Automatica 177 (2025) 112302

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Linear quadratic control of nonlinear systems with Koopman operator learning and the Nyström method ☆

Edoardo Caldarelli [a,e,*,1], Antoine Chatalic [c,f], Adrià Colomé [a], Cesare Molinari [c], Carlos Ocampo-Martinez [a,b], Carme Torras [a], Lorenzo Rosasco [c,d,e]

[a] Institut de Robòtica i Informàtica Industrial, CSIC – UPC, Barcelona, Spain
[b] Automatic Control Department (ESAII), Universitat Politècnica de Catalunya - BarcelonaTECH, Spain
[c] MaLGa Center – DIBRIS – Università di Genova, Genoa, Italy
[d] CBMM – Massachusets Institute of Technology, Cambridge, MA, USA
[e] Istituto Italiano di Tecnologia, Genoa, Italy
[f] CNRS, Univ. Grenoble-Alpes, GIPSA-lab, France

ARTICLE INFO

Article history:
Received 5 March 2024
Received in revised form 27 November 2024
Accepted 5 March 2025
Available online 10 April 2025

Keywords:
Koopman operator
Kernel methods
Nyström method
Linear quadratic regulator
Data-driven methods

ABSTRACT

In this paper, we study how the Koopman operator framework can be combined with kernel methods to effectively control nonlinear dynamical systems. While kernel methods have typically large computational requirements, we show how random subspaces (Nyström approximation) can be used to achieve huge computational savings while preserving accuracy. Our main technical contribution is deriving theoretical guarantees on the effect of the Nyström approximation. More precisely, we study the linear quadratic regulator problem, showing that the approximated Riccati operator converges at the rate $m^{-1/2}$, and the regulator objective, for the associated solution of the optimal control problem, converges at the rate m^{-1}, where m is the random subspace size. Theoretical findings are complemented by numerical experiments corroborating our results.

1. Introduction

Nonlinear dynamical systems are ubiquitous, and pose a great challenge in terms of system identification and control. A powerful approach to deal with nonlinear dynamical systems is provided by the *Koopman operator* framework (Bevanda, Sosnowski, & Hirche, 2021; Brunton, Budišić, Kaiser, & Kutz, 2022; Koopman, 1931; Mezić, 2021). In this approach, the nonlinear dynamical system is transformed through a set of nonlinear functions, called *observables*, so that the dynamics of the transformed states are *linear*, and can be used to reconstruct the states of the system. The Koopman operator, at the basis of this technique, was originally introduced in the context of autonomous dynamical systems (Koopman, 1931). However, as shown in the seminal work of Korda and Mezić (2018), the Koopman operator formulation allows to apply linear control techniques, which are well understood and efficient to compute. As discussed by Brunton et al. (2022) and Otto and Rowley (2021), the control input can be included in the Koopman framework by either defining a family of Koopman operators (one for each value of the control input) (Nüske, Peitz, Philipp, Schaller, & Worthmann, 2023; Peitz & Klus, 2019), or by directly extending the state space to include the control input as an additional state of the system. The latter perspective is the one we consider in this paper. A key challenge to apply the Koopman operator approach is to choose the suitable space of observable functions, both with and without considering the control input (Korda & Mezić, 2018). Common choices include splines (Korda & Mezić, 2018), polynomial or Fourier bases (Abraham, de la Torre, & Murphey, 2017), and neural networks (Hao, Huang, Pan, Wu, & Mou, 2024; Shi & Meng, 2022; Yin, Welle, & Kragic, 2022), see also Gibson, Calvisi, and Yee (2022), Kaiser, Kutz, and Brunton (2021) and Korda and Mezić (2020).

In this paper, we consider observables in a *reproducing kernel Hilbert space* (RKHS) (Aronszajn, 1950). Band limited functions, splines and Sobolev spaces are then special cases (Berlinet & Thomas-Agnan, 2011). This choice was already mentioned in Korda and Mezić (2018) and has been recently analyzed in Bevanda et al. (2024), Das and Giannakis (2020), Giannakis, Henriksen, Tropp, and Ward (2023), Khosravi (2023), Klus, Nüske, and

☆ The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Matthias A. Muller under the direction of Editor Alessandro Chiuso.

* Corresponding author.
E-mail addresses: edoardo.caldarelli@iit.it (E. Caldarelli), antoine.chatalic@cnrs.fr (A. Chatalic), acolome@iri.upc.edu (A. Colomé), molinari@dima.unige.it (C. Molinari), carlos.ocampo@upc.edu (C. Ocampo-Martinez), torras@iri.upc.edu (C. Torras), lrosasco@mit.edu (L. Rosasco).

[1] This work was done while the corresponding author was at: Institut de Robòtica i Informàtica Industrial, CSIC – UPC, Barcelona, Spain.

https://doi.org/10.1016/j.automatica.2025.112302

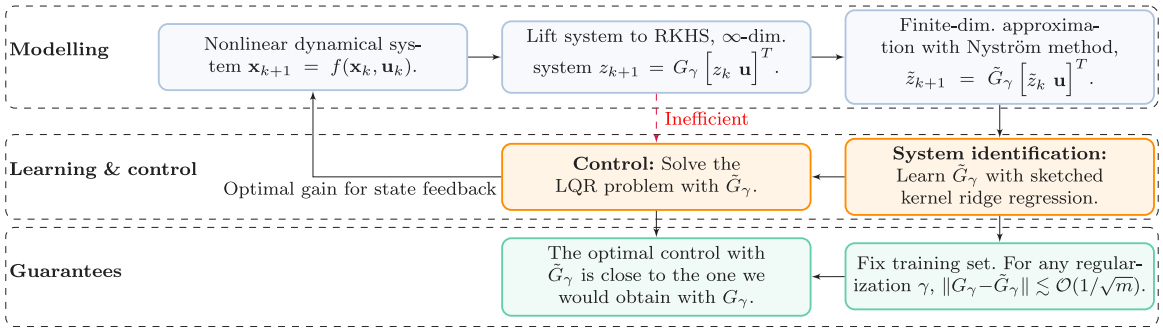

Fig. 1. Summary: given some controls and corresponding state trajectories of a nonlinear dynamical system, we use kernels to build a linear, data-driven model of the system. Kernels yield a computationally inefficient representation of the state space, due to the inversion of the kernel matrix, which we render computationally tractable using the Nyström method.

Hamzi (2020), Kostic et al. (2022) and Philipp, Schaller, Worthmann, Peitz, and Nüske (2024) for the forecasting and analysis of dynamical systems. Kernel methods are popular in machine learning (Schölkopf & Smola, 2002), as an RKHS has the advantage of being a possibly infinite dimensional space and corresponds to universal approximators (Steinwart & Christmann, 2008), while the associated estimators are nonparametric (Wasserman, 2006) and their computation reduces to finite-dimensional numerical problems (Schölkopf, Herbrich, & Smola, 2001). When efficiency is needed, further approximations are however required. In this paper, we use the so-called Nyström method, which can be interpreted as a dimensionality reduction technique. More precisely, we approximate the dynamics of the observable functions in the RKHS using projections on random and data-dependent finite-dimensional subspaces of functions (Nyström, 1930; Williams & Seeger, 2000). The effect of this approximation has been characterized for supervised machine learning, see, e.g., Musco and Musco (2017) and Rudi, Camoriano, and Rosasco (2015) and references therein, and more recently for dynamical system identification, see DeGennaro and Urban (2019) and Meanti, Chatalic, Kostic, Novelli, Pontil, and Rosasco (2023). Besides the Nyström method, as an alternative kernel approximation, Nüske and Klus (2023) study the case of random Fourier features for dynamical system identification.

In this paper, we combine the above ideas and develop an efficient and accurate control approach for nonlinear dynamical systems, based on using the Koopman operator framework together with kernels and the Nyström method. Our main technical contribution is the analysis of the approximation due to the Nyström method but, unlike previous works, we consider a control setting. We focus on the linear quadratic regulator (LQR) method, which has appealing theoretical properties, such as an analytic form of the optimal solution, and easily allows to deal with multi-input-multi-output systems. Our main result studies the impact of the Nyström approximation on the optimal control problem. Namely, in a fixed design analysis, we quantify the error introduced by the Nyström approximation to an empirical, kernel-based estimator of the nonlinear dynamics. We further show how this error propagates to the solution of the optimal control problem, if such an empirical estimator is employed in an LQR. Instrumental to our analysis are recent results on learning and controlling *linear* systems (Dean, Mania, Matni, Recht, & Tu, 2020; Mania, Tu, & Recht, 2019; Simchowitz, Mania, Tu, Jordan, & Recht, 2018). We complement our theoretical analysis with some numerical experiments. The combination of the Koopman approach and the LQR (Moyalan, Choi, Chen, & Vaidya, 2023) has been applied to challenging robot learning problems (Abraham & Murphey, 2019; Yin et al., 2022), also involving soft robots (Haggerty et al., 2023). Here, we assess the proposed control pipeline on a classic control benchmark, i.e., the Duffing oscillator (Korda & Mezić, 2020), and on the compelling problem of identifying the dynamics of cloth (Amadio, Delgado-Guerrero, Colomé, & Torras, 2023; Coltraro, Amorós, Alberich-Carramiñana, & Torras, 2022; Luque, Parent, Colomé, Ocampo-Martinez, & Torras, 2024; Zheng, Colomé, Sentis, & Torras, 2022). To conclude, we observe that concurrent to our work, Driessen (Driessen, 2023) also considers kernel-based methods and their Nyström approximation for control purposes, but without any theoretical guarantees, and learning a bilinear (Korda & Mezić, 2018; Moyalan et al., 2023) data-driven system.

As shown in Fig. 1, our contributions are the following:

- we show how the Koopman operator framework for controlled dynamical systems can be used to design linear control by lifting the state-space representation to an RKHS;
- we show how the Nyström method can be exploited to derive efficient computations;
- we prove finite sample bounds for the convergence of the Nyström data-driven system to the infinite dimensional state-space representation in RKHS;
- when using an LQR, we show how these error rates due to the Nyström approximation translate to the associated Riccati operator, and to the optimal control sequence computed by solving the LQR;
- we provide a publicly available and open-source implementation of the proposed Nyström-based system identification and optimal control algorithms[2] that we test on some illustrative examples.

The structure of this paper is as follows. In Section 2, we report the main technical background of our work. In Section 3, we discuss how to combine reproducing kernels and the Nyström method to obtain linear predictors for nonlinear systems. In Section 4 we discuss how to use such a predictor in an LQR. In Section 5, we present the core theoretical contributions of our work, while Section 6 contains a numerical evaluation of our identification algorithm and the associated LQR problem.

2. Background and notation

In this paper, we consider a nonlinear controlled discrete dynamical system, which is approximated for control purposes by a surrogate dynamical system whose dynamics are linear both in the lifted state and in the input variable. This type of heuristic approximation, previously considered for instance by Korda and Mezić (2018), has been empirically shown to be suited to nonlinear systems that are *affine* in the control input, and have no coupled terms between state and input variables. While this may be seen as a restriction of the class of dynamical systems that

[2] Available at https://github.com/LCSL/nys-koop-lqr.

can be modeled with our approximation, we can observe that it is already a fairly general model, see, e.g., the dynamical systems we study in Section 6, or the soft robot studied in Haggerty et al. (2023).

We denote the state vector as $\mathbf{x} \in \mathbb{R}^d$, the output as $\mathbf{y} \in \mathbb{R}^{n_y}$, and the control input vector as $\mathbf{u} \in \mathbb{R}^{n_u}$. For a sequence of control inputs $(\mathbf{u}_t)_{t\in\mathbb{N}}$, an initial state $\mathbf{x}_0$, and a function $f : \mathbb{R}^{d+n_u} \to \mathbb{R}^d$, we are interested in studying deterministic nonlinear dynamical systems described by the difference equation

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t), \tag{1}$$

where $t \in \mathbb{N}$ denotes the discrete-time instant. As mentioned before, the nonlinear map is considered to be affine in the control variable, and without coupled terms, i.e., for $g : \mathbb{R}^d \to \mathbb{R}^d$, $\mathfrak{B} : \mathbb{R}^{n_u} \to \mathbb{R}^d$, $f(\mathbf{x}, \mathbf{u}) = g(\mathbf{x}) + \mathfrak{B}\mathbf{u}$. The state and the input can be merged in an augmented *finite-dimensional* state, denoted by $\mathbf{w} = [\mathbf{x}^T, \mathbf{u}^T]^T \in \mathbb{R}^{d+n_u}$. We will use both notations interchangeably throughout the article. For a suitable space $\mathcal{A}$ of observable functions ξ, the Koopman operator $\mathcal{K} : \mathcal{A} \to \mathcal{A}$ can be defined as

$$(\mathcal{K}\xi)(\mathbf{w}_t) = \xi(\mathbf{w}_{t+1}), \quad \forall \xi \in \mathcal{A}. \tag{2}$$

The definition in (2) follows the Koopman with inputs and control (KIC) formalism (Proctor, Brunton, & Kutz, 2018). Note that when the control is prescribed by a state-feedback gain, this definition is related to the standard definition of the Koopman operator for the associated autonomous dynamics. An alternative definition of the Koopman operator involving sequences of controls propagated using the shift operator has also been proposed (Korda & Mezić, 2018). However, all definitions of the Koopman operator with control suffer from a common drawback, which is that they implicitly assume that a space of observable $\mathcal{A}$ invariant under the action of the Koopman operator indeed exists.

As we will discuss, a fundamental component of our learning framework is the notion of RKHS (Aronszajn, 1950). For an input space $\mathcal{X}$, an RKHS H is a Hilbert space of scalar functions on $\mathcal{X}$ for which there exists a $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$, the *reproducing kernel*, so that, for any $\chi \in \mathcal{X}$ and $f \in H$, it holds $k(\chi, \cdot) \in H$ and $f(\chi) = \langle f, k(\chi, \cdot)\rangle_H$. The latter notation denotes the inner product in the RKHS. An RKHS is a potentially infinite-dimensional space with universality properties. Some examples of reproducing kernels are given by the Gaussian kernel, or by the Matérn family of kernels, which can be linked to Sobolev spaces. Note that the Koopman operator is linear in the observable ξ. It can thus be seen as a convenient way to build a linear approximation of the original dynamical system, at the cost of manipulating the lifted state $\xi(\mathbf{w})$. In this paper, we take inspiration from the Koopman formalism in order to devise a data-driven approximation of the dynamics in the RKHS suitable to use with an LQR. We focus on the error induced by the use of the Nyström approximation, however we do not relate the learned dynamics in the RKHS to the definition of the Koopman operator with control in (2).

Notations. In the following, we use the notation $\|\cdot\|$ for the operator norm. We denote the Euclidean norm of a vector as $\|\cdot\|_2$. For any Hilbert space $\mathbb{H}$ we denote HS($\mathbb{H}$) the Hilbert space of Hilbert–Schmidt operators on $\mathbb{H}$ and $\|\cdot\|_{\text{HS}}$ the associated Hilbert–Schmidt norm. We denote A^* the adjoint of an operator A, and $A^\dagger$ the pseudo-inverse of A. Besides, $\sigma_{\min}(A)$ denotes the smallest singular value of A.[3]

3. Koopman system identification

As discussed previously, we focus on approximations of the original dynamical system which are linear in the control input. In the following, we will show how to model such a linearity in the control input. We will also detail the regression problems that are solved in order to retrieve the data-driven, kernel-based dynamics, and show how such kernel-based dynamics can be transformed to vector-valued dynamics by leveraging the Nyström method.

3.1. Choosing the Koopman lifting function

In order to define a suitable state and control transformation, we consider the RKHS $\mathcal{H}_1$ associated to a stationary positive definite kernel $k : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$. We let $\psi : \mathbb{R}^d \to \mathcal{H}_1$, $\psi(x) := k(x, \cdot)$. We define $\mathcal{H} := \mathcal{H}_1 \times \mathbb{R}^{n_u}$ and choose as our state and input transformation

$$\phi : \mathbb{R}^{d+n_u} \to \mathcal{H}, \ \phi(\mathbf{w}) := \begin{bmatrix} \psi(\mathbf{x}) \\ \mathbf{u} \end{bmatrix}. \tag{3}$$

The choice of the stationary kernel allows to lift the data through an infinite-dimensional nonlinear transformation. Moreover, ϕ is linear in the control input variable, which allows to use linear control techniques on the system of interest. However, ϕ is infinite-dimensional for many standard choices of kernel functions. We will show in Section 3.3 that a related finite-dimensional lifting function can be obtained by using a Nyström approximation of the kernel k.

3.2. Regression problem and corresponding solutions

We now explain how the function ϕ introduced in (3) can be leveraged to obtain a linear, data-driven surrogate dynamical system to be used in place of the original nonlinear one when designing the control law. The difference equation of this data-driven model can be estimated by least-squares regression, as we detail in the following.

Having a dataset of n training pairs[4] $((\mathbf{w}_i, \mathbf{w}_{i+1}))_{i=1,\dots,n}$ including control inputs and corresponding states, i.e. $\mathbf{w}_i := [\mathbf{x}_i^T, \mathbf{u}_i^T]^T$, we define the sampling operators for the system's state ($S_\mathbf{x}$) and control input ($S_\mathbf{u}$) as

$$S_\mathbf{x} : \mathcal{H}_1 \to \mathbb{R}^n, \ S_\mathbf{x} l := \tfrac{1}{\sqrt{n}}[l(\mathbf{x}_1), \dots, l(\mathbf{x}_n)]^T, \tag{4}$$

$$S_\mathbf{u} : \mathbb{R}^{n_u} \to \mathbb{R}^n, \ S_\mathbf{u}\mathbf{u} := \tfrac{1}{\sqrt{n}}[\langle \mathbf{u}_1, \mathbf{u}\rangle, \dots, \langle \mathbf{u}_n, \mathbf{u}\rangle]^T. \tag{5}$$

The operator $S_\mathbf{x}$ returns a renormalized vector of the evaluations of its input l at the points $(\mathbf{x}_i)_{1\leqslant i\leqslant n}$, while the operator $S_\mathbf{u}$ corresponds to sampling the linear function $\langle \mathbf{u}, \cdot\rangle$ associated to its input at the locations $(\mathbf{u}_i)_{1\leqslant i\leqslant n}$. Note that with these definitions, we can define the following compound sampling operator:

$$S : \mathcal{H}_1 \times \mathbb{R}^{n_u} \to \mathbb{R}^n, \ S\begin{bmatrix} l \\ \mathbf{u} \end{bmatrix} = S_\mathbf{x} l + S_\mathbf{u}\mathbf{u}. \tag{6}$$

Moreover, as discussed by Korda and Mezić (2018), when considering Koopman operator regression for controlled systems, the regression output can be limited to be the one-step-ahead state, i.e., we are not interested in forecasting the evolution of the control input variable. Thus, the sampling operator for the output training points $(\mathbf{x}_2, \dots, \mathbf{x}_{n+1})$ can be written as

$$Z_\mathbf{x} : \mathcal{H}_1 \to \mathbb{R}^n, \ Z_\mathbf{x} l = \tfrac{1}{\sqrt{n}}[l(\mathbf{x}_2), \dots, l(\mathbf{x}_{n+1})]^T. \tag{7}$$

[3] A comprehensive summary of the notation used in the article is available at https://github.com/LCSL/nys-koop-lqr/tree/main/proofs.

[4] We use this notation for simplicity, however in practice one could also put together samples obtained by sampling multiple different trajectories.

Empirical risk minimization. The dynamics can be estimated by solving the following problem over the set of linear operators from $\mathcal{H}$ to $\mathcal{H}_1$:

$$G_\gamma := \arg\min_{W:\mathcal{H}\to\mathcal{H}_1} \mathcal{R}(W) + \gamma\|W\|_{\mathrm{HS}}^2 \tag{8}$$

$$\text{where} \quad \mathcal{R}(W) := \frac{1}{n}\sum_{i=1}^{n} \|\psi(\mathbf{x}_{i+1}) - W\phi(\mathbf{w}_i)\|_{\mathcal{H}_1}^2. \tag{9}$$

Here $\gamma > 0$ is a regularization parameter, and we recall that ϕ is defined in (3). Note that this regression problem is not strictly speaking a Koopman regression problem, as the regression input and output spaces are different. The objective in (8) is continuous, coercive and strictly convex, and thus admits a unique minimizer. As shown in Appendix A, the risk function can be rewritten as $\mathcal{R}(W) = \|Z_{\mathbf{x}} - SW^*\|_{\mathrm{HS}}^2$ and thus the solution of (8) can be expressed as

$$G_\gamma = Z_{\mathbf{x}}^*(SS^* + \gamma I)^{-1}S. \tag{10}$$

Although we started the exposition with the definition of the Koopman operator for simplicity, the operator G_γ is more similar to a regularized conditional mean embedding (Muandet, Fukumizu, Sriperumbudur, & Schölkopf, 2017) or an embedded Perron–Frobenius operator (Klus, Schuster, & Muandet, 2020).

For any initial state $\mathbf{x}_0 \in \mathbb{R}^d$, this operator defines the following linear dynamics in $\mathcal{H}_1$:

$$\begin{cases} z_0 = \psi(\mathbf{x}_0), & \text{(a)} \\ z_{t+1} = G_\gamma \begin{bmatrix} z_t \\ \mathbf{u}_t \end{bmatrix}. & \text{(b)} \end{cases} \tag{11}$$

Affine dynamics. Note that the operator $G_\gamma : \mathcal{H} \to \mathcal{H}_1$ defined in (8) can be decomposed in two operators $A_\gamma : \mathcal{H}_1 \to \mathcal{H}_1$ and $B_\gamma : \mathbb{R}^{n_u} \to \mathcal{H}_1$, controlling respectively the parts of the dynamics due to the state and to the control input. More precisely, defining

$$A_\gamma = Z_{\mathbf{x}}^*(SS^* + \gamma I)^{-1}S_{\mathbf{x}}, \tag{12}$$

$$B_\gamma = Z_{\mathbf{x}}^*(SS^* + \gamma I)^{-1}S_{\mathbf{u}}, \tag{13}$$

the dynamics (11(b)) are equivalent to the autoregressive linear model

$$z_{t+1} = A_\gamma z_t + B_\gamma \mathbf{u}_t. \tag{14}$$

3.3. Nyström approximation

Given that the lifted state z is typically infinite-dimensional, we are now interested in designing a finite-dimensional approximation of the dynamics in (11(b)), which would be more useful for practical control purposes. In this section, we thus approximate the nonlinear kernel k using a Nyström approximation (Williams & Seeger, 2000). The approximation is based on the choice of two sets $\tilde{\mathbf{x}}_1^{\mathrm{in}}, \dots, \tilde{\mathbf{x}}_m^{\mathrm{in}}$ and $\tilde{\mathbf{x}}_1^{\mathrm{out}}, \dots, \tilde{\mathbf{x}}_m^{\mathrm{out}}$ of m points, called the input and output landmarks. Multiple approaches have been studied for landmark selection, but in this paper we consider the simple setting where the landmarks are either all drawn uniformly from the dataset, or the input landmarks are drawn uniformly from the dataset and the output landmarks are then taken one step ahead in time w.r.t. the input ones. Now, let us define $\tilde{\mathcal{H}}_{\mathrm{in}} := \mathrm{span}\{\psi(\tilde{\mathbf{x}}_1^{\mathrm{in}}), \dots, \psi(\tilde{\mathbf{x}}_m^{\mathrm{in}})\}$, $\tilde{\mathcal{H}}_{\mathrm{out}} := \mathrm{span}\{\psi(\tilde{\mathbf{x}}_1^{\mathrm{out}}), \dots, \psi(\tilde{\mathbf{x}}_m^{\mathrm{out}})\}$, and let $\Pi_{\mathrm{in},\mathbf{x}} : \mathcal{H}_1 \to \mathcal{H}_1, \Pi_{\mathrm{out},\mathbf{x}} : \mathcal{H}_1 \to \mathcal{H}_1$ be the orthogonal projectors, respectively onto $\tilde{\mathcal{H}}_{\mathrm{in}}$ and $\tilde{\mathcal{H}}_{\mathrm{out}}$. Moreover, we define $\Pi_{\mathrm{in}} : \mathcal{H} \to \mathcal{H}$ as $\Pi_{\mathrm{in}}\phi(\mathbf{w}) = \begin{bmatrix} \Pi_{\mathrm{in},\mathbf{x}}\psi(\mathbf{x}) \\ \mathbf{u} \end{bmatrix}$; as the control variable is already finite-dimensional, we indeed only need to project the lifted state in order to obtain a finite-dimensional approximation. Following Meanti et al. (2023), we define the Nyström approximation of G_γ as follows:

$$\begin{aligned} \tilde{G}_\gamma &:= \arg\min_{W:\mathcal{H}\to\mathcal{H}_1} \mathcal{R}(\Pi_{\mathrm{out},\mathbf{x}} W \Pi_{\mathrm{in}}) + \gamma\|W\|_{\mathrm{HS}}^2 \\ &= \Pi_{\mathrm{out},\mathbf{x}} Z_{\mathbf{x}}^*(S\Pi_{\mathrm{in}}S^* + \gamma I)^{-1} S\Pi_{\mathrm{in}}. \end{aligned} \tag{15}$$

The operator $\tilde{G}_\gamma$ can be used to define linear dynamics approximating the ones from (14), namely for any initial condition $\mathbf{x}_0 \in \mathbb{R}^d$:

$$\begin{cases} \tilde{z}_0 = \Pi_{\mathrm{out},\mathbf{x}}\psi(\mathbf{x}_0), & \text{(a)} \\ \tilde{z}_{t+1} = \tilde{G}_\gamma \begin{bmatrix} \tilde{z}_t \\ \mathbf{u}_t \end{bmatrix}. & \text{(b)} \end{cases} \tag{16}$$

The projection in the initial condition guarantees that all the states visited during the evolution of the system belong to $\tilde{\mathcal{H}}_{\mathrm{out}}$. The dynamics in (16(b)) are still linear in the control input, and the operator $\tilde{G}_\gamma$ could be decomposed into operators $(\tilde{A}_\gamma, \tilde{B}_\gamma)$, approximating the operators (A_γ, B_γ) defined in (12) and (13). We choose however to manipulate only $\tilde{G}_\gamma$ in the following to keep expressions more concise.

Vector-valued representation. The dynamics in (16(b)) are defined for evolving functions: despite being all restricted to a finite-dimensional subspace, the iterates $(\tilde{z}_t)_{t\in\mathbb{N}}$ still belong to a functional space. In order to retrieve a vector-valued state representation that is practically computable, we will thus look at the dynamics of the coordinates of these evolving functions in a basis of the subspace to which they belong. More precisely, we can define $\tilde{S}_{\mathrm{out}}$ as

$$\tilde{S}_{\mathrm{out}} : \mathcal{H}_1 \to \mathbb{R}^m, \ \tilde{S}_{\mathrm{out}}g = [g(\tilde{\mathbf{x}}_1^{\mathrm{out}}), \dots, g(\tilde{\mathbf{x}}_m^{\mathrm{out}})]^T. \tag{17}$$

Note that $K_{m,\mathrm{out}} := \tilde{S}_{\mathrm{out}}\tilde{S}_{\mathrm{out}}^* \in \mathbb{R}^{m\times m}$ is the Gram matrix of the stationary kernel k computed at the output Nyström centers $\tilde{\mathbf{x}}_i^{\mathrm{out}}$. Defining $U_{\mathrm{out},\mathbf{x}} = \tilde{S}_{\mathrm{out}}^*(K_{m,\mathrm{out}}^\dagger)^{1/2}$, it holds $U_{\mathrm{out},\mathbf{x}}U_{\mathrm{out},\mathbf{x}}^* = \Pi_{\mathrm{out},\mathbf{x}}$.

Note that for any initial condition of the form (16(a)), the Nyström states $\tilde{z}_t$ defined by (16(b)) belong to $\tilde{\mathcal{H}}_{\mathrm{out}}$ for all $t \in \mathbb{N}$. Thus, define $\tilde{\mathbf{z}}_t$ as the solution of the following finite-dimensional autoregressive dynamics

$$\begin{cases} \tilde{\mathbf{z}}_0 := U_{\mathrm{out},\mathbf{x}}^*\psi(\mathbf{x}_0), & \text{(a)} \\ \tilde{\mathbf{z}}_{t+1} = U_{\mathrm{out},\mathbf{x}}^* Z_{\mathbf{x}}^*(S\Pi_{\mathrm{in}}S^* + \gamma I)^{-1} S\Pi_{\mathrm{in}} \begin{bmatrix} U_{\mathrm{out},\mathbf{x}}\tilde{\mathbf{z}}_t \\ \mathbf{u}_t \end{bmatrix}. & \text{(b)} \end{cases} \tag{18}$$

Then, for

$$\tilde{z}_t := U_{\mathrm{out},\mathbf{x}}\tilde{\mathbf{z}}_t, \tag{19}$$

$\tilde{z}_t$ fulfills (16(a)) and (16(b)).

As we detail in Appendix D, the dynamics (18(b)) can be rewritten in terms of matrix products that can be computed efficiently, and only requires to invert an $m \times m$ matrix. Note that the dynamics of $\tilde{z}$ could similarly be expressed in any orthonormal basis of $\tilde{\mathcal{H}}_{\mathrm{out}}$, however working with $U_{\mathrm{out},\mathbf{x}}$ naturally yields expressions where the kernel matrix $K_{m,\mathrm{out}}$ appears, which is convenient for implementation.

State reconstruction. The lifted state $\tilde{\mathbf{z}}$ can be used to reconstruct the original state $\mathbf{x}$, as discussed in Brunton et al. (2022, Section 6). This goal can be achieved, e.g., by using a least-squares estimate, or by augmenting the lifted state with the original one. Here, we consider a regularized least squares estimate for the state reconstruction matrix C. For a given regularization parameter $\lambda > 0$, we define

$$C = \arg\min_{M\in\mathbb{R}^{d\times m}} \frac{1}{n}\sum_{i=1}^{n} \|\mathbf{x}_{i+1} - M\tilde{\mathbf{z}}_{i+1}\|_2^2 + \lambda\|M\|_{\mathrm{HS}}^2. \tag{20}$$

A closed-form expression for matrix C is also included in Appendix D. Note that the training control variables appear in (20) only through the definition of the dynamics, and it would also be possible to use instead an arbitrary dataset of sampled states.

4. Kernels and Koopman LQR

Once we have estimated the state-space representation of the dynamical system of interest, we can use linear predictive control techniques, as the Koopman approach transforms a non-linear system in a linear one (Korda & Mezić, 2018). In the following, we focus on the LQR, and more particularly on the infinite horizon LQR (Mania et al., 2019) due to its theoretically appealing properties. The key idea is that the optimization problem at the basis of the LQR is rendered tractable by using a finite dimensional embedding of the state and inputs, which can be achieved by sketching techniques, as discussed in Section 3.

LQR for exact dynamics. We first consider using the exact kernel k and the transformation ϕ. We consider a control objective that is quadratic in the lifted state and control inputs, via the weighting operators $Q : \mathcal{H}_1 \to \mathcal{H}_1$ and $R : \mathbb{R}^{n_u} \to \mathbb{R}^{n_u}$. Then, the LQR strategy requires to solve the following optimal control problem, for which a time-invariant analytical solution is available:

$$\min_{\mathbf{u}_0,\mathbf{u}_1,\dots} \lim_{T\to\infty} \sum_{i=0}^{T} \langle z_i, Qz_i\rangle_{\mathcal{H}_1} + \langle \mathbf{u}_i, R\mathbf{u}_i\rangle_{\mathbb{R}^{n_u}} \quad \text{s.t. (11(a)), (11(b)).} \tag{21}$$

LQR for approximated dynamics. When the dynamics are approximated by using the Nyström approach as in (16(b)), the optimal control problem becomes

$$\min_{\mathbf{u}_0,\mathbf{u}_1,\dots} \lim_{T\to\infty} \sum_{i=0}^{T} \langle \tilde{z}_i, Q\tilde{z}_i\rangle_{\mathcal{H}_1} + \langle \mathbf{u}_i, R\mathbf{u}_i\rangle_{\mathbb{R}^{n_u}} \quad \text{s.t. (16(a)), (16(b)).} \tag{22}$$

According to (19), the problem in (22) can equivalently be rewritten in the basis $U_{\text{out},\mathbf{x}}$ for weighting matrices $\tilde{Q} \in \mathbb{R}^{m\times m}$,

$$\tilde{Q} = (K^{\dagger}_{m,\text{out}})^{1/2}\tilde{S}_{\text{out}} Q \tilde{S}^{*}_{\text{out}} (K^{\dagger}_{m,\text{out}})^{1/2}, \tag{23}$$

and R, as follows:

$$\min_{\mathbf{u}_0,\mathbf{u}_1,\dots} \lim_{T\to\infty} \sum_{i=0}^{T} \langle \tilde{\mathbf{z}}_i, \tilde{Q}\tilde{\mathbf{z}}_i\rangle_{\mathbb{R}^m} + \langle \mathbf{u}_i, R\mathbf{u}_i\rangle_{\mathbb{R}^{n_u}} \quad \text{s.t. (18(a)), (18(b)).} \tag{24}$$

Based on (23), problems (22) and (24) are equivalent, and yield the same optimal control sequence in form of state feedback (Hager & Horowitz, 1976). Practically, one can choose $\tilde{Q} = C^{*}Q'C$ for some $Q' : \mathbb{R}^d \to \mathbb{R}^d$, and C is the reconstruction matrix defined in (20) in order to define an objective that can be interpreted as a penalization of the states. Denoting the optimal gain resulting from (22) as $\tilde{K} : \mathcal{H}_1 \to \mathbb{R}^{n_u}$, the optimal control law is $\mathbf{u}_k = \tilde{K}\tilde{z}_k$. From (19), we have that the state-feedback input can be rewritten as $\mathbf{u}_k = \tilde{K}U_{\text{out},\mathbf{x}}\tilde{\mathbf{z}}_k$, with $\tilde{K}U_{\text{out},\mathbf{x}} : \mathbb{R}^m \to \mathbb{R}^{n_u}$ being the solution of the LQR problem in (24). When dealing with the true nonlinear dynamics in Section 6, we consider a state-feedback control law of the form

$$\mathbf{u}_k = \tilde{K}\Pi_{\text{out},\mathbf{x}}\psi(\mathbf{x}_k), \tag{25}$$

where $\mathbf{x}_k$ is the true state of the system, to perform control in closed loop.

Overall pipeline. The content of this section can be summarized in a whole pipeline that can be used for linear identification and control of nonlinear systems. Such a pipeline consists of the following steps:

(1) sample Nyström input and output landmarks $\tilde{\mathbf{x}}^{\text{in}}_1, \dots, \tilde{\mathbf{x}}^{\text{in}}_m$ and $\tilde{\mathbf{x}}^{\text{out}}_1, \dots, \tilde{\mathbf{x}}^{\text{out}}_m$ uniformly from the training set;
(2) compute the operators in (18(b)) to obtain a data-driven linear system;
(3) use these linear dynamics to solve problem (24) and get the optimal state-feedback gain.

5. Theoretical analysis

In this section, we perform a theoretical analysis of the proposed system identification method, and assess its effect on the LQR problem. In particular, we will compare the Nyström-based approach with the data-driven model based on the exact kernel. Overall, we show that the Nyström-based model of (16(b)) is a provably accurate approximation of the dynamics in (11(b)), that can be safely used for control purposes, in place of the intractable infinite-dimensional model of (11(b)).

Layout. In Section 5.1, we state the assumptions at the basis of our analysis. In Section 5.2, we bound the error, introduced by the Nyström approximation, on operators A_γ and B_γ from (11(b)) (compound in the transition operator G_γ). In Section 5.3, we show that the Riccati operator obtained with the Nyström dynamics is close to the one obtained with the exact kernel. Finally, in Section 5.4, we use the aforementioned results to show that, when plugging the optimal control from (22) into the dynamics of (11(b)), the LQR objective function is close to the one obtained when solving (21) directly.

5.1. Hypotheses

In this section, we introduce the hypotheses of our theoretical derivations in the next sections.

Assumption 1 (*Bounded Kernel*)**.** The stationary kernel k is bounded, i.e., there exists a positive constant $\kappa < \infty$ such that $k(\mathbf{x}, \mathbf{x}') \leq k^2$ for any $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$.

Note that under Assumption 1, it holds in particular $\|Z_{\mathbf{x}}\| \leq \kappa$ and $\|S_{\mathbf{x}}\| \leq \kappa$. Indeed, for any x, it implies $\|\psi(x)\| = \sqrt{k(x,x)} \leqslant \kappa$, and thus for any $g \in \mathcal{H}_1$ with $\|g\| \leqslant 1$, it holds

$$\|Z_{\mathbf{x}}g\|_2 = n^{-1/2}\|[\langle g, \psi(\mathbf{x}_2)\rangle, \dots, \langle g, \psi(\mathbf{x}_{n+1})\rangle]^T\|_2 \leqslant n^{-1/2}\left(\sum_{i=2}^{n+1}\|g\|^2\kappa^2\right)^{1/2} \leqslant \kappa,$$

and a similar argument holds for $S_{\mathbf{x}}$.

The following assumptions are standard and guarantee that the LQR problems in (21) and (22) are well-posed and admit an analytical solution in the form of a static state-feedback gain (Hager & Horowitz, 1976). Minimal assumptions on A_γ, B_γ from (12) and (13), and C from (20), for stabilizability and detectability to hold, could be derived by leveraging the theoretical deployments reported, e.g., in Bensoussan, Da Prato, Delfour, and Mitter (2007).

Assumption 2 (*Stabilizability, Hager and Horowitz (1976)*)**.** The dynamical systems in (11(b)) and (16(b)) are *stabilizable*, i.e., $\exists M : \mathcal{H}_1 \to \mathbb{R}^{n_u}$ such that $\rho(A_\gamma + B_\gamma M) < 1$, and $\exists \tilde{M} : \mathcal{H}_1 \to \mathbb{R}^{n_u}$ such that $\rho(\tilde{A}_\gamma + \tilde{B}_\gamma \tilde{M}) < 1$, where $\rho(W)$ is the spectral radius of operator W.

Stabilizability assumes that the dynamical systems can be driven to zero with a suitable state-feedback gain. This is a weaker assumption than controllability (i.e., assuming that the system can be driven to any location via the choice of suitable feedback gain), that is often used in the analysis of the LQR algorithm (Mania et al., 2019). However, the controllability assumption in $\mathcal{H}_1$ is not satisfied by the Nyström dynamics, as the $\tilde{z}$ functions always live in the subspace $\tilde{\mathcal{H}}_{\text{out}}$, as we discussed in Section 3.3.

Assumption 3 (*Detectability, Hager and Horowitz (1976)*)**.** The dynamical systems in (11(b)) and (16(b)) are *detectable*, i.e.,

$$\exists t, s, b, d \geq 0 \text{ such that } \|A_\gamma^t z\|_{\mathcal{H}_1} \geq b\|z\|_{\mathcal{H}_1} \Rightarrow \left\langle z, \sum_{i=0}^{s} A_\gamma^{i*} Q A_\gamma^i z \right\rangle_{\mathcal{H}_1} \geq d \langle z, z\rangle_{\mathcal{H}_1},$$

$$\exists t, s, b, d \geq 0 \text{ such that } \|\tilde{A}_\gamma^t \tilde{z}\|_{\mathcal{H}_1} \geq b\|\tilde{z}\|_{\mathcal{H}_1} \Rightarrow \left\langle \tilde{z}, \sum_{i=0}^{s} \tilde{A}_\gamma^{i*} Q \tilde{A}_\gamma^i \tilde{z} \right\rangle_{\mathcal{H}_1} \geq d \langle \tilde{z}, \tilde{z}\rangle_{\mathcal{H}_1}.$$

Detectability assumes that, if unstable dynamics happen in the linear system, these must be observed (and taken into account in the design of the LQR objective function).

Assumption 4. Let Q, R be the weights of the LQR problem (21), and let P be a solution of the discrete algebraic Riccati equation

$$P = A^*PA - A^*PB(R + B^*PB)^{-1}B^*PA + Q. \tag{26}$$

Then, $\sigma_{\min}(P) \geq 1$.

As discussed by Mania et al. (2019), this technical assumption can be fulfilled by re-scaling Q and R accordingly. Indeed, if P is a solution of (26), then ηP, $\eta > 0$ is also a solution, provided that Q and R are multiplied by η.

5.2. Accuracy of the Nyström approximation of the transition operator

We now upper-bound the error (in operator norm) induced by the Nyström approximation on the transition operator G_γ. Although other works studied sketched estimators of the Koopman operator, our bound notably differs from El Ahmad, Brogat-Motte, Laforgue, and d'Alché Buc (2024) and Meanti et al. (2023) who consider different norms (Hilbert–Schmidt and operator norm but on different spaces) and dynamical systems without control.

Theorem 5 (*Convergence Rate for $\tilde{G}_\gamma - G_\gamma$*)**.** *Under Assumption 1, for any $\gamma > 0$, it holds with probability $1 - \delta$ that*

$$\|\tilde{G}_\gamma - G_\gamma\| \leqslant \left(\frac{\kappa}{\gamma} + \frac{1}{\gamma^{1/2}}\right) 4\kappa \sqrt{\frac{3}{m}\log\left(\frac{8m}{5\delta}\right)} + \frac{48\kappa^3}{\gamma^{3/2}}\frac{1}{m}\log\left(\frac{8m}{5\delta}\right). \tag{27}$$

Proof. The proof for this result is provided in Appendix B.1. □

Note that, if we rewrite the approximate dynamics $\tilde{z}_{t+1} = \tilde{G}_\gamma \phi(\mathbf{w}_t)$ from (16(b)) in the autoregressive form as

$$\tilde{z}_{t+1} = \tilde{A}_\gamma \psi(\mathbf{x}_t) + \tilde{B}_\gamma \mathbf{u}_t, \tag{28}$$

Theorem 5 automatically translates in bounds on the approximation of A_γ and B_γ. Indeed, under the assumptions of Theorem 5, the right-hand side of (27) is also an upper bound on $\|A_\gamma - \tilde{A}_\gamma\|$ and $\|B_\gamma - \tilde{B}_\gamma\|$. Formally, the operators $\tilde{A}_\gamma, \tilde{B}_\gamma$ in (28) can be defined from $\tilde{G}_\gamma$ as $\tilde{A}_\gamma = \tilde{G}_\gamma[I_{\mathcal{H}_1}, 0_{\mathcal{H}_1 \to \mathbb{R}^{n_u}}]^* : \mathcal{H}_1 \to \mathcal{H}_1$ and $\tilde{B}_\gamma = \tilde{G}_\gamma[0_{\mathbb{R}^{n_u} \to \mathcal{H}_1}, I_{\mathbb{R}^{n_u}}]^* : \mathbb{R}^{n_u} \to \mathcal{H}_1$. As a consequence, it holds $\|A_\gamma - \tilde{A}_\gamma\| = \|(G_\gamma - \tilde{G}_\gamma)[I_{\mathcal{H}_1}, 0_{\mathcal{H}_1 \to \mathbb{R}^{n_u}}]^*\| \leqslant \|G_\gamma - \tilde{G}_\gamma\|$, and a similar argument holds for B.

In the following, we will use for simplicity the notations $A := A_\gamma, B := B_\gamma, \tilde{A} := \tilde{A}, \tilde{B} := \tilde{B}$. However, all these operators implicitly depend on the choice of the regularization parameter γ.

5.3. Convergence analysis for the Riccati operator

In this section, we show that the Riccati operators for problems in (21) and (22) are ϵ-close in operator norm, provided that $\|G_\gamma - \tilde{G}_\gamma\| \leq \epsilon$. Note that, according to Theorem 5, we can set ϵ as function of m, namely $\epsilon = \left(\frac{\kappa}{\gamma} + \frac{1}{\gamma^{1/2}}\right) 4\kappa \sqrt{\frac{3}{m}\log\left(\frac{8m}{5\delta}\right)} + \frac{48\kappa^3}{\gamma^{3/2}}\frac{1}{m}\log\left(\frac{8m}{5\delta}\right)$. This fact means that we transfer the guarantees for G_γ from Section 5.2 to guarantees on the fundamental building block of the LQR solution, i.e., the Riccati operator.

Let

$$\begin{aligned} F(P, A, B) &= P - A^*[P - PB(R + B^*PB)^{-1}B^*P]A - Q \\ &= P - A^*P(I + BR^{-1}B^*P)^{-1}A - Q. \end{aligned}$$

Then, as proven by Hager and Horowitz (1976, Theorem 9), under the assumptions in Section 5.1, the LQR problems in (21) and (22) admit analytical solutions given by the static state-feedback gains

$$K : \mathcal{H}_1 \to \mathbb{R}^{n_u},\ K = -(R + B^T PB)^{-1}B^*PA, \tag{29}$$

$$\tilde{K} : \mathcal{H}_1 \to \mathbb{R}^{n_u},\ \tilde{K} = -(R + \tilde{B}^T \tilde{P}\tilde{B})^{-1}\tilde{B}^*\tilde{P}\tilde{A}, \tag{30}$$

where $P, \tilde{P} : \mathcal{H}_1 \to \mathcal{H}_1$ are the unique self-adjoint, positive semi-definite operators obtained by solving the following discrete-time algebraic Riccati equations (Hager & Horowitz, 1976):

$$F(P, A, B) = 0, \tag{31}$$

$$F(\tilde{P}, \tilde{A}, \tilde{B}) = 0. \tag{32}$$

A fundamental step is now to derive an error rate for the solutions, upper bounding the quantity $\|P - \tilde{P}\|$. To do so, we can state and prove an analogous proposition to Mania et al. (2019, Proposition 2), based on the fundamental results in Konstantinov, Petkov, and Christov (1993), and generalized to the case of operator dynamics.

Lemma 6 (*Convergence Rate for $\tilde{P} - P$*)**.** *Let the systems in (21) and (22) be stabilizable and detectable. Let $\sigma_{min}(P)$ be the smallest singular value of the Riccati operator associated to (21), and let $\tilde{P}$ be the Riccati operator for (22). Assume R is positive definite, and $\sigma_{min}(P) \geq 1$. Let $L = A + BK$, where K is the stabilizing Riccati gain defined in (29), let $\rho(L)$ be the spectral radius of L, and let $\tau(L, \zeta) = \sup\{\|L^k\|\zeta^{-k}, k \geq 0\}$, with ζ being a constant such that $\rho(L) \leq \zeta < 1$. Then, for $\|A - \tilde{A}\| \leq \epsilon$, $\|B - \tilde{B}\| \leq \epsilon$, and*

$$\epsilon < \min\{\|B\|, \frac{1}{12}\frac{1}{(\|L\| + 1)^2 + (\|P\| + 1)}\frac{(1 - \zeta^2)^2}{\tau(L, \zeta)^4} \cdot (\|A\| + 1)^{-2}(\|P\| + 1)^{-2} \cdot (\|B\| + 1)^{-3} \cdot (\|R^{-1}\| + 1)^{-2}\},$$

it holds that the error on the Riccati operator due to the Nyström approximation is upper bounded as follows:

$$\|P - \tilde{P}\| \leq 6\epsilon \frac{\tau(L, \zeta)^2}{1 - \zeta^2}(\|A\| + 1)^2(\|P\| + 1)^2 \cdot (\|B\| + 1)(\|R^{-1}\| + 1).$$

Remark 7. When talking about stabilizability, we are considering the data-driven dynamics described by (11(b)) and (16(b)). This fact means that we assume these systems are stabilizable in $\mathcal{H}_1$, not in the original state space $\mathbb{R}^d$.

Proof. The proof of this result is close to the one of Mania et al. (2019, Proposition 2) which covers the setting where A and B are matrices. Some technical adjustments are required due to the fact

that we work with operators. We report a sketch of the proof in Appendix C.[5] □

5.4. Convergence analysis of the LQR objective function

In the previous subsections, we showed that $\|A - \tilde{A}\| \leq \epsilon$, $\|B - \tilde{B}\| \leq \epsilon$ for a sufficiently large number of Nyström landmarks m. We further showed that this implies that $\|P - \tilde{P}\| \leq \mathcal{O}(\epsilon)$. Now, we are interested in assessing how suitable the Nyström dynamics are for solving the LQR problem, instead of using the intractable, infinite-dimensional dynamics associated to an exact kernel. In order to do so, we first solve the LQR problem from (22) based on the Nyström approximation of the dynamics. Then, we plug the optimal control sequence retrieved in this way in the exact kernel dynamics of (11(b)), and compare the LQR objective function with this control sequence, and the control sequence we would obtained by solving the problem in (21). This type of error analysis resembles the one in Mania et al. (2019), in the context of certainty equivalence. In order to do so, let us define two control sequences $(\tilde{\mathbf{u}}_0^{\text{opt}}, \tilde{\mathbf{u}}_1^{\text{opt}}, \dots)$, and $(\mathbf{u}_0^{\text{opt}}, \mathbf{u}_1^{\text{opt}}, \dots)$. Moreover, for a given initial condition $\hat{z}_0 = z_0$, let

$$\mathcal{J} := \lim_{T\to\infty} \sum_{i=0}^{T} \langle z_i, Qz_i\rangle_{\mathcal{H}_1} + \langle \mathbf{u}_i^{\text{opt}}, R\mathbf{u}_i^{\text{opt}}\rangle_{\mathbb{R}^{n_u}},$$
$$z_{i+1} = Az_i + B\mathbf{u}_i^{\text{opt}}, \mathbf{u}_i^{\text{opt}} = Kz_i, \tag{33}$$

$$\hat{\mathcal{J}} := \lim_{T\to\infty} \sum_{i=0}^{T} \langle \hat{z}_i, Q\hat{z}_i\rangle_{\mathcal{H}_1} + \langle \tilde{\mathbf{u}}_i^{\text{opt}}, R\tilde{\mathbf{u}}_i^{\text{opt}}\rangle_{\mathbb{R}^{n_u}},$$
$$\hat{z}_{i+1} = A\hat{z}_i + B\tilde{\mathbf{u}}_i^{\text{opt}}, \tilde{\mathbf{u}}_i^{\text{opt}} = \tilde{K}\hat{z}_i. \tag{34}$$

Then, we upper bound the error $\hat{\mathcal{J}} - \mathcal{J}$.

Theorem 8 (*Convergence Rate for $\hat{\mathcal{J}} - \mathcal{J}$*). *Let z_0 be the initial state of the dynamical system of interest. Let $\rho(L)$ be the spectral radius of the closed-loop operator $L = A + BK$, where K is given by (29). Moreover, let $\mathcal{J}$ be as in (33), and $\hat{\mathcal{J}}$ as in (34). Let ζ be a real number such that $\rho(L) \leq \zeta < 1$. Lastly, let $\Gamma = 1 + \max\{\|A\|, \|B\|, \|P\|, \|K\|\}$. Under Assumptions 1, 2, and 3, if $\|A - \tilde{A}\| \leq \epsilon$, $\|B - \tilde{B}\| \leq \epsilon$, assuming there exists some function g such that $\|P - \tilde{P}\| \leq g(\epsilon)$ for some value of ϵ chosen small enough so that $g(\epsilon) \leq \frac{1-\zeta}{6\|B\|\tau(L,\zeta)\Gamma^2}$, and $\sigma_{min}(R) \geq 1$, we have that*

$$\hat{\mathcal{J}} - \mathcal{J} \leq 36\sigma_{max}(R)\Gamma^9 g(\epsilon)^2\kappa^2 \frac{\tau(L,\zeta)^2}{1-\zeta^2}.$$

Proof. The proof of this result can be obtained by following the same steps of Mania et al. (2019, Theorem 1). Although deployed for A and B being matrices, these steps do not depend on the dimensionality of the state space, i.e., can be used also when considering function-valued dynamics, as we do in our work.[4] We begin by computing an upper bound on the error for the Riccati gain $\|K - \tilde{K}\|$. We then show that when this error is small enough, K stabilizes the system with exact kernel defined in (11(b)). Lastly, we use (Fazel, Ge, Kakade, & Mesbahi, 2018, Lemma 10) to upper bound the error on the objective functions. □

Remark 9. Considering the constants appearing in the bounds of this section, they can be interpreted as follows: κ is the upper bound on the kernel function's values, i.e., the kernel variance; γ, which is the regularization weight, could be heuristically tuned based on the variance of the noise affecting the system in (1). δ should in principle be chosen to be as small as possible, as our bound in Theorem 5 holds with probability at least $1 - \delta$. The other constants (e.g., the norms of the learned operators) depend on the data used for training.

6. Simulation results

In this section, we propose a numerical validation of the Nyström-based system identification and the subsequent LQR algorithm discussed in the previous sections. We will start in Section 6.1 with a proof-of-concept evaluation of the control law retrieved by our Nyström-based approach, on the nonlinear dynamics with known optimal control discussed in Guo, Houska, and Villanueva (2022). Then, we will consider the Duffing oscillator in Section 6.2, and we will study the application of cloth manipulation in Section 6.3.

In all our case-studies, the kernel lengthscale and the regularization parameters are determined with a grid search combined with 5-fold cross validation. For simplicity, we use the same regularization parameter for the estimate of the dynamics and the state reconstruction, i.e. $\lambda = \gamma$.

6.1. Proof-of-concept dynamics

In this section, we assess how the control law defined by (25), denoted as u, compares to the known optimal control $u^{\text{true, opt}}$ for an illustrative single-input dynamical system. This simulation serves as a proof-of-concept to show that the proposed control law, retrieved by an LQR-based methodology, is not too far from the true optimal control. The continuous dynamics defined in the following are discretized with the Runge–Kutta method.

We consider the nonlinear dynamics studied by Guo et al. (2022):

$$\dot{x} = -x^3 + u. \tag{35}$$

The known optimal control, for $Q = 1$ and $R = 1$, stabilizing the origin is $u^{\text{true, opt}} = x^3 - x\sqrt{1 + x^4}$. In the simulation, we started at $x_0 = 0.9$. The Nyström-based system was defined using the Matérn-5/2 kernel with lengthscale equal to 1 (i.e. $k(x, y) := (1 + \sqrt{5}r + 5r^2/3)\exp(-\sqrt{5}r)$ where $r = \|x - y\|$), and regularization constant equal to 10^{-6}. It was trained with 20 trajectories of length 2 s, with a sampling time of 0.01 s. The initial states for the training trajectories, and u_t were uniformly sampled from $[-1, 1]$. The Nyström-based LQR is solved with $Q = C^TC$ and $R = 1$. Given the size of the training dataset, we further compare both the Nyström-based control law and the true optimal control with the exact kernel dynamics from (11(b)). Note that the latter induces states belonging to $\text{ran}(Z_{\mathbf{x}})$, hence when starting from an initial state in $\text{ran}(Z_{\mathbf{x}})$ it can be seen as a special case of the Nyström dynamics (16(b)) by taking $\Pi_{\text{in}} = I$, and choosing $m = n$ and $\tilde{\mathbf{x}}_1^{\text{out}} = \mathbf{x}_2, \tilde{\mathbf{x}}_m^{\text{out}} = \mathbf{x}_{n+1}$ in the definition of $\tilde{S}_{\text{out}}$. By doing so, we get $\mathbf{z}_{t+1} = (Z_{\mathbf{x}}Z_{\mathbf{x}}^*)^{1/2}(SS^* + \gamma I)^{-1}S\begin{bmatrix} Z_{\mathbf{x}}^*[(Z_{\mathbf{x}}Z_{\mathbf{x}}^*)^\dagger]^{-1/2}\mathbf{z}_t, \\ \mathbf{u}_t \end{bmatrix}$, and $\mathbf{z}_0 := U_{\text{out},\mathbf{x}}^*\psi(\mathbf{x}_0)$.

Besides computing the infinite-horizon LQR cost, we assess the distance between the two control laws with this key performance indicator: $\text{RMSE}_{\%}^{\mathbf{u}} = \sqrt{\sum_{t=1}^{200}(u_t - u_t^{\text{true, opt}})^2 / \sum_{t=1}^{200}(u_t^{\text{true, opt}})^2} \cdot 100$. As shown in Fig. 2(a), the retrieved control law is close to the optimal one, achieving a final median $\text{RMSE}_{\%}^{\mathbf{u}}$ of approximately 13%, both for the exact and for the Nyström kernel. A qualitative comparison of the control laws involved in this simulation is offered in Fig. 2(b). Fig. 2(c) shows the optimal and data-driven control laws as a function of the state $x \in [-1, 1]$ of system (35). Even though the data-driven control laws have an error w.r.t. the optimal one near the boundaries of the training domain, Fig. 2(b) shows that convergence to the true optimal control law happens

[5] The full derivation of the result can be found at https://github.com/LCSL/nys-koop-lqr/tree/main/proofs.

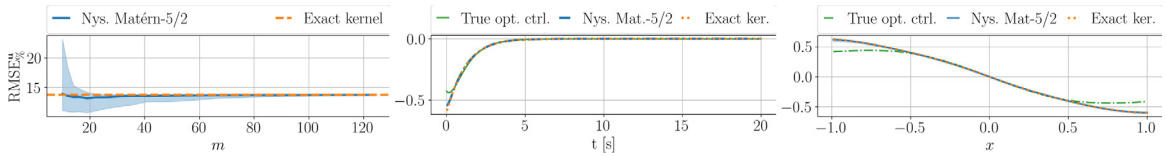

(a) Error on the control input. (b) Example control law for $x_0 = 0.9$. (c) Control law over domain.

Fig. 2. (a): Evaluation of the error between the control law defined in (25) and the true optimal control for the system in (35). Median, 15th and 85th percentile computed across 200 seeds. (b): A qualitative visualization of the optimal control retrieved, for $m = 100$. (c): A comparison between the true optimal control, the one defined in (25), and its version obtained with an exact kernel, on the state space of the system (35), for $m = 100$. For the Nyström approach, we show median, 15th and 85th percentile computed across 200 seeds.

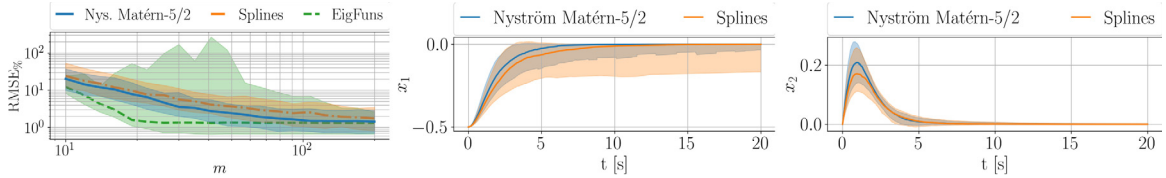

(a) Open-loop forecasts. (b) Stabilization of x_1. (c) Stabilization of x_2.

Fig. 3. (a): $\text{RMSE}_{\%}$ between the true trajectory and the one forecasted in open-loop for the Duffing oscillator, with three different feature representations (splines, eigenfunction approximation by Korda and Mezić (2020), and Nyström approximation of the Matérn-5/2 kernel), as a function of the dimensionality of the feature vector, m. Median, 15th and 85th percentile computed across 200 test trajectories with random initial conditions from the unit ball. (b)–(c): The LQR control strategy to stabilize towards the origin, with $m = 20$, starting from the initial conditions $[-0.5, 0.0]^T$. Most importantly, when the splines are used, in 2 cases the LQR gain yields unstable nonlinear dynamics (not included in the percentile range). Median, 15th and 85th percentile computed across 200 seeds.

Table 1
Infinite-horizon LQR costs for data-driven and optimal control laws, on proof-of-concept dynamics (35) (for the Nyström approximation we show **median**, 15th and 85th percentiles across 200 seeds).

Cost of (25), Nystr.	$m = 10$	**57.12**, 57.00, 58.19
	$m = 50$	**57.10**, 57.05, 57.10
	$m = 100$	**57.10**, 57.09, 57.10
Cost of (25), exact ker.		57.10
Cost of opt. ctrl.		56.74

fast even when starting in those regions (e.g., for $x = 0.9$). Lastly, Table 1 compares the infinite-horizon LQR cost achieved by the control law in (25), and by the optimal control for the task considered.

6.2. Duffing oscillator

We now consider a classic benchmark for nonlinear control, namely the *damped Duffing oscillator* (Korda & Mezić, 2020). The nonlinear dynamics are defined in continuous time domain and given by the following differential equations:

$$\dot{x}_1 = x_2, \qquad \dot{x}_2 = -0.5x_2 - x_1(4x_1^2 - 1) + 0.5u. \tag{36}$$

Similarly to Korda and Mezić (2020), the dynamics are discretized with the Runge–Kutta method to obtain a difference equation of the form of (1), with a discretization step of 0.01 s. For $u = 0$, this dynamical system is known to have an unstable equilibrium point at the origin, and two stable equilibria at $[-0.5, 0]^T$ and $[0.5, 0]^T$. The training set consists of 100 trajectories of length 5 s and without forcing, and 100 trajectories of length 2 s with forcing, all staring from inside the unit ball around the origin. For the forced data, the system is excited with inputs sampled uniformly at each time step from the interval $[-1, 1]$.

Firstly, we are interested in assessing how well the Nyström-based dynamics in (18(b)) approximate the ones of the true dynamical system in (36). To do so, we test the open-loop forecast of the dynamics subject to a square wave with unitary amplitude and frequency 3.33 Hz, starting from a random initial condition in the unit ball. This means that we excite the data-driven model of (16(b)) with the square wave, reconstruct the state via matrix C from (20), and compare the states reconstructed using C against the states visited by the true nonlinear dynamics. The length of the open-loop forecast is set to 2 s. The results can be observed in Fig. 3(a). The Nyström identification is compared against the splines introduced in Korda and Mezić (2018), with centers sampled uniformly at random from the unit circle, and against the eigenfunction-based method by Korda and Mezić (2020). For the latter approach, we opted for the implementation without optimization of the eigenvalues, which is more robust for large values of m, and practically equivalent to the one with optimization, based on the empirical evaluation carried out in Korda and Mezić (2020). The Nyström identification uses a Matérn-5/2 kernel, which seems empirically to perform better than other Matérn or RBF kernels in this setting. The lengthscale and variance are set to 1 by cross-validation . We set $\gamma = 10^{-6}$. Both the splines and the Nyström approach use the same samples for approximating the input and output spaces. The Nyström landmarks are sampled uniformly at random from the output training set. We consider the normalized root-mean-squared error (RMSE) between the actual states x and the forecasted ones $\tilde{x}$, defined as $\text{RMSE}_{\%} = \sqrt{[\sum_{i=1}^{2}\sum_{k=1}^{200}(\tilde{x}_{i,k} - x_{i,k})^2]/[\sum_{i=1}^{2}\sum_{k=1}^{200}x_{i,k}^2]} \cdot 100$, where $x_{i,k}$ denotes the ith dimension of the state at time step k. We can observe that all the representations converge in median to the same minimum error, but the Nyström method yields a smaller error when a few landmarks are used. Moreover, we can observe that the approach in Korda and Mezić (2020) suffers from a much larger variance, probably due to the short horizon of the unforced training trajectories employed (w.t.r. to 8 s used in Korda & Mezić, 2020).

The goal of the control strategy is to stabilize the system towards the origin $[0, 0]^T$, starting from the initial condition $[-0.5, 0.0]^T$, for $m = 20$. As discussed in Section 4, the optimal LQR gain is retrieved by solving (24), and the optimal

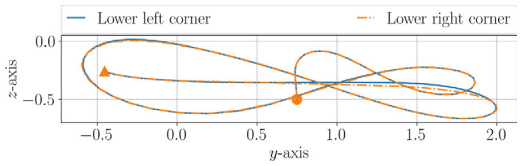

Fig. 4. An example trajectory of the left and right lower corners of the cloth in the y-z plane. The circle denotes the starting position, while the triangle the final one..

state-feedback law $u_k = \tilde{K}\Pi_{\text{out},\mathbf{x}}\psi(\mathbf{x}_k)$ is plugged in (36). The R weighting matrix is set to the identity, while $\tilde{Q} = C^TC$, where C is obtained by solving (20) both with the Nyström and the splines features. The values of the states x_1 and x_2 as a function of time, for a state-feedback control, are shown in Figs. 3(b) and 3(c). We can observe that the Nyström representation yields a better behavior of the closed-loop system, especially when considering the stabilization of the first state of the oscillator. Note that the method reported in Korda and Mezić (2020) builds a complex-valued data-driven system, which would yield to a complex-valued Riccati gain and control input. While it is certainly possible to project the complex-valued LQR-based control law retrieved in this way to be a real-valued signal, we observed in preliminary experiments that such a control law yields a suboptimal behavior, such as steady-state errors. Note that this issue is circumvented by Korda and Mezić (2020), who directly constrain the control to be a real-valued signal, and solve such an optimal control problem online with a model predictive control loop.

6.3. Cloth manipulation

Cloth manipulation is arguably one of the most active and challenging fields in robotic manipulation. Recent work has tried to address this problem in a model-based fashion, proposing approximate, data-driven methods that can identify the dynamics of cloth, and then combine such approximate models with real-time predictive controllers (Amadio et al., 2023; Luque et al., 2024; Zheng et al., 2022). Similarly to Amadio et al. (2023), in this work we consider a squared piece of cloth, modeled based on the simulator deployed by Coltraro et al. (2022). The training dataset is created by moving the upper corners of the cloth in a butterfly-like shape in the y-z plane (see Fig. 4), with different angles w.r.t. the x-y plane. The angle is sampled from the interval $[-60°, 60°]$.[6] The goal is to identify the cloth dynamics in this setup. The cloth is modeled as a square 8×8 mesh. The state is represented by the 3-dimensional position of the mesh points, and the control input is given by the variation in position of the upper corners. The control input is in $\mathbb{R}^6$, while the state is in $\mathbb{R}^{192}$. The sampling time is equal to 0.05 s.

We first assess the performance of the proposed approximate kernel-based identification on this dataset. In particular, the training set consists of 30 trajectories randomly sampled from a total of 40. The length of the trajectories is 5 s, and the sampling time is 0.05 s. We test the open-loop system forecasts in the remaining ten testing scenarios. We compare the Nyström approximation of the RBF kernel against the thin plate splines employed by Korda and Mezić (2018). Note that, in Korda and Mezić (2018) and in Section 6.2, the feature vector contains elements of the form $\|\mathbf{x} - \mathbf{x}_0\|^2 \log \|\mathbf{x} - \mathbf{x}_0\|$, for $\mathbf{x}_0$ being a center sampled uniformly from a box in $\mathbb{R}^d$. However, given the extremely large dimensionality of the cloth's state-space representation, the centers sampled in this way might be quite far from the real cloth configurations visited during the trajectories. We observed, in preliminary experiments, that this unbounded distance between centers and training points yields a highly numerically unstable behavior of splines for $m \geq 200$. For this reason, we propose to sample the splines' centers uniformly at random from the training data, making splines competitive for large values of m as well. Similarly, the Nyström landmarks are sampled uniformly at random from the output training set. The RBF kernel lengthscale is set to 10 and γ to 10^{-7} by cross-validation. A comparison is offered in Fig. 5(a), showing the *unnormalized* RMSE on the testing trajectories. While both feature representations converge to the same RMSE, the RBF representation converges faster and leads to a smaller error in the low-feature regime. Also, note that, given the high dimensionality of the state space, in this case the feature transformation can be seen as a dimensionality reduction technique.

In order to use an LQR approach, we define the control goal as moving the cloth to a target pose. Such a pose is given by the initial pose, rotated by 45° about the upper corners. This simulation captures how well the Koopman models the cloth's deformations, which can be used in fast, dynamic motions. In this case, the number of Nyström landmarks and spline centers is set to 100. Moreover, $\tilde{Q} = 0.0075C^TC$, and R is the identity matrix of suitable dimensions. As we can observe in Fig. 5(b), the Nyström method improves the regulation error, getting closer to the target pose. The better performance of the Nyström features is confirmed by Fig. 5(c). There, we can observe that the cloth gets closer to the target in a shorter amount of time. Note that due to gravity and to the deformable nature of cloth, the cloth will not stay in the target pose once reached, which is why the RMSE increases again for both methods after reaching the target. After the target is reached, we observe that the system enters a stable limit cycle with an RMSE at around 19 cm for the Nyström approach, and 18 cm for the splines. In this cycle, the cloth oscillates around a vertical position that is closer to the target one than the initial pose, for which the RMSE is approximately 30 cm. The limit cycle can be visualized, for instance, in Fig. 6.

As a final assessment, we can consider the following averaged running cost computed on the true nonlinear dynamics:

$$\mathfrak{C} = \frac{1}{T_{\max}} \sum_{t=0}^{T_{\max}} 0.0075(\mathbf{x}_t - \mathbf{r})^T(\mathbf{x}_t - \mathbf{r}) + \mathbf{u}_t^T\mathbf{u}_t, \tag{37}$$

where $\mathbf{r}$ is the target state, $T_{\max} = \text{T s}/0.05$ s, T is a variable time horizon, 0.05 s is the sampling time, and $\mathbf{u}_t$ is the control from (25), computed with splines or the Nyström features. This metric is related to the infinite-horizon LQR cost used to retrieve the state-feedback gain, since $\tilde{Q} = 0.0075C^TC$, and $\mathbf{x}_t \approx C\tilde{\mathbf{z}}_t$ ($\tilde{\mathbf{z}}$ indicating either the Nyström or the splines features). Considering this metric, we observe in Fig. 7 that the Nyström features yield a consistently smaller cost than splines.

7. Conclusions

In this paper, we have studied the combination of the Koopman operator framework with reproducing kernels and the Nyström method, to design linear control for nonlinear dynamical systems. In a novel theoretical analysis, we have linked novel error rates, due to the Nyström approximation, on the dynamical system's representation to its effect on the linear quadratic regulator problem. Lastly, we have evaluated the effectiveness of the proposed method both on two classic control benchmark, and on the task of dynamic cloth manipulation.

[6] We refer the interested reader to Figs. 4 and 5 by Amadio et al. (2023) for a clear visualization of the task considered here.

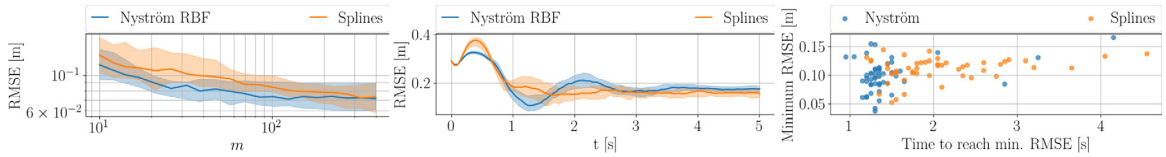

(a) Open-loop forecast. (b) Regulation error. (c) Time-to-min. RMSE.

Fig. 5. (a): the RMSE computed between the true cloth trajectory and the one forecasted in open loop, with two different feature representations (splines vs. Nyström approximation of the RBF kernel). Median, 15th and 85th percentile computed across 10 testing trajectories, sampled with 20 different seeds. (b): regulation error between the target pose of the cloth and the actual one. Median, 15th and 85th percentile computed across 50 seeds. (c): Scatter plot showing the time needed by each simulation of the cloth experiment to reach the minimum distance from the target.

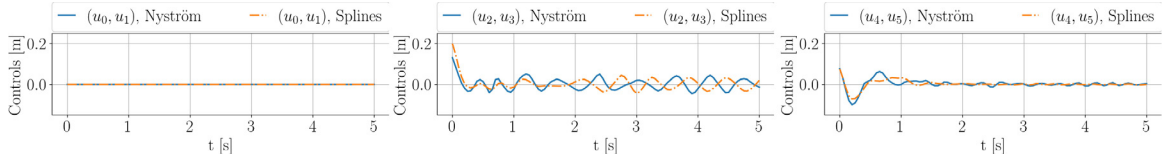

(a) Displacement along x-axis. (b) Displacement along y-axis. (c) Displacement along z-axis.

Fig. 6. An illustrative example of the components of the control signal, for the cloth manipulation simulation. After approaching the target, the system enters a stable limit cycle. The elements in the pair (u_i, u_{i+1}) are identical. This is because, for each coordinate, the training controls (which are the variations in position of the upper cloth's corners) are the same for the left and the right corner, as exemplified in Fig. 4. This bias is reflected in the trajectory obtained using the LQR gain.

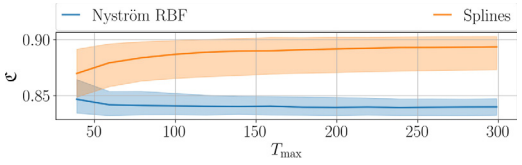

Fig. 7. Cost (37) as a function of the horizon $T_{\max}$..

Acknowledgments

E. Caldarelli, A. Colomé and C. Torras acknowledge support from the project CLOTHILDE ("CLOTH manIpulation Learning from DEmonstrations"), funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Advanced Grant agreement No 741930). E. Caldarelli acknowledges travel support from ELISE (GA No 951847). C. Molinari is part of "GNAMPA" (INdAM) and has been supported by the projects MIUR Excellence Department awarded to DIMA UniGe CUP D33C23001110001, AFOSR FA8655-22-1-7034, MIUR-PRIN 202244A7YL and PON "Ricerca e Innovazione" 2014–2020. C. Ocampo-Martinez acknowledges the support from the project SEAMLESS (PID2023-148840OB-I00) funded by MCIN/AEI. L. Rosasco acknowledges the financial support of the European Research Council, (grant SLING 819789), the European Commission, (Horizon Europe grant ELIAS 101120237), the US Air Force Office of Scientific Research (FA8655-22-1-7034), the Ministry of Education, University and Research (FARE grant ML4IP R205T7J2KP; grant BAC FAIR PE00000013 funded by the EU - NGEU) and the Center for Brains, Minds and Machines (CBMM).

Appendix A. Closed form for G_γ

We recall that we defined the risk in (9) as $\mathcal{R}(W) := \frac{1}{n}\sum_{i=1}^{n} \|\psi(\mathbf{x}_{i+1}) - W\phi(\mathbf{w}_i)\|^2_{\mathcal{H}_1}$.

Expression of the risk. For any orthonormal basis $(g_j)_{j\in\mathbb{N}}$ of $\mathcal{H}_1$, it holds

$$
\begin{aligned}
\mathcal{R}(W) &= \frac{1}{n}\sum_{i=1}^{n}\left[\sum_{j\in\mathbb{N}} \left\langle g_j, \psi(\mathbf{x}_{i+1}) - W\phi(\mathbf{w}_i)\right\rangle^2_{\mathcal{H}_1}\right] \\
&= \frac{1}{n}\sum_{j\in\mathbb{N}}\left[\sum_{i=1}^{n}\left(\left\langle g_j, \psi(\mathbf{x}_{i+1})\right\rangle_{\mathcal{H}_1} - \left\langle W^* g_j, \phi(\mathbf{w}_i)\right\rangle_{\mathcal{H}}\right)^2\right] \\
&= \frac{1}{n}\sum_{j\in\mathbb{N}} \|\sqrt{n} Z_{\mathbf{x}} g_j - \sqrt{n} S W^* g_j\|^2_{\mathbb{R}^n} \\
&= \left\|Z_{\mathbf{x}} - S W^*\right\|^2_{\mathrm{HS}}.
\end{aligned}
$$

Expression of G_γ. Recall $G_\gamma = \arg\min_{W:\mathcal{H}\to\mathcal{H}_1} \mathcal{R}(W) + \gamma\|W\|^2_{\mathrm{HS}}$. Computing the gradient of $\mathcal{R}(W)$ gives $\frac{\partial \mathcal{R}(W)}{\partial W} = -2Z^*_{\mathbf{x}} S + 2WS^*S$. The first-order optimality condition with additive Tikhonov regularization leads to

$$
\begin{aligned}
G_\gamma(S^*S + \gamma I) &= Z^*_{\mathbf{x}} S, \\
G_\gamma &= Z^*_{\mathbf{x}} S (S^*S + \gamma I)^{-1}.
\end{aligned} \tag{A.1}
$$

This gives the expression claimed in (10) by applying the Woodbury identity.

Appendix B. Convergence rate for $\tilde{G}_\gamma - G_\gamma$

We prove here the convergence rate for $\tilde{G}_\gamma - G_\gamma$ given in Theorem 5, using bounds on the Nyström approximation derived in Appendix B.2.

B.1. Proof of Theorem 5

Via (10) and (15), denoting $H_\gamma = SS^* + \gamma I$ the (weighted) regularized kernel matrix, it holds

$$
G_\gamma = Z^*_{\mathbf{x}} H_\gamma^{-1} S, \qquad \tilde{G}_\gamma = \Pi_{\mathrm{out},\mathbf{x}} Z^*_{\mathbf{x}} \tilde{H}_\gamma^{-1} S \Pi_{\mathrm{in}},
$$

where $\tilde{H}_\gamma^{-1} := (S\Pi_{\mathrm{in}} S^* + \gamma I)^{-1}$. In particular, we can observe that

$$
\begin{aligned}
\|\tilde{G}_\gamma - G_\gamma\| \leqslant\ & \|Z^*_{\mathbf{x}} H_\gamma^{-1} S(I - \Pi_{\mathrm{in}})\| \\
& + \|Z^*_{\mathbf{x}}(H_\gamma^{-1} - \tilde{H}_\gamma^{-1}) S\Pi_{\mathrm{in}}\| \\
& + \|(I - \Pi_{\mathrm{out},\mathbf{x}}) Z^*_{\mathbf{x}} \tilde{H}_\gamma^{-1} S \Pi_{\mathrm{in}}\|.
\end{aligned} \tag{B.1}
$$

The first addend can be split as follows:

$$\|Z_{\mathbf{x}}^* H_\gamma^{-1} S\Pi_{\text{in}}^\perp\| \leqslant \|Z_{\mathbf{x}}^* H_\gamma^{-1}\| \|S\Pi_{\text{in}}^\perp\| \leqslant \frac{\kappa}{\gamma}\|S\Pi_{\text{in}}^\perp\|. \tag{B.2}$$

By leveraging the identity $A^{-1} - B^{-1} = A^{-1}(B-A)B^{-1}$ for two invertible operators A and B, the second addend in (B.1) can be bounded as

$$\begin{aligned}\|Z_{\mathbf{x}}^*(H_\gamma^{-1} - \tilde{H}_\gamma^{-1})S\Pi_{\text{in}}\| &= \|Z_{\mathbf{x}}^* H_\gamma^{-1}(S\Pi_{\text{in}}S^* - SS^*)\tilde{H}_\gamma^{-1} S\Pi_{\text{in}}\| \\ &\leqslant \frac{\kappa}{\gamma}\|\Pi_{\text{in}}^\perp S^*\|^2 \|\tilde{H}_\gamma^{-1} S\Pi_{\text{in}}\|.\end{aligned}$$

Note that, using a polar decomposition, $S\Pi_{\text{in}} = (S\Pi_{\text{in}}S^*)^{1/2}U$ for some partial isometry U, and thus

$$\begin{aligned}\|\tilde{H}_\gamma^{-1} S\Pi_{\text{in}}\| &\leqslant \|(S\Pi_{\text{in}}S^* + \gamma I)^{-1/2}\| \\ &\quad \cdot \|(S\Pi_{\text{in}}S^* + \gamma I)^{-1/2}(S\Pi_{\text{in}}S^*)^{1/2}\| \\ &\leqslant \gamma^{-1/2},\end{aligned} \tag{B.3}$$

so that

$$\|Z_{\mathbf{x}}^*(H_\gamma^{-1} - \tilde{H}_\gamma^{-1})S\Pi_{\text{in}}\| \leqslant \frac{\kappa}{\gamma^{3/2}}\|\Pi_{\text{in}}^\perp S^*\|^2. \tag{B.4}$$

Eventually, using again (B.3) the third addend in (B.1) can be bounded as

$$\begin{aligned}\|\Pi_{\text{out},\mathbf{x}}^\perp Z_{\mathbf{x}}^* \tilde{H}_\gamma^{-1} S\Pi_{\text{in}}\| &\leqslant \|\Pi_{\text{out},\mathbf{x}}^\perp Z_{\mathbf{x}}^*\| \|\tilde{H}_\gamma^{-1} S\Pi_{\text{in}}\| \\ &\leqslant \gamma^{-1/2}\|\Pi_{\text{out},\mathbf{x}}^\perp Z_{\mathbf{x}}^*\|.\end{aligned} \tag{B.5}$$

Putting together (B.2), (B.4) and (B.5), we get $\|\tilde{G}_\gamma - G_\gamma\| \leqslant \frac{\kappa}{\gamma}\|\Pi_{\text{in}}^\perp S^*\| + \frac{\kappa}{\gamma^{3/2}}\|\Pi_{\text{in}}^\perp S^*\|^2 + \gamma^{-1/2}\|\Pi_{\text{out},\mathbf{x}}^\perp Z_{\mathbf{x}}^*\|$.

We recall that $\Pi_{\text{in}} = \begin{bmatrix} \Pi_{\text{in},\mathbf{x}} & 0 \\ 0 & I \end{bmatrix}$, so that $\|\Pi_{\text{in}}^\perp S^*\| = \|\Pi_{\text{in},\mathbf{x}}^\perp S_{\mathbf{x}}^*\|$, and both terms $\|\Pi_{\text{in},\mathbf{x}}^\perp S_{\mathbf{x}}^*\|$ and $\|\Pi_{\text{out},\mathbf{x}}^\perp Z_{\mathbf{x}}^*\|$ can be bounded using Lemma 11, which is itself derived from Rudi et al. (2015). We thus get

$$\begin{aligned}\|\tilde{G}_\gamma - G_\gamma\| &\leqslant \left(\frac{\kappa}{\gamma} + \frac{1}{\gamma^{1/2}}\right) 4\kappa\sqrt{\frac{3}{m}\log\left(\frac{8m}{5\delta}\right)} \\ &\quad + \frac{48\kappa^3}{\gamma^{3/2}}\frac{1}{m}\log\left(\frac{8m}{5\delta}\right).\end{aligned} \tag{B.6}$$

B.2. Error induced by the Nyström approximation

The following lemma is a restatement of Rudi et al. (2015, Lemma 6) with slightly different constants. Albeit being originally written in a random design scenario, the same result holds when conditioning on the data. In the following, we denote $\otimes$ the outer-product in $\mathcal{H}_1$, i.e. the rank-one operator defined as $(a \otimes b)c = \langle b, c\rangle_{\mathcal{H}_1} a$.

Lemma 10. *Let $x_1, \dots, x_n \in \mathbb{R}^d$ and denote $\mathcal{C} = \frac{1}{n}\sum_{i=1}^n \psi(x_i) \otimes \psi(x_i)$, where ψ is the canonical feature map associated to a kernel satisfying Assumption 1. Let $\tilde{x}_1, \dots, \tilde{x}_m$ be drawn uniformly from all partitions of cardinality m of $\{x_1, \dots, x_n\}$. Denoting Π_m the orthogonal projection onto $\text{span}(\psi(\tilde{x}_1), \dots, \psi(\tilde{x}_m))$, for any $\gamma \in]0, \|\mathcal{C}\|]$, it holds $\|(I - \Pi_m)(\mathcal{C} + \gamma I)^{1/2}\|^2 \leqslant 3\gamma$, with probability at least $1 - \delta$ provided $m \geq (2 + 5\frac{\kappa^2}{\gamma})\log\left(\frac{4\kappa^2}{\gamma\delta}\right)$.*

Proof. We apply exactly the same proof as in Rudi et al. (2015, Lemma 6) but conditioning on the data (i.e., in particular applying Rudi et al., 2015, Prop. 8 with the v_i drawn i.i.d. according to the empirical distribution and Q being the empirical covariance), with $\mathcal{N}_\infty(\gamma) = \kappa^2/\gamma$. Denoting $w = \log\left(\frac{4k^2}{\gamma\delta}\right)$, we end up with the bound $\|(I - \Pi_m)(\mathcal{C} + \gamma I)^{1/2}\|^2 \leqslant \frac{\gamma}{1-\beta(\gamma)}$, where $\beta(\gamma) \leqslant \frac{2w}{3m} + \sqrt{\frac{2wk^2}{\gamma m}}$. We now derive slightly better constants than in the original lemma in order to satisfy $\beta(\gamma) \leqslant 2/3$ (which gives the claimed result). Indeed, $\beta(\gamma) \leqslant 2/3 \Leftrightarrow m - \frac{3}{2}\sqrt{\frac{2wk^2}{\gamma}}\sqrt{m} - w \geqslant 0$. We solve this inequality as a second-order equation in $\sqrt{m}$, with discriminant $\Delta = w(\frac{9k^2}{2\gamma} + 4)$, which is positive whenever $\gamma < 4k^2\delta^{-1}$. A sufficient condition to satisfy $\beta(\gamma) \leqslant 2/3$ is thus $m \geqslant \left(\frac{3}{4}\sqrt{\frac{2wk^2}{\gamma}} + \frac{1}{2}\sqrt{w(\frac{9k^2}{2\gamma} + 4)}\right)^2$. Using the identity $2(a^2+b^2) \geqslant (a+b)^2$, we get the sufficient condition $m \geqslant w\left(5\frac{k^2}{\gamma} + 2\right)$. □

Lemma 11. *Under Assumption 1, provided that $\Pi_{\text{in},\mathbf{x}}$ and $\Pi_{\text{out},\mathbf{x}}$ are built by sampling uniformly m samples respectively from $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ and $\{\mathbf{x}_2, \dots, \mathbf{x}_{n+1}\}$, it holds with probability $1 - \delta$ $\max(\|\Pi_{in,\mathbf{x}}^\perp S_{\mathbf{x}}^*\|, \|\Pi_{out,\mathbf{x}}^\perp Z_{\mathbf{x}}^*\|) \leqslant 4\kappa\sqrt{\frac{3}{m}\log\left(\frac{8m}{5\delta}\right)}$.*

Proof. Using a polar decomposition and applying Lemma 10 with $m \geqslant (2 + 5\frac{\kappa^2}{\gamma})\log\left(\frac{4\kappa^2}{\gamma\delta}\right)$, it holds

$$\|\Pi_{\text{in},\mathbf{x}}^\perp S_{\mathbf{x}}^*\| = \|\Pi_{\text{in},\mathbf{x}}^\perp (S_{\mathbf{x}}^* S_{\mathbf{x}})^{1/2}\| \tag{B.7}$$

$$\leqslant \|\Pi_{\text{in},\mathbf{x}}^\perp (S_{\mathbf{x}}^* S_{\mathbf{x}} + \gamma I)^{1/2}\| \leqslant \sqrt{3\gamma}, \tag{B.8}$$

and a similar argument holds for $\|\Pi_{\text{out},\mathbf{x}}^\perp Z_{\mathbf{x}}^*\|$.

The condition $m \geqslant (2 + 5\frac{\kappa^2}{\gamma})\log\left(\frac{4\kappa^2}{\gamma\delta}\right)$ is in particular satisfied provided that $m/2 \geqslant \max(2, 5\frac{\kappa^2}{\gamma})\log\left(\frac{4\kappa^2}{\gamma\delta}\right)$, which holds taking $\gamma = \frac{16\kappa^2}{m}\log(\frac{4m}{5\delta})$. This yields the claimed result via a union bound. □

Appendix C. Proof sketch for Lemma 6

Let $L = A + BK$ be the closed-loop operator of the system specified by (11(b)), where K is the stabilizing gain defined in (29). The proof of the result, which closely follows the proof of Mania et al. (2019, Proposition 2), starts by defining a set $\mathcal{S}$ of perturbations of the Riccati operator P as $\mathcal{S} = \{X : \mathcal{H}_1 \to \mathcal{H}_1, X = X^*, P + X \succcurlyeq 0\}$. Then, we can derive a fixed-point operator equation

$$X = \Phi X, \tag{C.1}$$

where $\Phi X = \mathcal{T}^{-1}[F(X + P, A, B) - F(X + P, \tilde{A}, \tilde{B}) - \mathcal{R}X]$, $\mathcal{T}X = I - L^*XL$, and $\mathcal{R}X = L^*X[I + B^*PB(P + X)]^{-1}B^*PBXL$. Let us define $\mathcal{L}(\mathcal{H}_1)$ as the space of continuous linear operators from $\mathcal{H}_1$ to $\mathcal{H}_1$, and $\mathcal{L}(\mathcal{L}(\mathcal{H}_1))$ as the space of continuous linear operators from $\mathcal{L}(\mathcal{H}_1)$ to $\mathcal{L}(\mathcal{H}_1)$. The operator $\mathcal{T} : \mathcal{L}(\mathcal{H}_1) \to \mathcal{L}(\mathcal{H}_1)$ is invertible in $\mathcal{L}(\mathcal{L}(\mathcal{H}_1))$. This fact can be proven by showing that 1 is not in the spectrum of operator $\mathcal{D} : \mathcal{L}(\mathcal{H}_1) \to \mathcal{L}(\mathcal{H}_1)$, defined as $\mathcal{D}X = L^*XL$. This is achieved by upper bounding the spectral radius of $\mathcal{D}$ by means of Gelfand's formula, and by observing that the closed-loop dynamics are stable, i.e., L has spectral radius smaller than 1. Moreover, $\mathcal{T}^{-1}$ is bounded. This can be proven by showing that the Neumann series $\sum_{k=0}^\infty \mathcal{D}^k$ is absolutely convergent in operator norm, and therefore equal to $\mathcal{T}^{-1}$. Then, it can be shown that the unique solution of (C.1) belonging to $\mathcal{S}$ is $\tilde{P} - P$, since, by construction of the fixed point equation, it must hold $F(X + P, \tilde{A}, \tilde{B}) = 0$.

We can now define the following closed subset of $\mathcal{S}$, $\mathcal{S}_\nu = \{X : \|X\| \leq \nu, X = X^*, P + X \succcurlyeq 0\}$. Let $\|\tilde{A} - A\| \leq \epsilon$, $\|\tilde{B} - B\| \leq \epsilon$, where ϵ is the error rate in Theorem 5. Let $N := BR^{-1}B^*$. If we pick

$$\begin{aligned}\nu = \min\Bigg\{&\|N\|^{-1}, \frac{1}{2}, 6\,\epsilon\frac{\tau(L,\zeta)^2}{1-\zeta^2}(\|A\| + 1)^2(\|P\| + 1)^2 \\ &\cdot(\|B\| + 1)(\|R^{-1}\| + 1)\Bigg\},\end{aligned}$$

for $X, X_1, X_2 \in \mathcal{S}_v$, we can show that $\Phi X \in \mathcal{S}_v$, and $\exists \eta < 1$ such that $\|\Phi X_1 - \Phi X_2\| \leq \eta \|X_1 - X_2\|$, that is, Φ is a contraction operator. Then, we can apply Banach fixed-point theorem, meaning that the self-adjoint fixed point $\tilde{P} - P$ of Φ belongs to $\mathcal{S}_v$, i.e., the error on the Riccati operator due to the Nyström approximation has bounded operator norm as a function of ϵ.

Appendix D. Computable expression for Nyström vector dynamics and C

In this appendix, we report the computable expression for the operator appearing in (18(b)), as well as the expression of the reconstruction matrix C in (20). They are the formulas used in practice for the implementation of the method. To simplify the notation, for operators M, N, we denote a block-diagonal operator as $\mathcal{B}(M, N) = \begin{bmatrix} M & 0 \\ 0 & N \end{bmatrix}$. We introduce the following additional notation. In the sequel we refer to the nonlinear kernel k:

- $K_{nm,\text{out}}$ is the kernel between the regression outputs and the output landmarks;
- $K_{nm,\text{in}}$ is the kernel between the regression inputs (state only) and the input landmarks;
- $K_{m,\text{in}}$ is the kernel at the input landmarks;
- $K_{m,\text{in,out}}$ is the kernel between the input landmarks and the output landmarks.

Let us define the sampling operator for the input Nyström landmarks as $\tilde{S}_{\text{in}} : \mathcal{H}_1 \to \mathbb{R}^m$, $\tilde{S}_{\text{in}} g = [g(\tilde{\mathbf{x}}_1^{\text{in}}), \dots, g(\tilde{\mathbf{x}}_m^{\text{in}})]^T$. Note that $\text{ran}(\tilde{S}_{\text{in}}^*) = \tilde{\mathcal{H}}_{\text{in}}$. We denote $\tilde{S}_{\text{in}} = U \Sigma V^*$ its singular value decomposition, where $U : \mathbb{R}^t \to \mathbb{R}^m$, $\Sigma : \mathbb{R}^t \to \mathbb{R}^t$, $V : \mathbb{R}^t \to \tilde{\mathcal{H}}_1$. Here, Σ is diagonal and strictly positive. It holds $\Pi_{\text{in}} = \mathcal{B}(VV^*, I)$ and thus

$$
\begin{aligned}
&U_{\text{out},\mathbf{x}}^* Z_{\mathbf{x}}^* (S \Pi_{\text{in}} S^* + \gamma I)^{-1} S \mathcal{B}(\Pi_{\text{in},\mathbf{x}} U_{\text{out},\mathbf{x}}, I) \\
&= U_{\text{out},\mathbf{x}}^* Z_{\mathbf{x}}^* S \mathcal{B}(\tilde{S}_{\text{in}}^*, I) \mathcal{B}(U \Sigma^{-1}, I) \\
&\quad \cdot \left(\mathcal{B}(V^*, I) S^* S \mathcal{B}(V, I) + \gamma I \right)^{-1} \\
&\quad \cdot \mathcal{B}(\Sigma^{-1} U^*, I) \mathcal{B}(\tilde{S}_{\text{in}}, I) \mathcal{B}(U_{\text{out},\mathbf{x}}, I).
\end{aligned}
$$

Now, we can observe that $V^* V = I$. Moreover, $(\mathcal{B}(V^*, I) S^* S \mathcal{B}(V, I) + \gamma I)$ is full-rank, $\mathcal{B}(U \Sigma^{-1}, I)$ is full column-rank and $\mathcal{B}(\Sigma^{-1} U^*, I)$ is full row-rank, then by Ben-Israel and Greville (2003, Section 1.6) we have that

$$
\begin{aligned}
&\mathcal{B}(U \Sigma^{-1}, I) \left(\mathcal{B}(V^*, I) S^* S \mathcal{B}(V, I) + \gamma I \right)^{-1} \mathcal{B}(\Sigma^{-1} U^*, I) \\
&\quad = \left(\mathcal{B}(\tilde{S}_{\text{in}}, I) S^* S \mathcal{B}(\tilde{S}_{\text{in}}^*, I) + \gamma \mathcal{B}(\tilde{S}_{\text{in}} \tilde{S}_{\text{in}}^*, I) \right)^{\dagger}.
\end{aligned}
$$

This yields

$$
\begin{aligned}
&U_{\text{out},\mathbf{x}}^* \tilde{G}_{\gamma} \mathcal{B}(U_{\text{out},\mathbf{x}}, I) \\
&\quad = (K_{m,\text{out}}^{\dagger})^{1/2} K_{mn,\text{out}} \begin{bmatrix} K_{nm,\text{in}} & \sqrt{n} S_{\mathbf{u}} \end{bmatrix} \\
&\qquad \cdot \left(\begin{bmatrix} K_{mn,\text{in}} \\ \sqrt{n} S_{\mathbf{u}}^T \end{bmatrix} \begin{bmatrix} K_{nm,\text{in}} & \sqrt{n} S_{\mathbf{u}} \end{bmatrix} + \gamma n \mathcal{B}(K_{m,\text{in}}, I) \right)^{\dagger} \\
&\qquad \cdot \mathcal{B}(K_{m,\text{in,out}} (K_{m,\text{out}}^{\dagger})^{1/2}, I).
\end{aligned}
$$

Moreover, having stacked the regression outputs in a matrix $X_{+1} \in \mathbb{R}^{d \times n}$, the matrix C from (20) takes the form $C = X_{+1} K_{nm,\text{out}} (K_{mn,\text{out}} K_{nm,\text{out}} + \lambda n K_{m,\text{out}})^{-1} K_{m,\text{out}}^{1/2}$.

References

Abraham, I., & Murphey, T. D. (2019). Active learning of dynamics for data-driven control using koopman operators. *IEEE Transactions on Robotics*, *35*(5), 1071–1083.

Abraham, I., de la Torre, G., & Murphey, T. (2017). Model-based control using koopman operators. *Robotics: Science and Systems XIII*.

Amadio, F., Delgado-Guerrero, J. A., Colomé, A., & Torras, C. (2023). Controlled Gaussian process dynamical models with application to robotic cloth manipulation. *International Journal on Dynamics Control*, *11*, 3209–3219.

Aronszajn, N. (1950). Theory of reproducing kernels. *Transactions of the American Mathematical Society*, *68*(3), 337–404.

Ben-Israel, A., & Greville, T. N. E. (2003). *CMS books in mathematics, Generalized inverses: theory and applications* (2nd ed.). Springer.

Bensoussan, A., Da Prato, G., Delfour, M. C., & Mitter, S. K. (2007). *Representation and control of infinite dimensional systems (Vol. 1)*. Springer.

Berlinet, A., & Thomas-Agnan, C. (2011). *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media.

Bevanda, P., Beier, M., Lederer, A., Sosnowski, S., Hüllermeier, E., & Hirche, S. (2024). Koopman kernel regression. *Advances in Neural Information Processing Systems*, *36*.

Bevanda, P., Sosnowski, S., & Hirche, S. (2021). Koopman operator dynamical models: Learning, analysis and control. *Annual Reviews in Control*, *52*, 197–212.

Brunton, S. L., Budišić, M., Kaiser, E., & Kutz, J. N. (2022). Modern Koopman theory for dynamical systems. *SIAM Review*, *64*(2), 229–340.

Coltraro, F., Amorós, J., Alberich-Carramiñana, M., & Torras, C. (2022). An inextensible model for the robotic manipulation of textiles. *Applied Mathematical Modelling*, *101*, 832–858.

Das, S., & Giannakis, D. (2020). Koopman spectra in reproducing kernel Hilbert spaces. *Applied and Computational Harmonic Analysis*, *49*(2), 573–607.

Dean, S., Mania, H., Matni, N., Recht, B., & Tu, S. (2020). On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, *20*(4), 633–679.

DeGennaro, A. M., & Urban, N. M. (2019). Scalable extended dynamic mode decomposition using random kernel approximation. *SIAM Journal on Scientific Computing*, *41*(3), A1482–A1499.

Driessen, B. J. (2023). *Koopman operator regression based surrogate modelling for control*. (Master's thesis), Eindhoven University of Technology.

El Ahmad, T., Brogat-Motte, L., Laforgue, P., & d'Alché Buc, F. (2024). Sketch in, sketch out: Accelerating both learning and inference for structured prediction with kernels. In *International conference on artificial intelligence and statistics* (pp. 109–117). PMLR.

Fazel, M., Ge, R., Kakade, S., & Mesbahi, M. (2018). Global convergence of policy gradient methods for the linear quadratic regulator. In *International conference on machine learning* (pp. 1467–1476). PMLR.

Giannakis, D., Henriksen, A., Tropp, J. A., & Ward, R. (2023). Learning to forecast dynamical systems from streaming data. *SIAM Journal on Applied Dynamical Systems*, *22*(2), 527–558.

Gibson, A. J., Calvisi, M. L., & Yee, X. C. (2022). Koopman linear quadratic regulator using complex eigenfunctions for nonlinear dynamical systems. *SIAM Journal on Applied Dynamical Systems*, *21*(4), 2463–2486.

Guo, Y., Houska, B., & Villanueva, M. E. (2022). A tutorial on Pontryagin-Koopman operators for infinite horizon optimal control. In *2022 IEEE 61st conference on decision and control* (pp. 6800–6805). IEEE.

Hager, W. W., & Horowitz, L. L. (1976). Convergence and stability properties of the discrete Riccati operator equation and the associated optimal control and filtering problems. *SIAM Journal on Control and Optimization*, *14*(2), 295–312.

Haggerty, D. A., Banks, M. J., Kamenar, E., Cao, A. B., Curtis, P. C., Mezić, I., et al. (2023). Control of soft robots with inertial dynamics. *Science Robotics*, *8*(81), eadd6864.

Hao, W., Huang, B., Pan, W., Wu, D., & Mou, S. (2024). Deep Koopman learning of nonlinear time-varying systems. *Automatica*, *159*, Article 111372.

Kaiser, E., Kutz, J. N., & Brunton, S. L. (2021). Data-driven discovery of Koopman eigenfunctions for control. *Machine Learning Science and Technology*, *2*(3), Article 035023.

Khosravi, M. (2023). Representer theorem for learning Koopman operators. *IEEE Transactions on Automatic Control*, *68*(5), 2995–3010.

Klus, S., Nüske, F., & Hamzi, B. (2020). Kernel-based approximation of the Koopman generator and Schrödinger operator. *Entropy*, *22*(7), 722.

Klus, S., Schuster, I., & Muandet, K. (2020). Eigendecompositions of transfer operators in reproducing kernel Hilbert spaces. *Journal of Nonlinear Science*, *30*(1), 283–315.

Konstantinov, M. M., Petkov, P. H., & Christov, N. D. (1993). Perturbation analysis of the discrete Riccati equation. *Kybernetika*, *29*(1), 18–29.

Koopman, B. O. (1931). Hamiltonian systems and transformation in Hilbert space. *Proceedings of the National Academy of Sciences of the United States of America*, *17*(5), 315–318.

Korda, M., & Mezić, I. (2018). Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica*, *93*, 149–160.

Korda, M., & Mezić, I. (2020). Optimal construction of Koopman eigenfunctions for prediction and control. *IEEE Transactions on Automatic Control*, *65*(12), 5114–5129.

Kostic, V., Novelli, P., Maurer, A., Ciliberto, C., Rosasco, L., & Pontil, M. (2022). Learning dynamical systems via Koopman operator regression in reproducing kernel Hilbert spaces. *Advances in Neural Information Processing Systems*, *35*, 4017–4031.

Luque, A., Parent, D., Colomé, A., Ocampo-Martinez, C., & Torras, C. (2024). Model predictive control for dynamic cloth manipulation: Parameter learning and experimental validation. *IEEE Transactions on Control Systems Technology*, *32*(4), 1254–1270.

Mania, H., Tu, S., & Recht, B. (2019). Certainty equivalence is efficient for linear quadratic control. *Advances in Neural Information Processing Systems*, *32*, 10154–10164.

Meanti, G., Chatalic, A., Kostic, V., Novelli, P., Pontil, M., & Rosasco, L. (2023). Estimating Koopman operators with sketching to provably learn large scale dynamical systems. *Advances in Neural Information Processing Systems*.

Mezić, I. (2021). Koopman operator, geometry, and learning of dynamical systems. *Notices of the American Mathematical Society*, *68*(7), 1087–1105.

Moyalan, J., Choi, H., Chen, Y., & Vaidya, U. (2023). Data-driven optimal control via linear transfer operators: A convex approach. *Automatica*, *150*, Article 110841.

Muandet, K., Fukumizu, K., Sriperumbudur, B., & Schölkopf, B. (2017). Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning*, *10*(1–2), 1–141.

Musco, C., & Musco, C. (2017). Recursive sampling for the nystrom method. *Advances in Neural Information Processing Systems*, 3833–3845.

Nüske, F., & Klus, S. (2023). Efficient approximation of molecular kinetics using random Fourier features. *Journal of Chemical Physics*, *159*(7).

Nüske, F., Peitz, S., Philipp, F., Schaller, M., & Worthmann, K. (2023). Finite-data error bounds for Koopman-based prediction and control. *Journal of Nonlinear Science*, *33*(1), 14.

Nyström, E. J. (1930). Über Die Praktische Auflösung von Integralgleichungen mit Anwendungen auf Randwertaufgaben. *Acta Mathematics*, *54*, 185–204.

Otto, S. E., & Rowley, C. W. (2021). Koopman operators for estimation and control of dynamical systems. *Annual Review of Control, Robotics, and Autonomous Systems*, *4*(1), 59–87.

Peitz, S., & Klus, S. (2019). Koopman operator-based model reduction for switched-system control of PDEs. *Automatica*, *106*, 184–191.

Philipp, F. M., Schaller, M., Worthmann, K., Peitz, S., & Nüske, F. (2024). Error bounds for kernel-based approximations of the Koopman operator. *Applied and Computational Harmonic Analysis*, *71*, Article 101657.

Proctor, J. L., Brunton, S. L., & Kutz, J. N. (2018). Generalizing Koopman theory to allow for inputs and control. *SIAM Journal on Applied Dynamical Systems*, *17*(1), 909–930.

Rudi, A., Camoriano, R., & Rosasco, L. (2015). Less is more: Nyström computational regularization. *Advances in Neural Information Processing Systems*, *28*, 1657–1665.

Schölkopf, B., Herbrich, R., & Smola, A. J. (2001). A generalized representer theorem. In *International conference on computational learning theory* (pp. 416–426). Springer.

Schölkopf, B., & Smola, A. J. (2002). *Learning with kernels: Support vector machines, regularization, optimization, and beyond*. MIT Press.

Shi, H., & Meng, M. Q.-H. (2022). Deep Koopman operator with control for nonlinear systems. *IEEE Robotics and Automation Letters*, *7*(3), 7700–7707.

Simchowitz, M., Mania, H., Tu, S., Jordan, M. I., & Recht, B. (2018). Learning without mixing: Towards a sharp analysis of linear system identification. In *Conference on learning theory* (pp. 439–473). PMLR.

Steinwart, I., & Christmann, A. (2008). *Support vector machines*. Springer Science & Business Media.

Wasserman, L. (2006). *All of nonparametric statistics*. Springer Science & Business Media.

Williams, C., & Seeger, M. (2000). Using the Nyström method to speed up kernel machines. *Advances in Neural Information Processing Systems*, *13*, 682–688.

Yin, H., Welle, M. C., & Kragic, D. (2022). Embedding Koopman optimal control in robot policy learning. In *IEEE/RSJ international conference on intelligent robots and systems* (pp. 13392–13399). IEEE.

Zheng, C. X., Colomé, A., Sentis, L., & Torras, C. (2022). Mixtures of controlled Gaussian processes for dynamical modeling of deformable objects. In *Learning for dynamics and control conference* (pp. 415–426). PMLR.

Edoardo Caldarelli obtained the Bachelor Degree in Engineering of Computing Systems from Politecnico di Milano (2018), the Master of Science in Computer Science from ETH Zürich (2020), and the Doctoral Degree in Automatic Control, Robotics and Vision from Universitat Politècnica de Catalunya (2025), within the ELLIS Ph.D. program. He is currently a postdoctoral researcher at Istituto Italiano di Tecnologia. His research interests focus on theory and robotic applications of kernel-based, data-driven control algorithms.

Antoine Chatalic received the M.S. and Ph.D. degrees in computer science and signal processing from the University of Rennes 1 (France), respectively in 2017 and 2020. After a postdoctoral position at the University of Genoa, he joined the CNRS in 2024 and is currently working at Gipsa-lab (Grenoble, France). His research interests include large-scale kernel methods, compressive learning and sensing, sketching and dimensionality reduction, statistical learning theory, differential privacy.

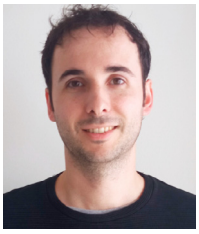

Adrià Colomé (https://www.iri.upc.edu/staff/acolome) is a Research Scientist at the Institut de Robòtica i Informàtica Industrial (CSIC-UPC) in Barcelona, specializing in robotics. He earned dual B.Sc. degrees in Mathematics and Industrial Engineering in 2009, followed by a M.Sc. (2011) and a Ph.D. (2017) in Automatic Control from the Technical University of Catalonia (UPC). His research interests span a range of robotics topics, including robot kinematics and dynamics, motion learning, variable impedance control, and the manipulation of deformable objects. Dr. Colomé has published numerous papers and contributed to various projects in these areas.

Cesare Molinari received his Bachelor's (2011) and Master's degrees (2014) in Mathematical Engineering from Politecnico di Milano, specializing in computational sciences and mathematical-physical modeling. In 2017, he earned a Ph.D. in Mathematics from a consortium of Chilean universities, focusing on algorithms for optimization in Hilbert spaces and optimal control of PDEs. Following his Ph.D., from 2018 he pursued postdoctoral research at ENSICaen (Université de Normandie, France) and from 2019 at the Italian Institute of Technology (IIT, Genoa, Italy), working on optimization algorithms, iterative regularization, and machine learning. In 2022, he was appointed as an Assistant Professor at the Mathematics Department (DIMA) of the University of Genoa, where he is part of the Machine Learning Genoa Center (MaLGa, LCSL Unit). His research focuses on the design and analysis of optimization algorithms for machine learning, with an emphasis on accuracy, stability, and computational efficiency.

Carlos Ocampo-Martinez received his Ph.D. degree in Control Engineering from the Universitat Politècnica de Catalunya (UPC) at Barcelona, Spain. From 2007 to 2010, he held postdoctoral positions at the University of Newcastle (Australia) and at the Institut de Robòtica i Informàtica Industrial, CSIC-UPC (IRI). Since 2011, he is with UPC, Automatic Control Department (ESAII). From 2014 to 2018, he has been also Deputy Director of IRI, a Joint Research Center of UPC and the Spanish National Research Council. His main research interests include constrained model predictive control, large-scale systems management (partitioning and non-centralized control), process control and industrial applications (mainly related to the key scopes of water and energy, and smart manufacturing under the IoT framework).

Carme Torras (www.iri.upc.edu/people/torras) is Research Professor at the Institut de Robòtica i Informàtica Industrial (CSIC-UPC) in Barcelona, where she heads a research group on assistive robotics. She has supervised 23 Ph.D. theses and led 17 European projects, among them her ERC Advanced Grant project CLOTHILDE – Cloth manipulation learning from demonstrations. Prof. Torras has participated in the review panels of the Swiss National Centre of Excellence in Robotics, Istituto Italiano di Tecnologia, German Research Foundation, and European Research Council. She has been Senior Editor of the IEEE Transactions on Robotics, and IEEE RAS Associate VP for Publications. Prof. Torras is IEEE, EurAI and ELLIS Fellow, member of Academia Europaea, the Royal Academy of Sciences and Arts of Barcelona, and the Royal Engineering Academy of Spain. Her research achievements have been recognized with several distinctions, among which the Catalan National Research Award (2020), and the Spanish National Research Prize in mathematics and ICT (2020), the highest research recognition in Spain. Committed to promoting ethics in the deployment of digital technologies, she has developed freely available online materials to teach a course on "Ethics in Social Robotics and AI" based on her science fiction novel The Vestigial Heart (MIT Press, 2018). She is vice-president of CSIC's Ethics Committee.

Lorenzo Rosasco is a professor at the University of Genova. He is a research affiliate at the Massachusetts Institute of Technology (MIT) and a visiting scientist at the Italian Technological Institute (IIT). He is a founder and serves as a coordinator of the Machine Learning Genova center (MaLGa) and the Laboratory for Computational and Statistical Learning, focusing on the theory, algorithms, and applications of machine learning. He obtained his Ph.D. in 2006 from the University of Genova and was a visiting student at the Center for Biological and Computational Learning at MIT, the Toyota Technological Institute at Chicago (TTI-Chicago), and the Johann Radon Institute for Computational and Applied Mathematics. From 2006 to 2013, he worked as a postdoc and research scientist at the Brain and Cognitive Sciences Department at MIT. He is a fellow at Ellis and serves as the co-director of the "Theory, Algorithms and Computations of Modern Learning Systems" program as well as the Ellis Genoa unit. Lorenzo has received several awards, including an ERC consolidator grant.