

Experimental Applications of the Koopman Operator in Active Learning for Control

Thomas A. Berrueta, Ian Abraham, and Todd Murphey

Abstract Experimentation as a setting for learning demands adaptability on the part of the decision-making system. It is typically infeasible for agents to have complete *a priori* knowledge of an environment, their own dynamics, or the behavior of other agents. In order to achieve autonomy in robotic applications, learning must occur incrementally, and ideally as a function of decision-making by exploiting the underlying control system. Most artificial intelligence techniques are ill-suited for experimental settings because they either lack the ability to learn incrementally or do not have information measures with which to guide their learning. This chapter examines the Koopman operator, its application in active learning, and its relationship to alternative learning techniques, such as Gaussian processes and kernel ridge regression. Additionally, examples are provided from a variety of experimental applications of the Koopman operator in active learning settings.

1 Introduction

Machine learning techniques describe complex patterns in data. Many learning methods, such as neural networks and convolutional networks, process large datasets in order to develop models that describe causal relationships between data. However, what happens when one must gather data as a part of a real-time process? How should one incorporate new data into models, and then leverage agency to further improve them? The answers to these questions form the basis of the robotic paradigm—*sense, plan, act*—and are at the heart of what defines autonomy.

In robotics, one is interested in endowing physical agents with the ability to reason about uncertainty and to *act* in order to achieve a goal. The experimental nature

Authors are with the Mechanical Engineering Dept. at Northwestern University, Evanston, IL

e-mail: tberrueta@u.northwestern.edu, i-abr@u.northwestern.edu,
t-murphey@northwestern.edu

of this process necessitates learning frameworks capable of incremental adaptation, such that an agent is able to gather new information and use it to its advantage. We refer to the process in which an agent selectively seeks out new information relevant to better achieving a goal as *active learning*.

Not all machine learning techniques are well-suited for active learning. Batch learning processes requiring large amounts of training data are not useful in real-time settings where learning is occurring at all times. For active learning purposes, we need the ability to incrementally update models to avoid recomputing the learned model from scratch, which can be prohibitive. In order to develop principled approaches for incremental information acquisition, we need to quantify the model’s information gain. We use the approximate Koopman operator as formulated in [41] and used in our earlier work [2] as a model of choice, and apply it in active learning settings.

While Koopman operators are the primary focus of this chapter, techniques such as generalized linear regression, kernel ridge regression, and Gaussian processes are often also applied as online learning tools in experiments. Specifying the theoretical and practical relationships between these methods is important to understanding the qualities of good active learning techniques. By developing a set of assumptions under which these methods are equivalent, one can transfer techniques developed for one method to another.

In Section 2, we describe the Koopman operator as a technique for active learning. In Section 3, we consider generalized linear regression, kernel ridge regression, and Gaussian processes as alternative techniques for experimental learning, and compare them to approximate Koopman operators, both in theoretical terms and computational implications. In Section 4, we illustrate the Koopman operator’s use in experimental settings. In Section 5, we conclude with a discussion of what makes the Koopman operator a good system representation in active learning settings.

2 Koopman Operator

The Koopman operator is an infinite-dimensional linear operator capable of describing the time-evolution of any observable dynamical system [22]. We consider discrete-time dynamical systems described by

$$\mathbf{x}_{k+1} = S(\mathbf{x}_k) = \mathbf{x}_k + \int_{t_k}^{t_k+\Delta t} F(\mathbf{x}(\tau))d\tau, \tag{1}$$

where the F is the vector field of the continuous dynamical system, which evolves on a state-space manifold \mathcal{M} such that $S : \mathcal{M} \rightarrow \mathcal{M}$. The right-hand side is a continuous-time system discretized with sampling interval Δt .

The Koopman operator is a linear operator that describes the time-evolution of basis functions instead of states. This shifts the system representation from state-space to an infinite-dimensional function space. Figure 1 shows the relationship




Fig. 1 Schematic of the relationship between a nonlinear dynamical system and its corresponding Koopman operator representation. Given a nonlinear system in n -dimensional state-space as shown in (a), the Koopman operator is a linear representation of the system in an infinite-dimensional function-space, shown in (b). The map Ψ lifts the original system onto an infinite-dimensional function-space. In this function-space, the Koopman operator, U , is a linear transformation representing the system dynamics.

between a dynamical system $S(\mathbf{x})$ and its corresponding Koopman operator U . For a detailed analysis of the Koopman operator and its properties, we refer the reader to Chapter 0, as well as [7].

2.1 EDMD Approximation

The Koopman operator itself is not a learning tool; it is an alternative representation of dynamical systems. However, numerically synthesizing an approximation of this representation in finite dimensions from data is a learning problem. In order to implement Koopman operators in computational applications, one generates finite-dimensional approximations. While certain systems can be represented by an exact, closed-form, finite-dimensional Koopman operator, this is generally not the case [20]. Dynamic Mode Decomposition (DMD) [33] and Extended Dynamic Mode Decomposition (EDMD) [41] are methods developed for approximating Koopman operators in finite dimensions and are commonly applied in fluid dynamics as analytical tools [26].

EDMD synthesizes a data-driven Koopman operator by developing a mapping that tries to span a finite-dimensional subspace of a given function space using a finite set of basis functions. Given basis functions $\Psi(\mathbf{x}) = [\psi_1(\mathbf{x}), \dots, \psi_N(\mathbf{x})]^T$, *s.t.* $\{\psi_i(\mathbf{x}) : \mathcal{M} \rightarrow \mathbb{R}, \forall i\}$ and a dataset of sequential observations $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_M]$, we can generate an approximate Koopman operator, \mathbf{U} , to describe the evolution of the lifted system

$$\Psi(\mathbf{x}_{k+1}) = \mathbf{U}\Psi(\mathbf{x}_k) + \varepsilon, \quad (2)$$

such that $\varepsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2)$, where \mathcal{N} is a multivariate Gaussian distribution with mean and variance $\mathbf{0}, \sigma^2 \in \mathbb{R}^N$. The corresponding approximate Koopman operator is given by the solution to the optimization

$$\min_{\mathbf{U}} \frac{1}{2} \sum_{k=1}^{M-1} \|\Psi(\mathbf{x}_{k+1}) - \mathbf{U}\Psi(\mathbf{x}_k)\|^2 \quad (3)$$

which has closed form solution

$$\mathbf{U} = \mathbf{A}\mathbf{G}^\dagger, \quad (4)$$

where \dagger denotes the Moore-Penrose pseudoinverse and the individual matrix components are

$$\begin{aligned} \mathbf{G} &= \mathbf{G}[M] = \frac{1}{M} \sum_{k=1}^{M-1} \Psi(\mathbf{x}_k)\Psi(\mathbf{x}_k)^T \\ \mathbf{A} &= \mathbf{A}[M] = \frac{1}{M} \sum_{k=1}^{M-1} \Psi(\mathbf{x}_{k+1})\Psi(\mathbf{x}_k)^T. \end{aligned} \quad (5)$$

2.2 Incremental Updates

The closed-form solution in (5) is of particular interest because it can be modified for incremental updates [1, 18]. We can define a set of difference equations describing an incremental update to the Koopman operator after a new observation

$$\begin{aligned} \mathbf{G}[M+1] &= \mathbf{G}[M] + \frac{1}{M+1} (\Psi(\mathbf{x}_M)\Psi(\mathbf{x}_M)^T - \mathbf{G}[M]) \\ \mathbf{A}[M+1] &= \mathbf{A}[M] + \frac{1}{M+1} (\Psi(\mathbf{x}_{M+1})\Psi(\mathbf{x}_M)^T - \mathbf{A}[M]), \end{aligned} \quad (6)$$

where we can calculate the updated Koopman operator with

$$\mathbf{U}[M+1] = \mathbf{A}[M+1](\mathbf{G}[M+1])^\dagger. \quad (7)$$

This formulation does not scale in complexity with the number of measurements since the pseudoinverse computation complexity scales with respect to the number of basis functions considered. The difference equation (6) describes an implicit cumulative average of all measurements. However, it is possible to perform incremental updates to the Koopman operator using other difference equations, such as moving average filters, exponential moving average filters or Kalman filters [16, 14].

2.3 Formulation for Control

In order to apply the Koopman operator in experimental settings, we can formulate the operator for real-time control [6]. By considering control inputs as a part of the measurement vector, we can split the basis functions $\Psi(\mathbf{x}, \mathbf{u})$ into $\Psi(\mathbf{x}, \mathbf{u}) = [\Psi_x(\mathbf{x})^T, \Psi_u(\mathbf{x}, \mathbf{u})^T]^T$, where $\{\Psi_x(\mathbf{x}) : \mathcal{M} \rightarrow \mathbb{R}^{N_x}\}$ and $\{\Psi_u(\mathbf{x}, \mathbf{u}) : \mathcal{M} \rightarrow \mathbb{R}^{N_u}\}$ such that $N = N_x + N_u$. As a result, the corresponding data-driven Koopman operator can be split into submatrices

$$\mathbf{U} = \begin{bmatrix} \mathbf{K}_x & \mathbf{K}_u \\ \mathbf{K}_{ux} & \mathbf{K}_{uu} \end{bmatrix}. \quad (8)$$

Using these submatrices we can formulate the time-evolution of state observables with control [2]

$$\Psi_x(\mathbf{x}_{k+1}) = \mathbf{K}_x \Psi_x(\mathbf{x}_k) + \mathbf{K}_u \Psi_u(\mathbf{x}_k, \mathbf{u}_k). \quad (9)$$

Additionally, when considering the evolution of observables that depend on input, the following equation holds

$$\Psi_u(\mathbf{x}_{k+1}, \mathbf{u}_{k+1}) = \mathbf{K}_{ux} \Psi_x(\mathbf{x}_k) + \mathbf{K}_{uu} \Psi_u(\mathbf{x}_k, \mathbf{u}_k). \quad (10)$$

It is worth noting a couple special cases of (9). For the case in which $\Psi_u(\mathbf{x}, \mathbf{u})$ is linear in \mathbf{u} the equation becomes

$$\Psi_x(\mathbf{x}_{k+1}) = \mathbf{K}_x \Psi_x(\mathbf{x}_k) + \hat{\mathbf{K}}_u \mathbf{u}_k, \quad (11)$$

where $\hat{\mathbf{K}}_u = \mathbf{K}_u \frac{\partial \Psi_u(\mathbf{x}_k, \mathbf{u}_k)}{\partial \mathbf{u}}$. Then for the case in which $\Psi_u(\mathbf{x}, \mathbf{u})$ is linear in \mathbf{u} and does not depend on the state the equation is

$$\Psi_x(\mathbf{x}_{k+1}) = \mathbf{K}_x \Psi_x(\mathbf{x}_k) + \mathbf{K}_u \mathbf{u}_k. \quad (12)$$

This formulation in (12) of the Koopman operator for control systems enables us to easily incorporate it into classical control schemes such as optimal control, or other model-predictive frameworks [23]. Alternative data-driven formulations of the Koopman operator for control have been described in the work of [20].

In optimal control, we seek to specify control laws that best drive systems towards a desired goal state with respect to an objective. For a discrete-time linear quadratic (LQ) control problem with goal state \mathbf{x}_d , we can define an objective function of the following form over the iterate k

$$J = \sum_{k=0}^{\infty} (\Psi_x(\mathbf{x}_k) - \Psi_x(\mathbf{x}_d))^T \mathbf{Q} (\Psi_x(\mathbf{x}_k) - \Psi_x(\mathbf{x}_d)) + \mathbf{u}_k^T \mathbf{R} \mathbf{u}_k, \quad (13)$$

where \mathbf{Q} and \mathbf{R} are positive semi-definite matrices that specify weights on system states and control effort, respectively. Deriving an optimal solution to this LQ problem leads to a feedback law

$$\mathbf{u}_{LQk} = -\mathbf{F}_{LQk} (\Psi_x(\mathbf{x}) - \Psi_x(\mathbf{x}_d)) \quad (14)$$

such that the optimal feedback gain \mathbf{F}_{LQK} is

$$\mathbf{F}_{LQK} = (\mathbf{R} + \mathbf{K}_u^T \mathbf{P} \mathbf{K}_u)^{-1} \mathbf{K}_u^T \mathbf{P} \mathbf{K}_x, \quad (15)$$

where \mathbf{P} is the solution to the discrete-time algebraic Ricatti equation. Through this process, we can solve optimal control problems with data-driven Koopman models [6, 20].

Additionally, given Koopman operator dynamics of the same form as (12), we note that system derivatives with respect to basis functions of state and control are linear:

$$\frac{d}{d\Psi_x}(\Psi_x(\mathbf{x}_{k+1})) = \mathbf{K}_x, \quad \frac{d}{d\mathbf{u}}(\Psi_x(\mathbf{x}_{k+1})) = \mathbf{K}_u. \quad (16)$$

The quality of a data-driven model is dependent on the amount of information captured about the underlying system. In order to determine how much information is captured by system models, we must develop a means of quantifying information with respect to the model.

2.4 Information Measures

Active learning is a process by which an agent can acquire more information about a task and improve its performance. Through experimental learning, we are interested in improving our system models such that we maximize task performance. In order to actively acquire information, we must develop information measures with respect to our system model. Fisher information provides a means of measuring the amount of information that a random variable carries about parameters [10]. Given a Koopman operator submatrix \mathbf{K}_x from a system specified as (9), we model each of its elements as a normally distributed parameter $\mathbf{K}_x^{(i,j)} \sim \mathcal{N}(\mu_{ij}, \Sigma_{ij}^2)$, and compute the Fisher information

$$I(\Psi_x(\mathbf{x})) = \frac{\partial \Gamma^T}{\partial \bar{k}} \Sigma^{-1} \frac{\partial \Gamma}{\partial \bar{k}}, \quad (17)$$

where $\bar{k} = \{\mathbf{K}_x^{(i,j)} : (i, j) \in \{1, \dots, N_x\} \times \{1, \dots, N_x\}\}$, $\Gamma \in \mathbb{R}^{N_x}$ is the measurement model defined by $\Gamma = \mathbf{K}_y \Psi_x(\mathbf{x})$, and $\Sigma \in \mathbb{R}^{N_x \times N_x}$ is assumed to be constant. The matrix Σ is the measurement covariance that can be empirically estimated from sample data *a priori*. The matrix $\frac{\partial \Gamma}{\partial \bar{k}} \in \mathbb{R}^{N_x \times (N_x \cdot N_x)}$ encodes measurement sensitivities to all elements of the Koopman operator. The sensitivity matrix takes the form

$$\frac{\partial \Gamma}{\partial \bar{k}} = \begin{bmatrix} \Psi_x(\mathbf{x})^T & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \Psi_x(\mathbf{x})^T & \dots & \mathbf{0} \\ \vdots & & \ddots & \\ \mathbf{0} & \dots & & \Psi_x(\mathbf{x})^T \end{bmatrix}, \quad (18)$$

where $\mathbf{0} \in \mathbb{R}^{1 \times N_x}$. We now optimize the trace of the Fisher information matrix, also known as T-optimality [36], to approximate information gain

$$\text{tr}(I(\Psi_x(\mathbf{x}))) = \sum_{i=1}^{N_x} \|\Psi_x(\mathbf{x})\|_{\Sigma_i^{-1}}^2. \quad (19)$$

We introduce (19) into a control objective by assigning a cost q to the inverse of the Fisher information

$$J = \sum_{k=1}^{\infty} q \text{tr}(I(\Psi_x(\mathbf{x}_k)))^{-1} + \|\Psi_x(\mathbf{x}_k) - \Psi_x(\mathbf{x}_d)\|_{\mathbf{Q}}^2 + \|\mathbf{u}_k\|_{\mathbf{R}}^2. \quad (20)$$

Using the cost functional in (20) we synthesize information maximizing trajectories and continuously improve our Koopman operator representation of the dynamics, as shown in [1]. The simplicity of the Koopman operator's measurement model Γ makes it easy to develop information measures with respect to the generated model.

2.5 Formulation for Hybrid Systems

Robotic systems often have discontinuous dynamics because of impacts and intermittent contact, leading to hybrid dynamics. Hybrid system dynamics are defined piecewise, where an indicator function determines which hybrid mode will evolve the current state. These dynamics are formulated as

$$S(\mathbf{x}) = \begin{cases} S_1(\mathbf{x}), & \text{if } \Phi(\mathbf{x}) = 1 \\ \vdots & \vdots \\ S_i(\mathbf{x}), & \text{if } \Phi(\mathbf{x}) = i \\ \vdots & \vdots \\ S_B(\mathbf{x}), & \text{otherwise} \end{cases} \quad (21)$$

where each $S_i(\mathbf{x})$ represents the dynamics of states $\mathbf{x} \in \mathbb{R}^n$ at mode $\Phi(\mathbf{x}) = i$. Additionally, at each hybrid mode boundary there is a discontinuous map between modes. Assuming that $\Phi(\mathbf{x})$ is known *a priori*, it is possible to compute a Koopman operator \mathbf{U}_i to represent each hybrid mode. We can specify a hybrid Koopman operator using

$$\bar{\mathbf{U}}(\mathbf{x}), \bar{\Psi}(\mathbf{x}) = \begin{cases} \mathbf{U}_1(\mathbf{x}), \Psi_1, & \text{if } \Phi(\mathbf{x}) = 1 \\ \vdots & \vdots \\ \mathbf{U}_i(\mathbf{x}), \Psi_i, & \text{if } \Phi(\mathbf{x}) = i \\ \vdots & \vdots \\ \mathbf{U}_B(\mathbf{x}), \Psi_B, & \text{otherwise} \end{cases} \quad (22)$$

where we characterize the discontinuous jumps between hybrid modes by collecting data at the mode boundaries and using $\Psi_i^*(\mathbf{x}_k^+) = \Psi_i^*(\mathbf{x}_k^-)\mathbf{U}_i^*$. Given that these jumps in state often take place over very small time scales, it may be difficult to gather enough data to provide a robust Koopman operator for that mapping. Active excitation methods might be necessary in order to properly characterize transitions. Alternative treatments of the Koopman operator have looked into modeling and analyzing switched [31], hybrid [15], and distributed systems [17].

The Koopman operator’s linearity, ability to update incrementally, and ease of generating information measures make it a good choice of data-driven model representation for experimental settings. However, the operator’s performance is dependent on the suitability of the chosen basis functions, which may not be known ahead of time.

Though the model representation is different, data-driven Koopman operator synthesis, as formulated in EDMD, solves a similar least-squares optimization as alternative regression models, such as generalized linear regression, kernel ridge regression, and Gaussian processes. Characterizing the conditions under which these models are equivalent to one another can help expand understanding of both Koopman operators synthesis and alternative learning methods. Under certain sets of assumptions, methods developed for Koopman operators synthesized via EDMD apply to the alternatives, and *vice versa*.

3 Alternative Methods

While the Koopman operator is neither a function approximation technique nor a regression tool, EDMD, as formulated in Section 2.1, is both [41]. In Section 2 we proposed that the EDMD approximation of the Koopman operator is well-suited for active learning. In principle, other function approximation techniques can be applied to generate finite Koopman operators. However, EDMD is the most common Koopman operator synthesis technique, and as such it is of interest to explicitly relate it to alternative machine learning techniques.

We show that under certain assumptions generalized linear regression, kernel ridge regression, and Gaussian processes solve the same underlying optimization problem as EDMD and can be formulated similarly. By understanding these relationships, we can apply techniques formally developed for one method to others, and *vice versa*. Through these equivalences we can extend the active learning techniques developed for the Koopman operator in Section 2 to other methods. We first provide a review of generalized linear regression, kernel ridge regression, and Gaussian processes, then establish the relationships between them.

3.1 Generalized Linear Regression

Linear regression and online linear regression have long been applied for statistical analysis of sequential models [34]. Linear regression models are typically limited to describing a function $\mathbf{y} = f(\mathbf{x})$, *s.t.* $\mathbf{x}, \mathbf{y} \in \mathbb{R}^S$ as linear combinations of inputs $\mathbf{y} = \mathbf{W}^T \mathbf{x}$, where we find $\mathbf{W} \in \mathbb{R}^{S \times S}$ to minimize an error metric. Generalized linear regression extends this approach to consider nonlinear functions of state, such that we are performing a linear regression in a given feature space defined by the basis functions, where the features are the basis functions. Given a dataset with inputs $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_M]$, outputs $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_M]$, and a set of basis functions $\Psi(\mathbf{x}) = [\psi_1(\mathbf{x}), \dots, \psi_N(\mathbf{x})]^T$ *s.t.* $\{\psi_i(\mathbf{x}) : \mathbb{R}^S \rightarrow \mathbb{R}, \forall i\}$, we consider the set of candidate functions defined by linear combinations of the basis functions subject to zero-mean Gaussian noise ε with variance σ^2

$$\mathbf{y} = \mathbf{W}^T \Psi(\mathbf{x}) + \varepsilon. \quad (23)$$

We can minimize over the least-squares error functional

$$\min_{\mathbf{W}} \frac{1}{2} \sum_{k=1}^{M-1} \|\mathbf{y}_k - \mathbf{W}^T \Psi(\mathbf{x}_k)\|^2 \quad (24)$$

to obtain a closed-form solution using maximum likelihood estimation [5]. We define the design matrix as

$$\Phi = [\Psi(\mathbf{x}_1), \dots, \Psi(\mathbf{x}_M)]^T \in \mathbb{R}^{M \times N} \quad (25)$$

which allows us to express the closed-form solution of (24) as

$$\mathbf{W} = (\Phi^T \Phi)^{-1} \Phi^T \mathbf{Y}, \quad (26)$$

where $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_M]^T$. We note that (26) is equivalent to

$$\mathbf{W} = \sum_{k=1}^{M-1} (\Psi(\mathbf{x}_k) \Psi(\mathbf{x}_k)^T)^{-1} (\Psi(\mathbf{x}_k)^T \mathbf{y}_k). \quad (27)$$

We also note that for autoregressive models where $\mathbf{y}_k = \Psi(\mathbf{x}_{k+1})$

$$\Psi(\mathbf{x}_{k+1}) = \mathbf{W}^T \Psi(\mathbf{x}_k) + \varepsilon. \quad (28)$$

3.2 Kernel Ridge Regression

Kernel ridge regression reframes the generalized linear regression problem by considering infinite-dimensional feature spaces specified by kernels. Features are ob-

served properties or characteristics of a phenomenon [5]. A feature space is then the space formed by all features considered. A kernel is a generalized inner product between vectors in a feature space. Finite-dimensional kernels can be specified directly in terms of a choice of basis functions

$$k(\mathbf{x}, \mathbf{x}') = \langle \Psi(\mathbf{x}), \Psi(\mathbf{x}') \rangle, \quad (29)$$

where the basis functions are a nonlinear transformation into a feature space. Kernels can implicitly define inner products in infinite-dimensional spaces without having to directly project coordinates into the feature space. The most common infinite-dimensional kernel of choice is the radial basis function kernel, also known as the Gaussian kernel [5]

$$k(\mathbf{x}, \mathbf{x}') = e^{-\frac{1}{2\sigma^2} \|\mathbf{x} - \mathbf{x}'\|^2}. \quad (30)$$

The sense in which the Gaussian kernel represents an inner product in an infinite-dimensional space can be understood by examining its Taylor series expansion. The Gaussian kernel is shown below as an infinite-dimensional polynomial kernel, with $\sigma^2 = \frac{1}{2}$ for simplicity:

$$\begin{aligned} k(\mathbf{x}, \mathbf{x}') &= e^{-\|\mathbf{x} - \mathbf{x}'\|^2} \\ &= e^{-(\mathbf{x}^T \mathbf{x} + 2(\mathbf{x}^T \mathbf{x}') + (\mathbf{x}')^T \mathbf{x}')} \\ &= (e^{-\mathbf{x}^T \mathbf{x}})(e^{(\mathbf{x}')^T \mathbf{x}'}) \sum_{n=0}^{\infty} \frac{(2\mathbf{x}^T \mathbf{x}')^n}{n!} \\ &= C(\mathbf{x}, \mathbf{x}') \sum_{n=0}^{\infty} \frac{2^n}{n!} \langle \mathbf{x}, \mathbf{x}' \rangle^n, \end{aligned} \quad (31)$$

where $C(\mathbf{x}, \mathbf{x}')$ is some constant value dependent on \mathbf{x} and \mathbf{x}' that does not affect the Taylor expansion. Equation (31) shows that the Gaussian kernel is an inner product over the infinite-dimensional space of polynomial functions, where polynomial functions of the input vector are considered as features.

Ridge regression, also known as L_2 -norm regularized generalized linear regression, extends the least-squares residual minimization problem described in (24) by including a regularization term over the model weights. Solutions with low weight magnitudes are given preference to prevent overfitting of the model. For a dataset with inputs $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_M]$ and outputs $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_M]$, the ridge regression optimization is given by

$$\min_{\mathbf{W}} \frac{1}{2} \sum_{k=1}^{M-1} \|\mathbf{y}_k - \mathbf{W}^T \Psi(\mathbf{x}_k)\|^2 + \frac{\lambda}{2} \|\mathbf{W}\|_F^2, \quad (32)$$

where $\lambda \geq 0$ is a regularization parameter specifying a cost on the magnitude of the model weights, and $\|\cdot\|_F$ represents the Frobenius matrix norm.

Using the design matrix Φ in (25), we can write down a closed-form solution to the minimization in (32)

$$\mathbf{W} = (\Phi^T \Phi + \lambda \mathbf{I})^{-1} \Phi^T \mathbf{Y}. \quad (33)$$

We refer to this solution as the primal solution to the ridge regression estimation problem.

Kernel ridge regression is based on an alternative representation of the solution to the ridge regression problem that enables application with infinite-dimensional kernels. The primal solution to the ridge regression problem scales in complexity with the number of basis functions used, which makes it unable to consider high-dimensional feature spaces. However, we can manipulate (33) to get an equivalent representation, which we will refer to as the dual solution

$$\mathbf{W} = \Phi^T (\Phi \Phi^T + \lambda \mathbf{I})^{-1} \mathbf{Y}. \quad (34)$$

We define $\mathbf{B} = (\mathbf{G}_r + \lambda \mathbf{I})^{-1} \mathbf{Y}$ where \mathbf{G}_r is the Gram matrix, such that $\mathbf{W} = \Phi^T \mathbf{B}$. The Gram matrix is a matrix containing all inner products between points in a dataset as specified by a given kernel $k(\mathbf{x}, \mathbf{x}')$, such that $\mathbf{G}_r^{(i,j)} = k(\mathbf{x}_i, \mathbf{x}_j)$, $\forall i, j \in \{1, \dots, M\}$. In this case, $\mathbf{G}_r = \Phi \Phi^T$ due to the choice of kernel in (29). If we consider \mathbf{W} as the primal weight matrix, we can refer to \mathbf{B} as the dual weight matrix and use it to formulate our estimate as a function of the kernel choice [27]

$$\begin{aligned} \hat{\mathbf{y}} &= \Psi(\mathbf{x}')^T \mathbf{W} & (35) \\ &= \Psi(\mathbf{x}')^T \Phi^T \mathbf{B} \\ &= \mathbf{k}(\mathbf{x}')^T \mathbf{B} \\ &= \mathbf{k}(\mathbf{x}')^T (\mathbf{G}_r + \lambda \mathbf{I})^{-1} \mathbf{Y}, & (36) \end{aligned}$$

where $\mathbf{k}(\mathbf{x}') = \Phi \Psi(\mathbf{x}') = [k(\mathbf{x}_1, \mathbf{x}'), \dots, k(\mathbf{x}_M, \mathbf{x}')]^T$. The kernel ridge regression function prediction scales in complexity with the number of data points, and as such is not suitable for incremental learning. Equation (36) enables us to analyze the models in infinite-dimensional feature spaces using kernels, which is beneficial in different settings. This kernel-based representation is how kernel ridge regression is most often applied. However, kernel ridge regression can also be formulated recursively and applied in online learning [9, 12].

3.3 Gaussian Process Regression

Until this point, we have only considered non-probabilistic models for regression in active learning. In some settings, randomness and uncertainty are inherently present in state measurements and as such require models that incorporate uncertainty into their predictions. Stochastic processes, such as Gaussian processes, specialize in describing systems that probabilistically evolve over time [14].

A Gaussian process is a stochastic process where any collection of random variables drawn are jointly Gaussian. This is to say that any finite set of random variable observations will be distributed according to $f(\mathbf{x}_1, \dots, \mathbf{x}_M) \sim \mathcal{N}(\mu, \Sigma)$, *s.t.* $\mu \in$

\mathbb{R}^M , $\Sigma \in \mathbb{R}^{M \times M}$. We specify a Gaussian process by defining its mean and covariance functions, $m(\mathbf{x})$ and $k(\mathbf{x}, \mathbf{x}')$. Covariance functions describe a similarity measure between features in an infinite-dimensional function space, and are equivalent to positive-definite kernel functions [32]. Given a Gaussian process $f(\mathbf{x})$

$$\begin{aligned} f(\mathbf{x}) &\sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \\ m(\mathbf{x}) &= \mathbb{E}[f(\mathbf{x})] \\ k(\mathbf{x}, \mathbf{x}') &= \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))^T (f(\mathbf{x}') - m(\mathbf{x}'))]. \end{aligned} \quad (37)$$

Gaussian process regression specifies an approximated function by expressing a predictive distribution over a function space based on observations of realized random variables. This predictive distribution $p(f(\mathbf{x}')|\mathbf{X}, \mathbf{Y})$ is also known as a posterior, and can be computed given inputs $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_M]$ and outputs $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_M]$ through a Bayesian inference framework

$$\begin{aligned} \text{posterior} &= \frac{\text{prior} \times \text{likelihood}}{\text{marginal likelihood}} \\ p(f(\mathbf{x})|\mathbf{X}, \mathbf{Y}) &= \frac{p(f(\mathbf{x}))p(\mathbf{Y}|\mathbf{X}, f(\mathbf{x}))}{\int p(\mathbf{Y}|f(\mathbf{x}), \mathbf{X})p(f(\mathbf{x})|\mathbf{X})df(\mathbf{x})}. \end{aligned} \quad (38)$$

By exploiting properties of the Gaussian distribution one can avoid explicitly computing the posterior in (38). We use the fact that conditioning and marginalizations of Gaussian distributions are other Gaussians with closed-form solutions [32]. Without imposing any prior belief on the function space this problem is ill-posed. We will consider processes of the form $\mathbf{y}_k = f(\mathbf{x}_k) + \varepsilon$ with $\varepsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2)$. As a prior, we restrict the set of functions considered to $\mathbf{y} = \mathbf{W}^T \Psi(\mathbf{x}) + \varepsilon$ such that $\mathbf{W} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\mathbf{W}})$ with zero mean for notational simplicity. We can see that the mean and covariance functions are then

$$\begin{aligned} m(\mathbf{x}) &= \mathbb{E}[f(\mathbf{x})] = \mathbb{E}[\mathbf{W}^T] \Psi(\mathbf{x}) + \mathbb{E}[\varepsilon] = \mathbf{0} \\ k(\mathbf{x}, \mathbf{x}') &= \mathbb{E}[f(\mathbf{x})^T f(\mathbf{x}')] = \Psi(\mathbf{x})^T \mathbb{E}[\mathbf{W}\mathbf{W}^T] \Psi(\mathbf{x}') + \mathbb{E}[\varepsilon^T \varepsilon] = \Psi(\mathbf{x})^T \Sigma_{\mathbf{W}} \Psi(\mathbf{x}') + \sigma^2. \end{aligned} \quad (39)$$

Given a design matrix as defined in (25), one can use the mean and covariance functions to define a Gaussian processes $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$. To generate a predictive distribution, we then define a joint Gaussian distribution between the input dataset's design matrix Φ and a new observation $\Psi(\mathbf{x}')$

$$\begin{bmatrix} f(\mathbf{X}) \\ f(\mathbf{x}') \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{k}(\Phi, \Phi) & \mathbf{k}(\Phi, \Psi(\mathbf{x}')) \\ \mathbf{k}(\Phi, \Psi(\mathbf{x}'))^T & k(\Psi(\mathbf{x}'), \Psi(\mathbf{x}')) \end{bmatrix}\right), \quad (40)$$

where the design matrix autocovariance is $\mathbf{k}(\Phi, \Phi) \in \mathbb{R}^{M \times M}$, and the covariance between the design matrix and the new observation $\mathbf{k}(\Phi, \Psi(\mathbf{x}')) \in \mathbb{R}^M$. We then express these covariances in closed form

$$\begin{bmatrix} f(\mathbf{X}) \\ f(\mathbf{x}') \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \Phi \Sigma_{\mathbf{W}} \Phi^T + \sigma^2 \mathbf{I} & \Phi \Sigma_{\mathbf{W}} \Psi(\mathbf{x}') + \sigma^2 \mathbf{1} \\ (\Phi \Sigma_{\mathbf{W}} \Psi(\mathbf{x}'))^T + \sigma^2 \mathbf{1}^T & \Psi(\mathbf{x}')^T \Sigma_{\mathbf{W}} \Psi(\mathbf{x}') + \sigma^2 \end{bmatrix}\right), \quad (41)$$

where $\mathbf{1} \in \mathbb{R}^M$, and \mathbf{I} is the identity matrix. Equation (41) allows us to calculate a predictive posterior distribution for new observations of $f(\cdot)$ by marginalizing over the design matrix Φ and outputs \mathbf{Y} [39]. The predictive posterior distribution is then of the form

$$\begin{aligned} f(\mathbf{x}') | \Phi, \mathbf{Y} &\sim \mathcal{N}(\mathbb{E}[f(\mathbf{x}')], \mathbb{V}[f(\mathbf{x}')]) \\ \mathbb{E}[f(\mathbf{x}')] &= \Psi(\mathbf{x}')^T \Sigma_{\mathbf{W}} \Phi^T (\Phi \Sigma_{\mathbf{W}} \Phi^T + \sigma^2 \mathbf{I})^{-1} \mathbf{Y} \\ \mathbb{V}[f(\mathbf{x}')] &= \Psi(\mathbf{x}')^T \Sigma_{\mathbf{W}} \Psi(\mathbf{x}') - \Psi(\mathbf{x}')^T \Sigma_{\mathbf{W}} \Phi^T (\Phi \Sigma_{\mathbf{W}} \Phi^T + \sigma^2 \mathbf{I})^{-1} \Phi \Sigma_{\mathbf{W}} \Psi(\mathbf{x}'). \end{aligned} \quad (42)$$

This posterior distribution enables us to generate predictions for new observations, as well as specify a covariance for each point [32].

Although we have constructed our derivation using a finite-dimensional basis representation, Gaussian processes are typically implemented using infinite-dimensional kernels. By following the same procedure shown in Section 3.2, kernelization is trivial, and results in the following predictive distribution in terms of a specified kernel:

$$\begin{aligned} f(\mathbf{x}') | \Phi, \mathbf{Y} &\sim \mathcal{N}(\mathbb{E}[f(\mathbf{x}')], \mathbb{V}[f(\mathbf{x}')]) \\ \mathbb{E}[f(\mathbf{x}')] &= m(\mathbf{x}') + \mathbf{k}(\mathbf{x}')^T (\mathbf{G}_{\mathbf{r}} + \sigma^2 \mathbf{I})^{-1} (\mathbf{Y} - m(\mathbf{X})) \\ \mathbb{V}[f(\mathbf{x}')] &= k(\mathbf{x}', \mathbf{x}') - \mathbf{k}(\mathbf{x}')^T (\mathbf{G}_{\mathbf{r}} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}(\mathbf{x}'), \end{aligned} \quad (43)$$

where we have added the mean function $m(\mathbf{x})$ for generality. The case in which $m(\mathbf{x}) = \mathbf{0}$ is referred to as a *centered* Gaussian process. This kernelized representation of the Gaussian process is not well-suited for active learning settings. In order to overcome this limitation, there has been work in deriving local models for real-time applications [29].

3.4 Relationships Between Learning Techniques

Now that we have reviewed the foundations of generalized linear regression, kernel ridge regression, and Gaussian processes, we can examine their relationships to EDMD approximations of Koopman operators, as well as to one another. We develop assumptions for which these techniques are equivalent to one another.

3.4.1 Generalized Linear Regression

For finite-dimensional bases, generalized linear regression is most closely related to EDMD. By limiting the functions we are willing to model to autoregressive func-

tions, we can make the substitution $\mathbf{y}_k = \Psi(\mathbf{x}_{k+1})$ in (27). Then, if we note that $\mathbf{A}^\dagger = \mathbf{A}^{-1}$ when \mathbf{A} is square, we can restate (27) in relation to the Koopman operator matrix components in (5)

$$\begin{aligned} \mathbf{W} &= \sum_{k=1}^{M-1} (\Psi(\mathbf{x}_k)\Psi(\mathbf{x}_k)^T)^\dagger (\Psi(\mathbf{x}_k)^T\Psi(\mathbf{x}_{k+1})) \\ &= ((\mathbf{G}[M])^\dagger)^T (\mathbf{A}[M])^T = (\mathbf{A}[M](\mathbf{G}[M])^\dagger)^T \\ &= \mathbf{U}^T. \end{aligned} \tag{44}$$

Thus, generalized linear regression models are equivalent to finite-dimensional Koopman operators synthesized through EDMD for the same set of basis functions. This should not be surprising since Koopman operator synthesis via DMD [33, 38] and EDMD [41] is formulated as a regression problem.

The correspondence between these techniques indicates that all methods developed in Section 2 apply to generalized linear regression, and *vice versa*. This means that generalized linear regression models can be applied in control, represent hybrid systems, update incrementally, and have information measures easily generated, which makes them suitable for experimental settings.

3.4.2 Kernel Ridge Regression

Kernel ridge regression models are closely related to generalized linear regression models, and consequently Koopman operators. Ridge regression solves a slightly different optimization problem than EDMD, as shown in (32). However, for the case of $\lambda = 0$ the ridge regression optimization is equivalent to that of generalized linear regression and EDMD, as shown in (24).

For finite dimensional kernels of the form $k(\mathbf{x}, \mathbf{x}') = \langle \Psi(\mathbf{x}), \Psi(\mathbf{x}') \rangle$, *s.t.* $\Psi(\mathbf{x}) \in \mathbb{R}^n$, we have shown that the solutions in (34) and (33) are equivalent in Section 3.2. For $\lambda = 0$, (33) is equivalent to the generalized linear regression solution in (26), which we have shown in the previous section to be the same as that of EDMD. This is to say that the kernel ridge regression estimate can be equivalent to the Koopman operator's estimate derived in Section 2.1 for finite-dimensional kernels with a regularization coefficient of zero. In this limited representation, kernel ridge regression models can be computed in real time. However, it is important to note that kernel ridge regression without infinite-dimensional kernels or regularization is not how the technique is typically implemented, and does not take advantage of the compactness of kernel representations for high-dimensional data.

As shown by the work of [42], Koopman operators can be synthesized using infinite-dimensional kernel methods at the expense of growing computational complexity with additional samples. In this kernel-based representation, the Koopman operator estimate is equivalent to (36) for $\lambda = 0$. While the choice of finite-dimensional or infinite-dimensional kernel impacts online learning, the regulariza-

tion coefficient does not. Hence, regularization coefficients may be used to promote sparsity in Koopman operators, which may be desirable in some settings.

3.4.3 Gaussian Processes

Despite Gaussian processes being stochastic, they are closely related to kernel ridge regression. By examining the Gaussian processes' mean function in (43) we see that it is equivalent to the kernel ridge regression prediction in (35) given that we set $\sigma^2 = \lambda$, $\Sigma_{\mathbf{W}} = \mathbf{I}$, and assume the Gaussian process is centered, with \mathbf{W} as defined in (34). In order for the equivalence to hold between typical kernel ridge regression models and Gaussian processes, we ignore the process covariance estimate and sequential Bayesian model updates. However, recent work in Bayesian kernel ridge regression models has shown that they are equivalent to Gaussian processes when the Gaussian process covariance estimate is fixed [40].

The correspondence between Gaussian process regression models and Koopman operator synthesis follows from EDMD's relationship to kernel ridge regression as described in the previous section. By setting $\sigma^2 = 0$ and $\Sigma_{\mathbf{W}} = \mathbf{I}$ the mean estimate from the Gaussian process regression matches the EDMD solution. However, by doing this we remove all non-determinism from the function approximation which makes it fundamentally *not* a Gaussian process. We note that the role σ^2 plays is very similar to that of λ in kernel ridge regression, which indicates that it may be incorporated in the EDMD Koopman operator synthesis process. By increasing σ^2 one would prevent overfitting in the solution by promoting sparsity, and consequently admit more variance. Typical implementations of Gaussian processes scale in complexity with the number of data observed, and as such are not suitable for settings that demand incremental learning. Nonetheless, one can generate information measures for Gaussian processes with stationary mean and covariance functions [11].

Alternatively, generalized linear regression, kernel ridge regression, and DMD can be derived through a Bayesian inference framework [5, 37], which suggests that EDMD can be subject to a Bayesian treatment as well. It would be of interest to show a direct correspondence between Bayesian representations of the Koopman operator and Gaussian processes because of their relationship to neural networks. Given a Gaussian prior over model weights, [28] showed that a single-layer neural network with infinite bandwidth approximates a Gaussian process as a result of the central limit theorem. This work was recently extended to describe the connection between Gaussian processes and deep neural networks in [24]. Moreover, neural networks have been applied both to Koopman operator synthesis [43], as well as basis function dictionary learning [25]. The relationships outlined in this section between Gaussian processes, neural networks and Koopman operators suggest a deeper connection between these methods which remains an open problem.

	Generalized Linear Regression	Kernel Ridge Regression	Gaussian Processes
Koopman Operators (EDMD)	$\mathbf{U} = \mathbf{W}^T$	$\lambda = 0$ $k(\mathbf{x}, \mathbf{x}') = \langle \Psi(\mathbf{x}), \Psi(\mathbf{x}') \rangle$	$\sigma^2 = 0, \Sigma_{\mathbf{W}} = \mathbf{I}$ $k(\mathbf{x}, \mathbf{x}') = \langle \Psi(\mathbf{x}), \Psi(\mathbf{x}') \rangle$ Ignore covariance estimate

Table 1 Assumptions under which each alternative method presented in Section 3 is equivalent to the EDMD estimate of a finite Koopman operator. As formulated, generalized linear regression models are equivalent to the synthesized Koopman operator matrix. The kernel ridge regression prediction is equivalent to that of EDMD when the regression regularization coefficient is 0, and only for finite-dimensional kernels. The mean function estimate of a Gaussian process is equivalent to an EDMD Koopman operator prediction when $\sigma^2 = 0, \Sigma_{\mathbf{W}} = \mathbf{I}$, for finite-dimensional kernels.

3.4.4 Summary

The assumptions under which generalized linear regression, kernel ridge regression, and Gaussian processes produce the same solutions to an estimation problem as EDMD are summarized in Table 1. Approximate Koopman operators as formulated in Section 2.1 are equivalent to generalized linear regression models. Kernel ridge regression models can be made equivalent to a Koopman operator model under a set of simplifying assumptions for finite kernel representations, as well as for infinite-dimensional kernels by sacrificing online computation. Finally, the relationship between Gaussian processes and EDMD approximations of Koopman operators follows from its correspondence with kernel ridge regression in deterministic settings.

Although Gaussian processes and kernel ridge regression are often applied in real-time settings, in their typical formulations they are unable to be updated incrementally and are ill-suited for active learning problems. However, generalized linear regression models are compatible with experimental settings because of their direct correspondence with EDMD. The correspondences between methods shown throughout this section imply that there should exist settings in which learning techniques such as Gaussian processes and deep neural networks can be made to learn incrementally and with respect to an information measure. Additionally, the relationship between Gaussian processes and EDMD suggest that there should be a way to formulate Koopman operator synthesis to incorporate uncertainty estimates for predictions.

3.5 Example

Though the methods presented in this section are formulated through alternative means, they all solve similar optimization problems. The primary difference between these methods is representation, which can lead to *the formal optimization being the same, but the implementation being different*. Under a set of assumptions

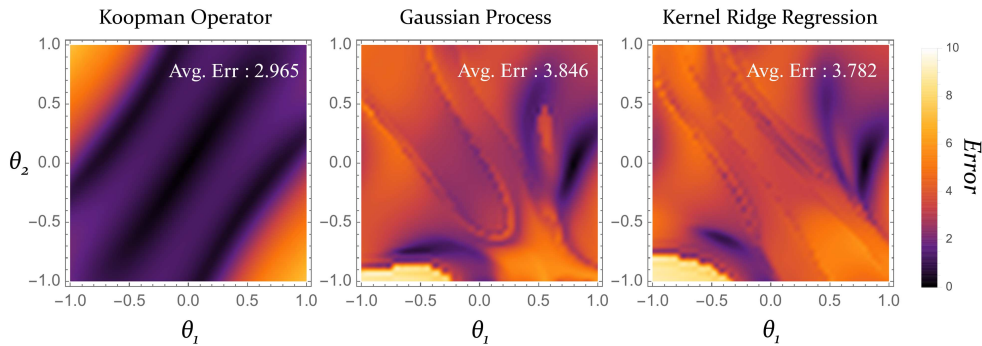


Fig. 2 Performance comparison between EDMD, Gaussian processes, and kernel ridge regression. Each model was trained on the same dataset collected from a 4s trajectory of the free dynamics of a double pendulum system with initial condition $(\theta_1, \theta_2, \dot{\theta}_1, \dot{\theta}_2) = (0.8, 0, 0, 0)$. Then, we calculated the integrated mean squared error over a 3s prediction from each model of the double pendulum dynamics over the entire $\{(\theta_1, \theta_2) : [-1, 1] \times [-1, 1]\}$ domain. Standard implementations of each model were used so as to compare performance under typical usage.

we are able to relate the model representations of kernel ridge regression, Gaussian processes, and EDMD. Although this is technically the case, these assumptions do not represent typical use cases for these methods.

As a comparison between standard implementations of these methods, we generate three independent data-driven models of the dynamics of a double pendulum system—one with Koopman operators, one with Gaussian processes, and one with kernel ridge regression. We do not implement generalized linear regression because we showed it to produce an equivalent model representation to the EDMD solution in Section 3.4.1. We collected a training dataset consisting of a single 4s trajectory taken from the double pendulum free dynamics sampled at 100Hz with an initial condition of $(\theta_1, \theta_2, \dot{\theta}_1, \dot{\theta}_2) = (0.8, 0, 0, 0)$. We use the training dataset to instantiate a Koopman operator, kernel ridge regression model, and a Gaussian process model. The Koopman operator model is calculated using a second-order polynomial basis of the double pendulum states, while both the kernel ridge regression model and the Gaussian process use the Gaussian kernel shown in (30). The regularization and variance parameters were selected by the Scikit-Learn software package [30] via Bayesian hyperparameter optimization [35]. Each model is then used to predict double pendulum trajectories for a horizon of 3s from each (θ_1, θ_2) initial condition over the $\{(\theta_1, \theta_2) : [-1, 1] \times [-1, 1]\}$ domain, with zero initial velocity. Then, we calculate the integrated mean square error between each model’s prediction and the actual dynamics over the entire horizon for each initial condition.

Figure 2 depicts the results of the comparison. The Koopman operator model prediction has the lowest average error of the three models. We observe that both the Gaussian process model and kernel ridge regression model prediction errors are lowest at the initial conditions of the training trajectory $(\theta_1, \theta_2) = (0.8, 0)$, while the Koopman operator model generalizes more easily over the domain. The purpose



Fig. 3 Sphero SPRK robot used for active learning of system dynamics on sand.

of this comparison is to show that while these methods can be shown to be closely related, standard implementations will produce very different results.

4 Experimental Applications

In this section, we demonstrate experimental applications of the Koopman operator in a variety of examples. We highlight applications in which active learning enables rapid system identification through information-driven incremental model updates. Additionally, we show how the Koopman operator’s ability to represent hybrid systems enables the data-driven modeling of finite automata in experimental settings.

4.1 Learning Sphero SPRK Dynamics in Sand

We are interested in demonstrating the use of a Koopman operator applied to learning complex nonlinear dynamics in a real-time experimental setting. We use the Sphero SPRK shown in Fig. 3, which is a programmable robotic toy shaped like a small ball that can be teleoperated using a controller or app. It is actuated through an internal mechanism designed for rolling on flat ground. If we change the system dynamics by switching the rolling surface to sand the SPRK is incapable of rolling with its default controller. To enable the SPRK to roll on sand, we use the Koopman operator for active identification of the robot’s dynamics on sand via real-time information maximization as described in Section 2.4, and shown in [1].

We use a nonlinear model-predictive controller known as Sequential Action Control (SAC) [3] to generate information maximizing trajectories with respect to the information objective (20), which we use to update our Koopman operator model in real time using (7). Over a 20s period, we synthesize information maximizing trajectories in order to characterize system dynamics on sand. After this period, we

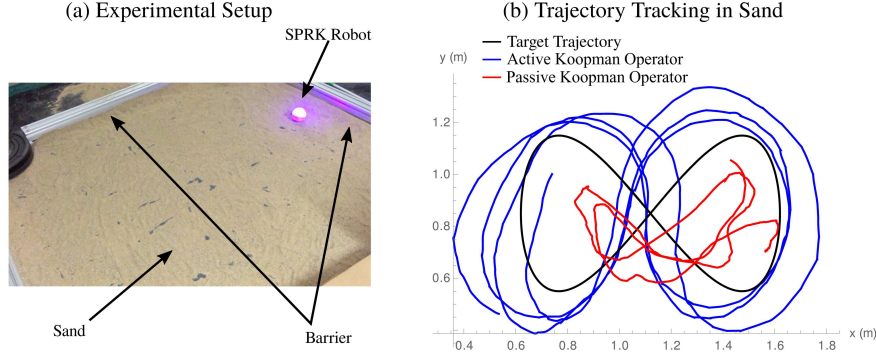


Fig. 4 Experimental application of active learning with Koopman operators for online learning of the Sphero SPRK dynamics on sand. (a) Experimental setup consisting of a boxed area covered in sand, with an overhead Xbox Kinect using OpenCV reporting state information. (b) Results from trajectory tracking experiments in sand. The Sphero SPRK was used to track a "figure 8 trajectory" (black), using both an active learning approach with a continuously updating Koopman operator (blue), and a pre-computed passive Koopman operator (red). While the fixed model establishes a baseline performance, the active learning Koopman operator model continuously improves tracking performance as the experiment progresses.

switch objectives to a tracking objective such as (13) where we attempt to trace a figure 8 trajectory in a sandbox while continuing to update the Koopman operator. We compare the actively learned Koopman operator performance to a passive Koopman operator. The passive Koopman operator was computed from 20s of data from tracing the figure 8 trajectory with open-loop control.


The experiment consists of a rectangular sandbox, with an Xbox Kinect mounted overhead equipped with OpenCV providing odometry information. The system states collected are $\mathbf{x} = [x, y, \dot{x}, \dot{y}]^T$ where (x, y) are coordinates in the Kinect field of vision and the corresponding velocities are estimated from finite differencing, and we define our control signals within this representation as $\mathbf{u} = [u_x, u_y]^T$. The set of basis functions used to model the SPRK robot were

$$\begin{aligned}\Psi_x(\mathbf{x}) &= [x, y, \dot{x}, \dot{y}, 1, \dot{x}^2, \dot{y}^2, \dot{x}^2\dot{y}, \dots, \dot{x}^3\dot{y}^3]^T \in \mathbb{R}^{18} \\ \Psi_u(\mathbf{u}) &= [u_x, u_y]^T \in \mathbb{R}^2,\end{aligned}\tag{45}$$

which are linear combinations of third order polynomial basis functions of the velocity states. The basis functions of control were chosen to be linear in \mathbf{u} and not dependent on state.

In Fig. 4 we show the experimental results resulting from the active and passive Koopman operator learning approaches. The online learning approach is much more effective at tracking the desired trajectory than the fixed learning approach. Moreover, we notice that as the experiment progresses, the active Koopman operator trajectories continue to improve. The Koopman operator synthesized through active learning required no *a priori* knowledge of the system dynamics, bootstrapped initial model guess, or specification of granular media physics. Despite the fact that

Fig. 5 Quadcopter vehicle model configuration in body-fixed frame for active learning of vehicle dynamics under emergency conditions. The quadcopter dynamics are defined such that one motor has an unspecified malfunction and can only work at 80% capacity.



this problem could have been solved by other means, this experimental application is facilitated by the Koopman operator model representation.

4.2 Rapid Quadcopter Stabilization with Active Learning

The Koopman operator's ability to actively acquire information in real-time allows us to approach problems that demand rapid reactive efforts. For settings in which large disturbances permanently affect system dynamics, we must be able to quickly develop new models to adapt to the disturbance while still achieving an underlying goal. We demonstrate the Koopman operator's quick adaptability by stabilizing a simulated quadcopter with a motor malfunction. We define the quadcopter malfunction as a single motor operating at 80% capacity. However, the malfunction is unmodeled and unknown to the Koopman operator ahead of time, which means it must be learned in real time in order to stabilize the quadcopter and prevent it from falling, as shown in [1]. The quadcopter configuration is shown in Fig. 5, and the system dynamics used to carry out this simulated experiment are

$$\begin{aligned}
 \dot{\mathbf{h}} &= \mathbf{h} \begin{bmatrix} \hat{\omega} & \mathbf{v} \\ \mathbf{0} & 0 \end{bmatrix} \\
 \mathbf{J}\dot{\omega} &= \mathbf{M} + \mathbf{J}\omega \times \omega \\
 \dot{\mathbf{v}} &= \frac{1}{m}F\mathbf{e}_3 - \omega \times \mathbf{v} - g\mathbf{R}^T\mathbf{e}_3,
 \end{aligned} \tag{46}$$

where the dynamics are defined on the Lie group $\mathbf{h} = (\mathbf{R}, \mathbf{p}) \in SE(3)$ in body-fixed coordinates, where $\hat{\omega}$ is the angular velocity tensor of the vector ω , \mathbf{v} is the vector of linear velocities. The matrix \mathbf{J} is the inertia tensor, m is the quadcopter mass, g is acceleration due to gravity and \mathbf{e}_3 is a body-fixed vector shown in Fig. 5.

The inputs to the system are a control force F and control torque \mathbf{M} , each of which depend on the control vector $\mathbf{u} = [u_1, u_2, u_3, u_4]$ according to

$$F = k_l(u_1^2 + u_2^2 + u_3^2 + u_4^2)$$

$$\mathbf{M} = \begin{bmatrix} k_l l(u_2^2 - u_4^2) \\ k_l l(u_3^2 - u_1^2) \\ k_m(u_1^2 - u_2^2 + u_3^2 - u_4^2) \end{bmatrix},$$

with k_l , k_m , l as model parameters. Details regarding implementation of the quadcopter dynamics are covered in [13].

The set of basis functions chosen directly embed the quadcopter dynamics, such that the Koopman operator has a bootstrapped model of the nominal system dynamics

$$\Psi_x(\mathbf{x}) = [\mathbf{a}_g, \boldsymbol{\omega}, \mathbf{v}, g(\mathbf{v}, \boldsymbol{\omega})]^T \in \mathbb{R}^{18}$$

$$\Psi_u(\mathbf{u}) = \mathbf{u} \in \mathbb{R}^4,$$

where $\mathbf{a}_g \in \mathbb{R}^3$ is the body-centered gravity vector, and \mathbf{v} , $\boldsymbol{\omega}$ are as defined above. The additional basis functions $g(\mathbf{v}, \boldsymbol{\omega})$ consist of all elements of the outer product between the linear and angular velocity vectors, and were chosen without any *a priori* knowledge of the motor malfunction specification

$$g(\mathbf{v}, \boldsymbol{\omega}) = [v_1 \omega_1, v_1 \omega_2, v_1 \omega_3, \\ v_2 \omega_1, v_2 \omega_2, v_2 \omega_3, \\ v_3 \omega_1, v_3 \omega_2, v_3 \omega_3].$$

The experimental procedure consists of setting the quadcopter in free-fall from some initial condition and for a period of 1s learning a Koopman operator online with respect to a specified objective. Active learning trials applied an information maximizing objective during the learning period, then switched to a stabilizing objective afterward. Passive learning trials followed a stabilizing objective for the entirety of each trial. After the 1s learning period we did not continue to update the operator online for either set of trials. A set of 100 Monte-Carlo trials over uniformly distributed initial conditions $\in \mathbb{R}^3$ were carried out. Figure 6 summarizes the experimental

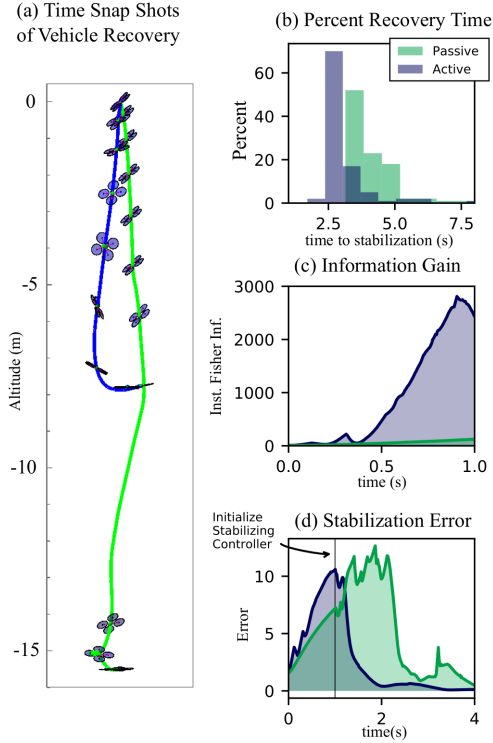


Fig. 6 Quadcopter vehicle free-fall simulation where one motor has an unmodeled malfunction. The goal is to learn to stabilize the quadcopter under two learning approaches. The first is a passive learning approach (green) where a Koopman operator is learned online over 1s while attempting to stabilize the quadcopter. The second is an active learning approach (blue) where a Koopman operator is generated while optimizing Fisher information, then the system switches objectives to achieve stabilization.

results. Figure 6A depicts timestamps from a pair of quadcopter trajectories under active and passive learning.

In Fig. 6B we display the results from the Monte-Carlo trials. The active Koopman operator stabilized the quadcopter within $2.5s$ in over 65% of trials, while the passive operator took another full second on average. Figure 6C shows the instantaneous information gain over the $1s$ learning period for both passive and active learning approaches, where the active learning approach acquires orders of magnitude more information than the passive approach in a selected trial. Finally, Fig. 6D depicts the stabilization error for a particular trial with passive and active learning. The active learning trial stabilizes the quadcopter more quickly than the passive learning approach and with lower error after the stabilizing objective switch.

Through this experiment, we have shown that the Koopman operator with active learning can be used in time-critical applications. We introduced an unmodeled motor malfunction into the quadcopter dynamics and learned a Koopman operator in real time in order to stabilize the system. We showed that by first executing information maximizing trajectories prior to attempting to stabilize we outperform passive online learning techniques.

4.3 Learning SLIP Hybrid Dynamics

Hybrid systems can be difficult to model because of inherent discontinuities in the system dynamics. These discontinuities often cannot be represented by a single model. Experimental learning in uncertain environments in the real world may demand learning techniques that model such discontinuities. In particular, field robotics often concerns itself with modeling locomotive systems, such as bipedal or quadrupedal walkers. Walking robots deal with the discontinuities of impact at each step of their locomotive cycle, and as such must have a way to reason about them. The simplest dynamical walker model is the spring-loaded inverted pendulum (SLIP) model, shown in Fig. 7.

The SLIP system dynamics are split into two hybrid modes: stance and flight. The transition between these hybrid modes is described by an indicator function, often referred to as a guard equation, which specifies the set of dynamics evolving the system at a given coordinate in state-space. The system states are $\mathbf{x} = [x, y, \dot{x}, \dot{y}, x_f]^T$, where (x, y) indicates the position of the mass in the configuration specified in Fig. 7, and x_f is the position of the SLIP model's foot. The model's control inputs are $\mathbf{u} = [u_s, u_f]^T$. The dynamics of the SLIP model are

$$f_{stance} = \begin{bmatrix} \dot{x} \\ \dot{y} \\ (k(l_0 - l(\mathbf{x})) + u_s) \frac{x - x_f}{ml(\mathbf{x})} \\ (k(l_0 - l(\mathbf{x})) + u_s) \frac{y}{ml(\mathbf{x})} - g \\ 0 \end{bmatrix}, \quad f_{flight} = \begin{bmatrix} \dot{x} \\ \dot{y} \\ 0 \\ -g \\ \dot{x} + u_f \end{bmatrix}, \quad (47)$$




Fig. 7 SLIP model configuration. The SLIP system obeys a set of hybrid dynamics corresponding to the separate stance and flight modes.

where l_0 is the resting length of the spring, k its stiffness, and $l(\mathbf{x}) = \sqrt{(x - x_f)^2 + y^2}$ is its variable length [21]. Then, the indicator function is $\Phi_{SLIP}(\mathbf{x}) = \text{sign}(1 - \frac{l_0}{l(\mathbf{x})}) \in \{-1, 1\}$. The complete hybrid dynamics are

$$f_{SLIP}(\mathbf{x}, \mathbf{u}) = \begin{cases} f_{stance}(\mathbf{x}, \mathbf{u}), & \text{if } \Phi_{SLIP}(\mathbf{x}) = -1 \\ f_{flight}(\mathbf{x}, \mathbf{u}), & \text{otherwise} \end{cases}. \quad (48)$$

We use the hybrid formulation of the Koopman operator as specified in Section 2.5, as well as the analytical indicator function to actively learn a hybrid Koopman operator representation of the SLIP model in simulation. For the first 10s we have a simulated SLIP hopper follow an information maximizing trajectory, where each hybrid Koopman mode incrementally updates according to the indicator function. Following the active learning period, the objective changes and the model is made to follow a forward hopping velocity of $4.2m/s$. Results are shown in Fig. 8 where the hybrid Koopman model first maximizes information about its dynamics for 10s, and then it moves forward tracking a trajectory. Figure 9 depicts the SLIP hopper's switch in objectives in terms of the model's forward velocity.

The Koopman operator's hybrid formulation enables modeling of discontinuous dynamical systems, which is of great importance in the field of locomotion, but also field robotics at large. Furthermore, the experiment presented an active learning framework for characterizing hybrid dynamical systems with known guard equations that could be extended to other walking robots.





Fig. 8 Active identification of simulated SLIP system dynamics. The SLIP system first obeys an information maximizing objective for 10s, and then the objective is changed to following a forward velocity trajectory.

Fig. 9 SLIP model objectives. First, the SLIP hopper follows an information maximizing objective for 10s while learning a Koopman representation of its dynamics. Then, the SLIP hopper uses this model to carry out a forward trajectory.



4.4 Generating Finite Automata via DSS

In Section 4.3 we discussed the importance of modeling hybrid system dynamics in experimental field robotics. Although we were able to generate a hybrid Koopman model for the SLIP hopper, we required *a priori* specification of the system’s indicator function. The indicator function is not generally known ahead of time. Thus, it is of interest to develop a systematic approach to identifying hybrid system dynamics, as well as the underlying indicator function specifying state-space boundaries between hybrid modes.

Hybrid dynamical systems are often represented as finite automata with discrete and continuous components. While the structure of the hybrid automaton is discrete, each of its nodes is a continuous dynamical system. The transitions are specified by the indicator function and guard equations. This graphical model is a compact representation of the system dynamics [19]. Dynamical System Segmentation (DSS), proposed in [4], is an algorithm that generates state-space partitions of dynamical systems from data, thereby generating a hybrid automaton. For hybrid systems, DSS can identify the dynamics of each hybrid mode and generate an indicator function from data. For non-hybrid systems, DSS synthesizes a finite automaton representation of the dynamics from data. The state-space continuous dynamics are partitioned into a set of distinct dynamics governing a given region of the state-space manifold.

Given a dataset $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_M]$ from some dynamical system and a set of basis functions $\Psi(\mathbf{x}) \in \mathbb{R}^N$, DSS synthesizes a set of Koopman operators from subsets of the dataset $\mathbf{U}_i = \text{GenerateKoopman}(\mathbf{X}_{a:b})$, *s.t.* $\mathbf{X}_{a:b} = [\mathbf{x}_a, \dots, \mathbf{x}_b]$. The choice of method for splitting the dataset X into subsets is left to the user. This set of W Koopman operators $\mathcal{U} = \{\mathbf{U}_1, \dots, \mathbf{U}_W\}$ are each a local estimate of the system dynamics for some neighborhood of the state-space manifold. However, given that this set \mathcal{U} may have redundancies we are interested in distilling a minimal set of Koopman operators to represent the system dynamics. By considering each Koopman operator $\mathbf{U}_i \in \mathbb{R}^{N \times N}$ as points in \mathbb{R}^{N^2} space, we can use a clustering algorithm to construct a set of exemplar operators $\bar{\mathcal{U}}$ from the set \mathcal{U} . We suggest nonparametric clustering algorithms, such as HDBSCAN [8], in order to avoid presupposing a number of clusters. We will consider the minimal set of operators $\bar{\mathcal{U}}$ as a node set in a finite automaton. To derive the edge set $\mathcal{E} = \{(\mathbf{U}_i, \mathbf{U}_j), \dots\}$, we generate a support vector machine (SVM), $\Pi(\Psi(\mathbf{x}))$, using the labeled dataset \mathbf{X} and determine state-space




Fig. 10 Segmentation of a non-hybrid dynamical system into a hybrid automaton represented as a graph. Each node in the graph is a distinct dynamical system represented by a Koopman operator governing the system dynamics over a region of the underlying state-space manifold. The partitions of the manifold are specified by an SVM model $\Pi(\Psi(x))$.

transitions between nodes [30]. This graphical model $\mathbb{G} = (\bar{\mathcal{U}}, \mathcal{E})$ is the output of DSS. Figure 10 depicts a notional example of DSS applied to some dynamical system, as well as the relationship between the graph and the partitioned state-space manifold.

To demonstrate the DSS algorithm’s performance in identifying hybrid systems we applied DSS to a dataset of a simulated SLIP hopper as formulated in Section 4.3 following a forward constant velocity trajectory. Figure 11 depicts a successful identification of the bimodal SLIP dynamics purely from data without *any* prior knowledge of the system dynamics. The data-driven SVM indicator function $\Pi(\Psi(\mathbf{x}))$ correctly captures the SLIP system transitions between flight and stance modes.

Additionally, we segment non-hybrid systems to showcase the resulting finite descriptions of state-space continuous systems. We use the well-studied cart-pendulum inversion problem as an example. The cart-pendulum system, shown in Fig. 12, is actuated about its horizontal axis and has $\theta = 0$ defined at the unstable equilibrium point. We used the same nonlinear model-predictive controller as in Section 4.1, SAC, to synthesize a set of 30 trajectories with the goal of inverting and stabilizing the cart-pendulum system. We then applied DSS to this dataset and found that the pendulum inversion task could be segmented into three hybrid modes corresponding to swing-up, slow-down, stabilization dynamics. The synthesized finite automaton




Fig. 11 Dynamical system segmentation of the SLIP model system described in Section 4.3. The DSS algorithm without any *a priori* knowledge of the SLIP system dynamics is able to discern the two hybrid modes of flight and stance, as well as the indicator function that switches between modes. The plot shows the position of the SLIP hopper’s mass color coded according to the hybrid mode it is currently in. Green corresponding to flight, orange to stance.

resulting from DSS is shown in Fig. 13. Each node is accompanied by its corresponding partition of the original dataset shown in the $(\theta, \dot{\theta})$ phase plane.

Dynamical system segmentation is an extension of the hybrid dynamical formulation of the Koopman operator described in Section 2.5, and showcased in Section 4.3. DSS generates finite representations of dynamical systems from data without any prior knowledge of the system. As an experimental tool, as was shown in [4], DSS is capable of discerning complex relationships from motion data and formulating them in a low-dimensional representation. Identification of discontinuous phenomena from data is of great importance in experimental settings where environmental interactions are often difficult to model, and the Koopman operator provides us with the ability to formulate algorithms to address this problem.




Fig. 12 Cart-pendulum system used for DSS, with its unstable equilibrium defined at $\theta = 0$.

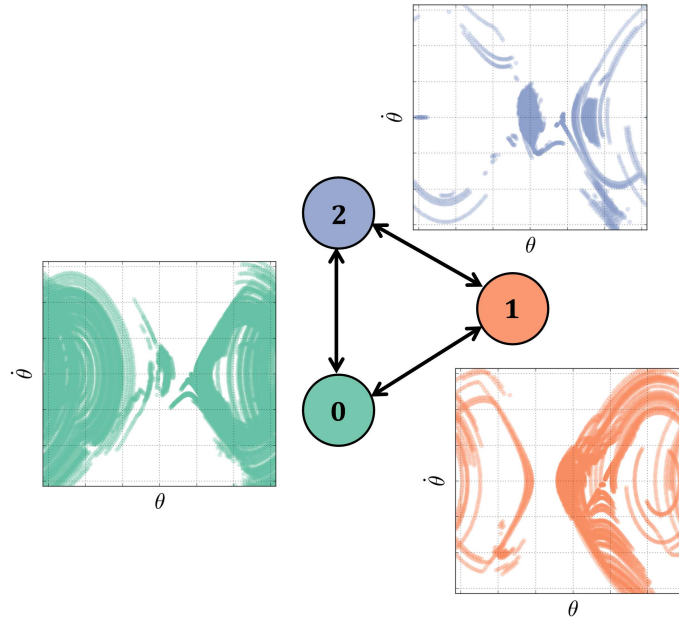


Fig. 13 Application of DSS to a set of 30 control solutions to the cart-pendulum inversion problem. Each node in the generated graph represents a Koopman operator that describes the cart-pendulum nonlinear dynamics over a particular region of the state-space manifold. Mode 0 corresponds to energy pumping and swing-up, mode 1 corresponds to energy removal and slow-down, and mode 2 corresponds to stabilization.

5 Conclusion

Experimental learning is necessary for real-time contexts where the environment and internal dynamics can change rapidly. The learning techniques we apply should reflect that these settings demand incremental learning in time-varying environments, and most machine learning techniques do not. We have proposed the EDMD-synthesized finite-dimensional Koopman operator as a technique well-suited to experimental settings and compared it to similar techniques to characterize their relationships. Moreover, we provide a set of conditions that specify the connection between EDMD and alternative learning methods. We highlight the synthesized Koopman operator's capabilities in a variety of experimental applications. Modeling real processes requires techniques capable of purposeful learning, and we propose the Koopman operator as a suitable model representation for these settings.

Acknowledgements This work was supported by the National Science Foundation under grant CBET-1637764. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. Abraham, I., Murphey, T.: Active learning of dynamics and data-driven control. *IEEE Transactions on Robotics* (2019)
2. Abraham, I., Torre, G., Murphey, T.: Model-based control using Koopman operators. *Robotics: Science and Systems* (2017)
3. Ansari, A., Murphey, T.D.: Sequential action control: Closed-form optimal control for nonlinear and nonsmooth systems. *IEEE Transactions on Robotics* **32** (2017)
4. Berrueta, T.A., Pervan, A., Fitzsimons, K., Murphey, T.D.: Dynamical system segmentation for information measures in motion. *IEEE Robotics and Automation Letters* **4**(1), 169–176 (2019)
5. Bishop, C.: *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer (2006)
6. Brunton, S., Brunton, B., Proctor, J., Kutz, J.: Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLOS ONE* **11**(2), 1–19 (2016)
7. Budišić, M., Mohr, R., Mezić, I.: Applied Koopmanism. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **22**(4), 047,510 (2012)
8. Campello, R., Moulavi, D., Sander, J.: Density-based clustering based on hierarchical density estimates. In: *Advances in Knowledge Discovery and Data Mining*, pp. 160–172. Springer (2013)
9. Chen, B.W., Abdullah, N.N.B., Park, S., Gu, Y.: Efficient multiple incremental computation for kernel ridge regression with Bayesian uncertainty modeling. *Future Generation Computer Systems* **82**, 679 – 688 (2018)
10. Chirikjian, G.S.: *Stochastic Models, Information Theory, and Lie Groups, Volume 2*. Springer (2012)
11. Dias, J., Leitao, J.: A method for computing the information matrix of stationary Gaussian processes. In: *1996 8th European Signal Processing Conference (EUSIPCO 1996)*, pp. 1–4 (1996)
12. Engel, Y., Mannor, S., Meir, R.: The kernel recursive least squares algorithm. *IEEE Transactions on Signal Processing* **52**(8), 2275–2285 (2004)
13. Fan, T., Murphey, T.: Online feedback control for input-saturated robotic systems on Lie groups. *Robotics: Science and Systems* (2016)
14. Gallager, R.: *Stochastic Processes: Theory for Applications*. Cambridge University Press (2013)
15. Govindarajan, N., Arbabi, H., van Blerghien, L., Matchen, T., Tegling, E., Mezić, I.: An operator-theoretic viewpoint to non-smooth dynamical systems: Koopman analysis of a hybrid pendulum. In: *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 6477–6484 (2016)
16. Hamilton, J.: *Time Series Analysis*. Princeton University Press (1994)
17. Heersink, B., Warren, M., Hoffmann, H.: Dynamic mode decomposition for interconnected control systems. *arXiv* (2017). DOI arxiv:1709.02883
18. Hemati, M.S., Williams, M.O., Rowley, C.W.: Dynamic mode decomposition for large and streaming datasets. *Physics of Fluids* **26**(11), 111,701 (2014)
19. Henzinger, T.A.: The theory of hybrid automata. In: *Proceedings of the 11th Annual IEEE Symposium on Logic in Computer Science*, pp. 278–283 (1996)
20. Kaiser, E., Kutz, J.N., Brunton, S.L.: Data-driven discovery of Koopman eigenfunctions for control. In: *arXiv* (2017). DOI arxiv:1707.01146
21. Kalinowska, A., Fitzsimons, K., Dewald, J.P., Murphey, T.D.: Online user assessment for minimal intervention during task-based robotic assistance. In: *Robotics: Science and Systems* (2018)
22. Koopman, B.: Hamiltonian systems and transformation in Hilbert space. *Proceedings of the National Academy of Sciences* **17**(5), 315–318 (1931)
23. Korda, M., Mezić, I.: Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica* **93**, 149 – 160 (2018)

24. Lee, J., Bahri, Y., Novak, R., Schoenholz, S.S., Pennington, J., Sohl-Dickstein, J.: Deep neural networks as Gaussian processes. In: arXiv (2018). DOI arxiv:1711.00165
25. Li, Q., Dietrich, F., Bollt, E.M., Kevrekidis, I.G.: Extended dynamic mode decomposition with dictionary learning: A data-driven adaptive spectral decomposition of the Koopman operator. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **27**(10), 103,111 (2017)
26. Mezić, I.: Analysis of fluid flows via spectral properties of the Koopman operator. *Annual Review of Fluid Mechanics* **45**(1), 357–378 (2013). DOI 10.1146/annurev-fluid-011212-140652
27. Murphy, K.P.: *Machine Learning: A Probabilistic Perspective*. The MIT Press (2012)
28. Neal, R.M.: Priors for infinite networks. In: *Bayesian Learning for Neural Networks*, pp. 29–53. Springer (1996)
29. Nguyen-Tuong, D., Peters, J., Seeger, M.: Local Gaussian process regression for real time online model learning and control. In: *Proceedings of the 21st International Conference on Neural Information Processing Systems*, pp. 1193–1200 (2008)
30. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
31. Peitz, S., Klus, S.: Koopman operator-based model reduction for switched-system control of PDEs. In: arXiv (2017). DOI arXiv:1710.06759
32. Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press (2005)
33. Schmid, P.J.: Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics* **656**, 5–28 (2010)
34. Seal, H.L.: Studies in the history of probability and statistics. XV: The historical development of the Gauss linear model. *Biometrika* **54**(1/2), 1–24 (1967)
35. Shahriari, B., Swersky, K., Wang, Z., Adams, R., de Freitas, N.: Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE* **104**(1), 148–175 (2016)
36. Srivastava, J., Anderson, D.: A comparison of the determinant, maximum root, and trace optimality criteria. *Communications in Statistics* **3**(10), 933–940 (1974)
37. Takeishi, N., Kawahara, Y., Tabei, Y., Yairi, T.: Bayesian dynamic mode decomposition. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, pp. 2814–2821 (2017)
38. Tu, J., Rowley, C., Luchtenburg, D., Brunton, S., Kutz, J.: On dynamic mode decomposition: Theory and applications. *Journal of Computational Dynamics* **1**, 391 (2014). DOI 10.3934/jcd.2014.1.391
39. Umlauf, J., Beckers, T., Kimmel, M., Hirche, S.: Feedback linearization using Gaussian processes. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pp. 5249–5255 (2017)
40. Van Vaerenbergh, S., Fernandez-Bes, J., Elvira, V.: On the relationship between online Gaussian process regression and kernel least mean squares algorithms. In: *2016 IEEE International Workshop on Machine Learning for Signal Processing* (2016)
41. Williams, M., Kevrekidis, I., Rowley, C.: A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science* **25**, 1307–1346 (2015)
42. Williams, M.O., Rowley, C.W., Kevrekidis, I.G.: A kernel-based method for data-driven Koopman spectral analysis. *Journal of Computational Dynamics* **2**, 247–265 (2015)
43. Yeung, E., Kundu, S., Hodas, N.: Learning deep neural network representations for Koopman operators of nonlinear dynamical systems. arXiv (2017). DOI arxiv:1708.06850