbb{1}^T \end{bmatrix} \begin{bmatrix} \mathbf{c}^+ \\ \mathbf{c}^- \end{bmatrix} \right)^2$. Now we can similarly notice that $A(\mathbf{c}^+ - \mathbf{c}^-) = \begin{bmatrix} A & -A \end{bmatrix} \begin{bmatrix} \mathbf{c}^+ \\ \mathbf{c}^- \end{bmatrix}$ Putting those two observations together, we get now that

$$\begin{aligned} \|A\mathbf{c} - \mathbf{b}\|_2^2 + \mu^2 \|\mathbf{c}\|_1^2 &= \left\| \begin{bmatrix} A & -A \end{bmatrix} \begin{bmatrix} \mathbf{c}^+ \\ \mathbf{c}^- \end{bmatrix} + \begin{bmatrix} \mu\mathbb{1}^T & \mu\mathbb{1}^T \end{bmatrix} \begin{bmatrix} \mathbf{c}^+ \\ \mathbf{c}^- \end{bmatrix} - \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \end{bmatrix} \right\|_2^2 \\ &= \left\| \begin{bmatrix} A & -A \\ \mu\mathbb{1}^T & \mu\mathbb{1}^T \end{bmatrix} \begin{bmatrix} \mathbf{c}^+ \\ \mathbf{c}^- \end{bmatrix} - \begin{bmatrix} \mathbf{b} \\ 0 \end{bmatrix} \right\|_2^2. \end{aligned}$$

□

With this lemma established, one can rewrite the EDMD Square LASSO problem as the usual Square LASSO problem to use the theory of compressed sensing established previously. The EDMD Square LASSO can be rewritten as

$$\min_{A^+, A^- \in \mathbb{C}^{l \times n}} \left\| \begin{bmatrix} \mathbf{\Psi}_n^\tau & -\mathbf{\Psi}_n^\tau \\ \mu\mathbb{1}^T & \mu\mathbb{1}^T \end{bmatrix} \begin{bmatrix} \mathbf{A}^+ \\ \mathbf{A}^- \end{bmatrix} - \begin{bmatrix} \mathbf{\Phi}_n^\tau \\ \mathbf{0} \end{bmatrix} \right\|_2^2 = \sum_{i=1}^{l} \min_{\mathbf{a}_i^+, \mathbf{a}_i^- \in \mathbb{C}^n} \left\| \begin{bmatrix} \mathbf{\Psi}_n^\tau & -\mathbf{\Psi}_n^\tau \\ \mu\mathbb{1}^T & \mu\mathbb{1}^T \end{bmatrix} \begin{bmatrix} \mathbf{a}_i^+ \\ \mathbf{a}_i^- \end{bmatrix} - \begin{bmatrix} \mathbf{b}_i \\ \mathbf{0} \end{bmatrix} \right\|_2^2$$

where $\mathbf{a}_i = \mathbf{a}_i^+ - \mathbf{a}_i^-$, and $\mathbf{b}_i$ is the i-th column of $\mathbf{\Phi}_n^\tau$, which we can recognize as a collection of square LASSO problems

$$\min_{a_i \in \mathbf{C}^m} \|\mathbf{b}_i - \mathbf{\Psi}_n^\tau \mathbf{a}_i\|_2^2 + \mu^2 \|\mathbf{a}_i\|_1^2.$$

Once again, the EDMD square LASSO problem is equivalent to solving l independent Square LASSO problems (of the same form that was introduced earlier by letting $\lambda = \mu^2$). For the QCBP problem, we will consider the case when $\sigma = 0$, that is when the constraint becomes $\Phi_n^\tau = A\Psi_n$. This is a strong assumption, but from empirical results, it might be relaxed to draw a more general link between the usual QCBP problem and the QCBP-EDMD problem. If we keep this assumption for now, then we can rewrite the QCBP problem as

$$\begin{aligned} &\min_{A \in \mathbb{C}^{l \times n}} \|A\|_F^2 \text{ subject to } \mathbf{\Phi}_n^\tau = A\mathbf{\Psi}_n \\ &= \min_{A \in \mathbb{C}^{l \times n}} \|A^*\|_F^2 \text{ subject to } (\mathbf{\Phi}_n^\tau)^* = (\mathbf{\Psi}_n)^* A^* \\ &= \min_{\mathbf{a}_1, \ldots, \mathbf{a}_l \in \mathbb{C}^n} \sum_{i=1}^{l} \|\mathbf{a}_i\|_2^2 \text{ subject to } \mathbf{b}_i = \mathbf{\Psi}_n^* \mathbf{a}_i \quad \forall i = 1, \ldots l \\ &= \sum_{i=1}^{l} \min_{\mathbf{a}_i \in \mathbb{C}^n} \|\mathbf{a}_i\|_2^2 \text{ subject to } \mathbf{b}_i = \mathbf{\Psi}_n^* \mathbf{a}_i \quad \forall i = 1, \ldots l, \end{aligned}$$

where $\mathbf{a}_i, \mathbf{b}_i$ are the i-th columns of $A^*, (\mathbf{\Phi}_n^\tau)^*$ respectively. Hence, the (noiseless) EDMD-QCBP problem can be solved by solving l independent (noiseless) QCBP problems.

Remark 4.1. *The assumption that $\sigma = 0$ is a strong one, but it is powerful. Without it, the link between the QCBP-EDMD problem and the usual QCBP problem is much harder to establish, as it is not easy to show that the solution of the independent problems can be put together to obtain the solution of the QCBP-EDMD problem.*

Theorem 4.1 (Sparsity guarantee for the LASSO decoder)**.** *Let $\psi, \boldsymbol{\phi}$ be two dictionaries satisfying Assumption 2.1. Moreover, assume $\boldsymbol{\psi}$ is orthogonal so that ψ_n^τ is a measurement matrix obtained from sampling a BOS with constant B, and let $(K_{mn}^\tau)^\#$ be the solution of the EDMD-LASSO or EDMD-Square LASSO problem. Then, if*

$$n \geq Bc\delta^{-2}(4s+1)L(4s+1),$$

with $L(s)$ as in Proposition 2.7, with probability $1-\epsilon$, $\mathbf{\Psi}_n$ has the RIP property of order $4s+1$ with constant δ and the solution $(K_{mn}^\tau)^\#$ satisfies

$$\|(K_{mn}^\tau)^\#\|_0 \leq 4ls.$$

Proof. Since $\mathbf{\Psi}_n$ is a measurement matrix from a bounded orthonormal system with constant B, by Proposition 2.7 it has the RIP of order $4s+1$ and constant δ. Then since the EDMD LASSO and EDMD Square LASSO problem are equivalent to solving l independent LASSO or Square LASSO problems, we can apply Theorem 2.8 to each of the l problems. So each column of $(K_{mn}^\tau)^\#$ has sparsity at most $4s$, and

$$\|(K_{mn}^\tau)^\#\|_0 \leq 4ls.$$

This concludes the proof. □

While this theorem does not provide a bound on the reconstruction error, it does provide a guarantee in terms of the sparsity of the solution. If we work in a setting where we know than the answer should have a certain sparsity, then we know that the LASSO EDMD decoder will enforce this structure into the approximation.

For the QCBP-EDMD problem, we can derive a recovery guarantee in terms of the best s-term approximation of $\mathcal{K}_m^\tau$.

Theorem 4.2 (Recovery of QCBP-EDMD)**.** *Let $\mathbf{\Phi}_n^\tau, \psi_n$ be as in Theorem 4.1. Let $(K_{mn}^\tau)^\#$ be the solution of the QCCBP-EDMD problem. Then, if*

$$n \geq Bc\delta^{-2}(2s)L(2s),$$

the reconstruction error satisfies

$$\|K_n^\tau - (K_{mn}^\tau)^\#\|_F \leq C\sqrt{\sum_{i=1}^{l} \frac{\sigma_s(\mathbf{a}_i)_1}{\sqrt{s}}} \tag{4.6}$$

with probability $1-l\epsilon$, where $\mathbf{a}_i$ is the i-th row of K_n^τ the matrix representation of $\mathcal{K}^\tau$ and $\sigma_s(\cdot)_1$ is the best s-term approximation error.

Proof. The proof is a direct application of Theorem 2.5 to each of the l independent QCBP problems. The probability that the result holds follows from the union bound. □

This theorem provides a guarantee that our approximation of the Koopman operator will be close to the true Koopman operator in the sense that it is comparable to the best possible approximation that is s-sparse. If the original operator is, this theorem tells us that we can recover it with no error since $\sigma_s(\mathbf{a}_i) = 0$. We note that even though our recovery results for the QCBP-EDMD problem do not apply when $\sigma \neq 0$, in practice, we still consider $\sigma > 0$ for feasibility and numerical stability.

4.2 Numerical experiments

In this section, we present numerical experiments to test the recovery of the Koopman operator using compressed sensing. First we will tune the parameters in the context of noisy data $\mathbf{\Psi}_n = \{\psi_j(\boldsymbol{x}_i + \epsilon\mu)\}$, $\mathbf{\Phi}_n^\tau = \{\phi_i(\mathbf{y}_i + \epsilon\mu)\}$ where $\epsilon > 0$, and μ is drawn from the normal distribution with mean 0 and standard deviation 1. Once we have tuned the parameters, we will then evaluate the recovery of the Koopman operator in the undersampled regime. We will follow the experimental setup introduced in [32], where data is gathered along multiple trajectories to recover the dynamics, we denote the number of points per trajectories per experiment by q. In order to solve those optimization problems, we used the COSMO library in Julia with Convex.jl [21, 38].

In all of these experiments, we evaluate recovery by measuring the residuals $\|\mathbf{\Phi}_n^\tau - \mathcal{K}\mathbf{\Psi}_n\|_F^2$, we use dictionaries containing Legendre polynomials of degree up to 5 and the dimension of the Lorenz-96 system used was $d = 10$.

4.2.1 Tuning parameters

In this subsection, we look for the optimal value of the tuning parameters in the presence of noisy data. The data consist of 20 trajectories, we sample 5 points per trajectory with noise ϵ and time step $\tau = 0.01$. The number of total data points is quite small, but this allows us to explore thoroughly the parameter space to find the most optimal tuning parameter as each experience is quick. The results of these experiments are outlined in Figure 4.1. For the QCBP problem, we observe that the optimal value of σ seems to be $\sigma \approx 10^{-4}$, when low level noise is present, it seems that we should favor a small value for σ as it performs better in that regime, however, for a larger noise level (in the figure, $\epsilon = 0.01$), we need a larger value of σ in order for the problem to be feasible. For the LASSO problem, the optimal value of λ is approximately $\lambda \approx 10^{-2}$, and for the Square LASSO problem, the optimal value of λ seems to be $\lambda \approx 10^{-1}$. The LASSO problems appear to have a specific window of values for λ that yield good results. This makes sense as a value of λ too small will not enforce sparsity, while a value too large will not provide a good approximation of the Koopman operator.

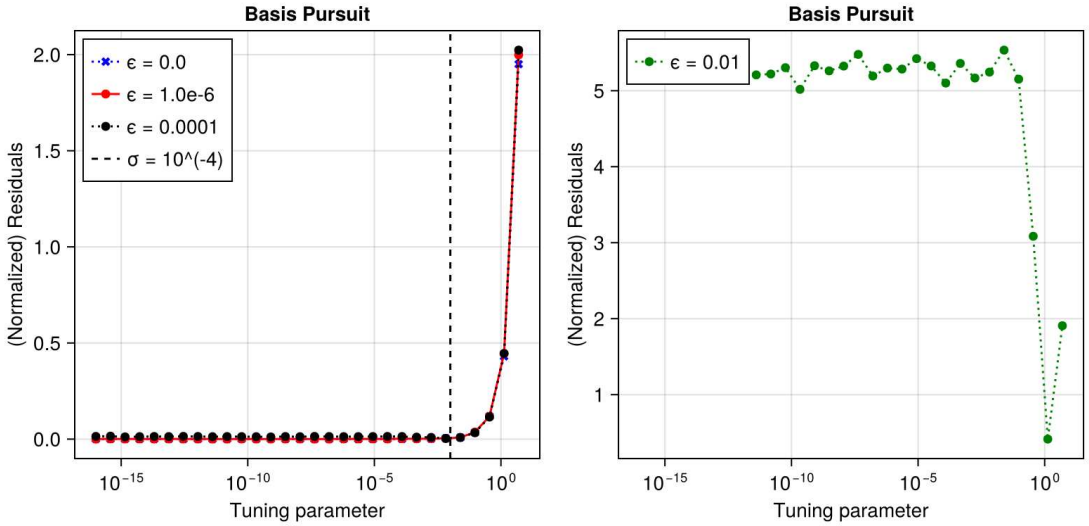

(a) Tuning parameter experiments for σ for QCBP problem

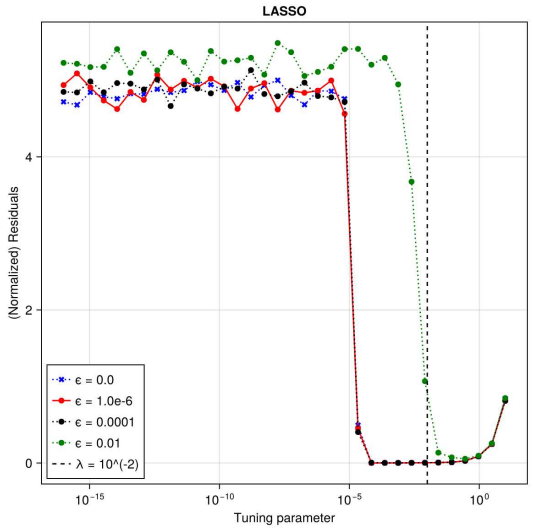

(b) Tuning parameter experiments for λ for the LASSO problem

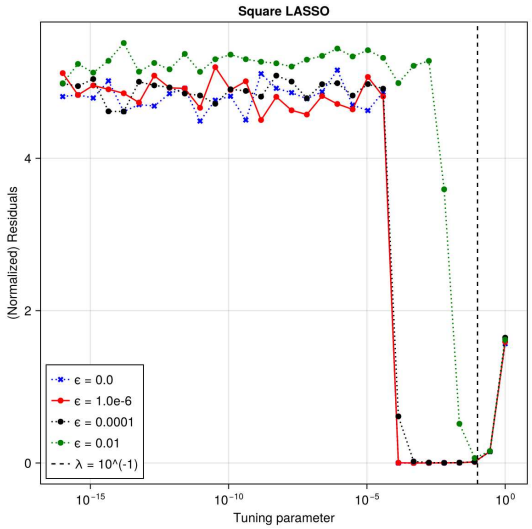

(c) Tuning parameter experiments for λ for the Square LASSO problem

Figure 4.1: Tuning parameter experiments for QCBP, LASSO, and Square LASSO. Residuals are computed with respect to a validation set, and not the training set.

4.2.2 Recovery of the Koopman operator for the Lorenz-96 model

In the previous subsection, we have tuned the hyper-parameters σ and λ, so we can use them to test the recovery of the Koopman operator. To do this, we evaluate the residuals $\|\mathbf{\Phi}_n^\tau - \mathcal{K}\mathbf{\Psi}_n\|_F^2$ as our error. The results are outlined in Figure 4.2. We note that in all cases, we seem to recover the Koopman operator. The LASSO approach is the one with the least residuals. It seems like the method from which we sample the data is quite important, as experiments with fewer bursts seem to achieve better recovery, this is probably due to the

fact that this system is chaotic, meaning that longer but fewer trajectories may explore the phase space more efficiently than smaller but more numerous ones.

To compare with the results outlined in Section 3.2, we note that this technique scales better with dimension, as EDMD is affected by the curse of dimensionality, where by construction, the compressed sensing techniques are not. We also note that the compressed sensing techniques are robust to noise in the data. For these reasons, the compressed sensing approach is a good fit in cases where data is scarce, and polluted with noise. However, for this to be effective, the assumption of sparsity is crucial. In the case of the Lorenz-96 model, we have shown that the Koopman operator is sparse, but this is not always necessary the case.

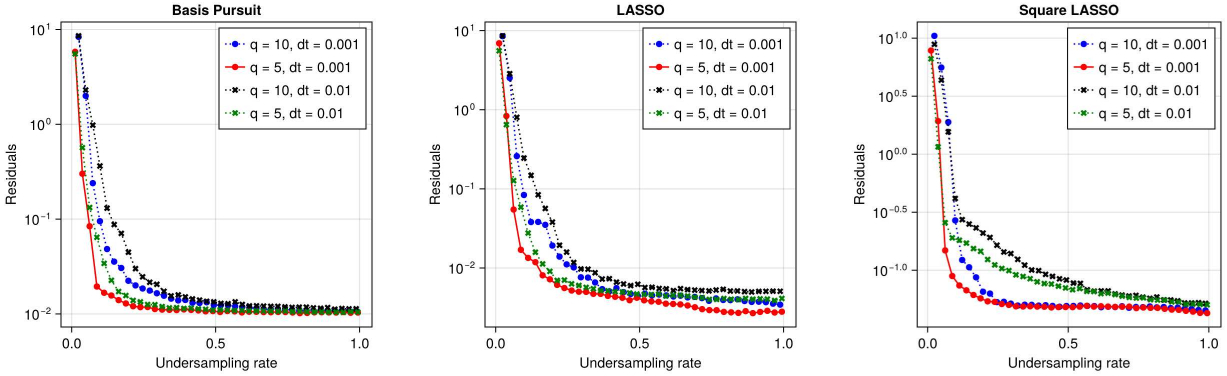

(a) Convergence of the QCBP problem for $\sigma = 10^{-4}$

(b) Convergence of the LASSO problem for $\lambda = 10^{-2}$

(c) Convergence of the Square LASSO problem for $\lambda = 10^{-1}$

Figure 4.2: Recovery experiments for QCBP, LASSO, and Square LASSO. Residuals are computed with respect to a validation set, and not the training set. The undersampling rate is defined as the ratio of number of sample points to the number of unknowns.

Chapter 5

Conclusions and future work

In this chapter, we briefly summarize the work done in this thesis, reflect on the results obtained and outline some future work that could be done to further investigate the theory of EDMD and its applications.

5.1 Conclusions

The work in this thesis has focused on bridging the gap between the theories of least squares approximation and compressed sensing and the theory of Extended Dynamic Mode Decomposition (EDMD) for the Koopman operator. We were able to provide convergence rates relying on minimal assumptions, that is, orthonormality of the expansion dictionary for a rate of $\frac{1}{\sqrt{n}}$ (Theorem 3.2). Moreover, we proved a faster-than-polynomial convergence under a holomorphy assumptions of the observables and dynamical systems considered (Theorem 3.3). Additionally, we provided theoretical recovery guarantees in the undersampled regime thanks to compressed sensing techniques such as LASSO and QCBP, if sparsity can be assumed (see Theorem 4.1 and Theorem 4.2).

Our theoretical work was complemented by a thorough numerical experimentation on chaotic dynamical systems, both deterministic and stochastic, discrete and continuous. We were able to observe the predicted rate in all those cases. We observed the importance of the choice of dictionary and sampling regime. A clever choice of sampling measure can outperform Monte Carlo sampling, which is the idea behind Christofell sampling.

5.2 Future work

In this thesis, we have focused on bridging the theoretical gap between the theory of least squares, compressed sensing and EDMD. However, our numerical experiments suggest that there is still room to investigate the performance of EDMD in the cases of stochastic systems and higher dimensional systems. We also observed that the choice of dictionary and sampling measure does not seem to have much impact on the performance of EDMD. It would be interesting to investigate how to adapt our recovery guarantees to drop the assumption of orthonormality of the dictionary $\boldsymbol{\psi}$ with perhaps a weighted EDMD approach.

In the context of compressed sensing, it could be interesting to seek further numerical validation of Theorem 4.1 on the sparsity of the approximate Koopman operator and Theorem 4.2 on the recovery guarantee for the QCBP problem on different dynamical systems.

Bibliography

[1] Ben Adcock. *Optimal sampling for least-squares approximation*. 2024. arXiv: 2409.02342 [stat.ML]. URL: https://arxiv.org/abs/2409.02342.

[2] Ben Adcock and Simone Brugiapaglia. *Monte Carlo is a good sampling strategy for polynomial approximation in high dimensions*. Nov. 5, 2023. arXiv: 2208.09045[cs, math]. URL: http://arxiv.org/abs/2208.09045 (visited on 07/09/2024).

[3] Ben Adcock, Simone Brugiapaglia, and Clayton G. Webster. *Sparse Polynomial Approximation of High-Dimensional Functions*. Computational Science & Engineering. SIAM, Jan. 2022. ISBN: 978-1-61197-687-8. DOI: 10.1137/1.9781611976885.ch5. URL: https://epubs.siam.org/doi/10.1137/1.9781611976885.ch5 (visited on 07/22/2024).

[4] Ben Adcock, Nick Dexter, and Sebastian Moraga. "Optimal approximation of infinite-dimensional holomorphic functions II: Recovery from i.i.d. pointwise samples". In: *Journal of Complexity* 89 (2025), p. 101933. ISSN: 0885-064X. DOI: https://doi.org/10.1016/j.jco.2025.101933. URL: https://www.sciencedirect.com/science/article/pii/S0885064X25000111.

[5] Ben Adcock and Anders C. Hansen. *Compressive Imaging: Structure, Sampling, Learning*. Cambridge University Press, 2021.

[6] H. M. Antia. *Numerical Methods for Scientists and Engineers*. Vol. 2. Texts and Readings in Physical Sciences. Gurgaon: Hindustan Book Agency, 2012. ISBN: 978-93-80250-40-3 978-93-86279-52-1. DOI: 10.1007/978-93-86279-52-1. URL: http://link.springer.com/10.1007/978-93-86279-52-1 (visited on 05/13/2025).

[7] Heinz H. Bauschke and Walaa M. Moursi. *An Introduction to Convexity, Optimization, and Algorithms*. SIAM, pp. 39–43. DOI: 10.1137/1.9781611977806.ch8. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611977806.ch8. URL: https://epubs.siam.org/doi/abs/10.1137/1.9781611977806.ch8.

[8] Aaron Berk, Simone Brugiapaglia, and Tim Hoheisel. "LASSO Reloaded: A Variational Analysis Perspective with Applications to Compressed Sensing". In: *SIAM Journal on Mathematics of Data Science* 5.4 (2023), pp. 1102–1129. DOI: 10.1137/22M1498991.

[9] Jason J. Bramburger. *Data-Driven Methods for Dynamic Systems*. Philadelphia, PA: SIAM, 2024. DOI: 10.1137/1.9781611978162. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611978162. URL: https://epubs.siam.org/doi/abs/10.1137/1.9781611978162.

[10] Jason J. Bramburger and Giovanni Fantuzzi. "Auxiliary Functions as Koopman Observables: Data-Driven Analysis of Dynamical Systems via Polynomial Optimization". In: *Journal of Nonlinear Science* 34.1 (Oct. 30, 2023), p. 8. ISSN: 1432-1467. DOI: 10.1007/s00332-023-09990-2. URL: https://doi.org/10.1007/s00332-023-09990-2 (visited on 07/09/2024).

[11] Jason J. Bramburger and Giovanni Fantuzzi. "Data-driven discovery of invariant measures". In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 480.2286 (2024), p. 20230627. DOI: 10.1098/rspa.2023.0627. URL: https://royalsocietypublishing.org/doi/abs/10.1098/rspa.2023.0627.

[12] Steven L. Brunton and J. Nathan Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. 2nd ed. Cambridge University Press, 2022.

[13] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. "Discovering governing equations from data by sparse identification of nonlinear dynamical systems". In: *Proceedings of the National Academy of Sciences* 113.15 (2016), pp. 3932–3937. DOI: 10.1073/pnas.1517384113. eprint: https://www.pnas.org/doi/pdf/10.1073/pnas.1517384113. URL: https://www.pnas.org/doi/abs/10.1073/pnas.1517384113.

[14] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. *Stable Signal Recovery from Incomplete and Inaccurate Measurements*. 2005. DOI: https://doi.org/10.1002/cpa.20124. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpa.20124. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.20124.

[15] Giuseppe Alessio D'Inverno, Kylian Ajavon, and Simone Brugiapaglia. *Surrogate models for diffusion on graphs via sparse polynomials*. 2025. arXiv: 2502.06595 [math.NA]. URL: https://arxiv.org/abs/2502.06595.

[16] D.L. Donoho. "Compressed sensing". In: *IEEE Transactions on Information Theory* 52.4 (2006), pp. 1289–1306. DOI: 10.1109/TIT.2006.871582.

[17] Daniel Fassler. *Finite-data Error Bounds for Approximating the Koopman Operator*. Version 1.0.0. Apr. 2025. URL: https://github.com/Falsorr/EDMD_Master.

[18] Simon Foucart. *Mathematical Pictures at a Data Science Exhibition*. Cambridge University Press, 2022.

[19] Simon Foucart. "The sparsity of LASSO-type minimizers". In: *Applied and Computational Harmonic Analysis* 62 (2023), pp. 441–452. ISSN: 1063-5203. DOI: https://doi.org/10.1016/j.acha.2022.10.004.

[20] Simon Foucart and Holger Rauhut. *An Invitation to Compressive Sensing*. New York, NY: Springer New York, 2013, pp. 1–39. ISBN: 978-0-8176-4948-7. DOI: 10.1007/978-0-8176-4948-7_1. URL: https://doi.org/10.1007/978-0-8176-4948-7_1.

[21] Michael Garstka, Mark Cannon, and Paul Goulart. "COSMO: A Conic Operator Splitting Method for Convex Conic Problems". In: *Journal of Optimization Theory and Applications* 190.3 (2021), pp. 779–810. DOI: 10.1007/s10957-021-01896-x. URL: https://doi.org/10.1007/s10957-021-01896-x.

[22] Walter Gautschi and Gabriele Inglese. "Lower bounds for the condition number of Vandermonde matrices". In: *Numerische Mathematik* 52.3 (May 1, 1987), pp. 241–250. ISSN: 0945-3245. DOI: 10.1007/BF01398878. URL: https://doi.org/10.1007/BF01398878 (visited on 05/06/2025).

[23] Ilon Joseph. "Koopman–von Neumann approach to quantum simulation of nonlinear classical dynamics". In: *Phys. Rev. Res.* 2 (4 Oct. 2020), p. 043102. DOI: 10.1103/PhysRevResearch.2.043102. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.2.043102.

[24] Bernard O. Koopman. "Hamiltonian Systems and Transformation in Hilbert Space." In: *Proc Natl Acad Sci U S A* 17(5) (1931), pp. 315–318.

[25] Milan Korda and Igor Mezić. "Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control". In: *Automatica* 93 (2018), pp. 149–160. ISSN: 0005-1098. DOI: https://doi.org/10.1016/j.automatica.2018.03.046.

[26] Ulrich Krengel. "On the speed of convergence in the ergodic theorem". In: *Monatshefte für Mathematik* 86.1 (Mar. 1, 1978), pp. 3–6. ISSN: 1436-5081. DOI: 10.1007/BF01300052. URL: https://doi.org/10.1007/BF01300052.

[27] Andrzej Lasota and Michael C. Mackey. *Chaos, Fractals, and Noise*. Red. by J. E. Marsden and L. Sirovich. Vol. 97. Applied Mathematical Sciences. New York, NY: Springer, 1994. ISBN: 978-1-4612-8723-0 978-1-4612-4286-4. DOI: 10.1007/978-1-4612-4286-4. URL: http://link.springer.com/10.1007/978-1-4612-4286-4 (visited on 05/12/2025).

[28] Igor Mezić. "Analysis of Fluid Flows via Spectral Properties of the Koopman Operator". In: *Annual Review of Fluid Mechanics* 45.Volume 45, 2013 (2013), pp. 357–378. ISSN: 1545-4479. DOI: https://doi.org/10.1146/annurev-fluid-011212-140652.

[29] B. K. Natarajan. "Sparse Approximate Solutions to Linear Systems". In: *SIAM Journal on Computing* 24.2 (1995), pp. 227–234. DOI: 10.1137/S0097539792240406. URL: https://doi.org/10.1137/S0097539792240406.

[30] Feliks Nüske, Sebastian Peitz, Friedrich Philipp, Manuel Schaller, and Karl Worthmann. "Finite-Data Error Bounds for Koopman-Based Prediction and Control". In: *Journal of Nonlinear Science* 33.1 (Nov. 23, 2022), p. 14. ISSN: 1432-1467. DOI: 10.1007/s00332-022-09862-1. URL: https://doi.org/10.1007/s00332-022-09862-1 (visited on 03/20/2025).

[31] Art B. Owen. *Monte Carlo theory, methods and examples*. https://artowen.su.domains/mc/, 2013.

[32] Hayden Schaeffer, Giang Tran, and Rachel Ward. "Extracting Sparse High-Dimensional Dynamics from Limited Data". In: *SIAM Journal on Applied Mathematics* 78.6 (2018), pp. 3279–3295. DOI: 10.1137/18M116798X.

[33] Terence Tao. *Topics in Random Matrix Theory*. Vol. 132. American Mathematical Soc., 2018. ISBN: 0-8218-7430-6.

[34] René Thomas. "Deterministic Chaos Seen in Terms of Feedback Circuits: Analysis, Synthesis, "Labyrinth Chaos"". In: *International Journal of Bifurcation and Chaos* 09.10 (1999), pp. 1889–1905. eprint: https://doi.org/10.1142/S0218127499001383.

[35] Giang Tran and Rachel Ward. "Exact Recovery of Chaotic Systems from Highly Corrupted Data". In: *Multiscale Modeling & Simulation* 15.3 (2017), pp. 1108–1129. DOI: 10.1137/16M1086637.

[36] Joel A. Tropp. "User-Friendly Tail Bounds for Sums of Random Matrices". In: *Foundations of Computational Mathematics* 12.4 (2012), pp. 389–434. ISSN: 16153375. DOI: 10.1007/s10208-011-9099-z. (Visited on 05/13/2025).

[37] Jonathan H. Tu, Clarence W. Rowley, Dirk M. Luchtenburg, Steven L. Brunton, and J. Nathan Kutz. "On dynamic mode decomposition: Theory and applications". In: *Journal of Computational Dynamics* 1.2 (2014), pp. 391–421. ISSN: 2158-2491. DOI: 10.3934/jcd.2014.1.391. URL: https://www.aimsciences.org/article/id/1dfebc20-876d-4da7-8034-7cd3c7ae1161.

[38] Madeleine Udell, Karanveer Mohan, David Zeng, Jenny Hong, Steven Diamond, and Stephen Boyd. "Convex Optimization in Julia". In: *SC14 Workshop on High Performance Technical Computing in Dynamic Languages* (2014). arXiv: 1410.4821 [math-oc].

[39] Caroline Wormell. "Orthogonal Polynomial Approximation and Extended Dynamic Mode Decomposition in Chaos". In: *SIAM Journal on Numerical Analysis* 63.1 (2025), pp. 122–148. DOI: 10.1137/23M1597873. URL: https://doi.org/10.1137/23M1597873.

[40] Rishikesh Yadav and Alexandre Mauroy. *Approximation of the Koopman operator via Bernstein polynomials*. 2025. arXiv: 2403.02438 [math.DS]. URL: https://arxiv.org/abs/2403.02438.