
Learning Koopman Invariant Subspaces for Dynamic Mode Decomposition

Naoya Takeishi[§], Yoshinobu Kawahara^{†,‡}, Takehisa Yairi[§]

[§]Department of Aeronautics and Astronautics, The University of Tokyo

[†]The Institute of Scientific and Industrial Research, Osaka University

[‡]RIKEN Center for Advanced Intelligence Project

{takeishi,yairi}@ailab.t.u-tokyo.ac.jp, ykawahara@sanken.osaka-u.ac.jp

Abstract

Spectral decomposition of the Koopman operator is attracting attention as a tool for the analysis of nonlinear dynamical systems. Dynamic mode decomposition is a popular numerical algorithm for Koopman spectral analysis; however, we often need to prepare nonlinear observables manually according to the underlying dynamics, which is not always possible since we may not have any *a priori* knowledge about them. In this paper, we propose a fully data-driven method for Koopman spectral analysis based on the principle of *learning Koopman invariant subspaces* from observed data. To this end, we propose minimization of the residual sum of squares of linear least-squares regression to estimate a set of functions that transforms data into a form in which the linear regression fits well. We introduce an implementation with neural networks and evaluate performance empirically using nonlinear dynamical systems and applications.

1 Introduction

A variety of time-series data are generated from nonlinear dynamical systems, in which a state evolves according to a nonlinear map or differential equation. In summarization, regression, or classification of such time-series data, precise analysis of the underlying dynamical systems provides valuable information to generate appropriate features and to select an appropriate computation method. In applied mathematics and physics, the analysis of nonlinear dynamical systems has received significant interest because a wide range of complex phenomena, such as fluid flows and neural signals, can be described in terms of nonlinear dynamics. A classical but popular view of dynamical systems is based on state space models, wherein the behavior of the trajectories of a vector in state space is discussed (see, e.g., [1]). Time-series modeling based on a state space is also common in machine learning. However, when the dynamics are highly nonlinear, analysis based on state space models becomes challenging compared to the case of linear dynamics.

Recently, there is growing interest in operator-theoretic approaches for the analysis of dynamical systems. Operator-theoretic approaches are based on the Perron–Frobenius operator [2] or its adjoint, i.e., the Koopman operator (composition operator) [3], [4]. The Koopman operator defines the evolution of observation functions (observables) in a function space rather than state vectors in a state space. Based on the Koopman operator, the analysis of nonlinear dynamical systems can be lifted to a linear (but infinite-dimensional) regime. Consequently, we can consider modal decomposition, with which the global characteristics of nonlinear dynamics can be inspected [4], [5]. Such modal decomposition has been intensively used for scientific purposes to understand complex phenomena (e.g., [6]–[9]) and also for engineering tasks, such as signal processing and machine learning. In fact, modal decomposition based on the Koopman operator has been utilized in various engineering tasks, including robotic control [10], image processing [11], and nonlinear system identification [12].

One of the most popular algorithms for modal decomposition based on the Koopman operator is dynamic mode decomposition (DMD) [6], [7], [13]. An important premise of DMD is that the target dataset is generated from a set of observables that spans a function space invariant to the Koopman operator (referred to as Koopman invariant subspace). However, when only the original state vectors are available as the dataset, we must prepare appropriate observables *manually* according to the underlying nonlinear dynamics. Several methods have been proposed to utilize such observables, including the use of basis functions [14] and reproducing kernels [15]. Note that these methods work well only if appropriate basis functions or kernels are prepared; however, it is not always possible to prepare such functions if we have no *a priori* knowledge about the underlying dynamics.

In this paper, we propose a *fully data-driven* method for modal decomposition via the Koopman operator based on the principle of *learning Koopman invariant subspaces* (LKIS) from scratch using observed data. To this end, we estimate a set of parametric functions by minimizing the residual sum of squares (RSS) of linear least-squares regression, so that the estimated set of functions transforms the original data into a form in which the linear regression fits well. In addition to the principle of LKIS, an implementation using neural networks is described. Moreover, we introduce empirical performance of DMD based on the LKIS framework with several nonlinear dynamical systems and applications, which proves the feasibility of LKIS-based DMD as a fully data-driven method for modal decomposition via the Koopman operator.

2 Background

2.1 Koopman spectral analysis

We focus on a (possibly nonlinear) discrete-time autonomous dynamical system

$$\mathbf{x}_{t+1} = \mathbf{f}(\mathbf{x}_t), \quad \mathbf{x} \in \mathcal{M}, \quad t \in \mathbb{T} = \{0\} \cup \mathbb{N}, \quad (1)$$

where \mathcal{M} denotes the state space and $(\mathcal{M}, \Sigma, \mu)$ represents the associated probability space. In dynamical system (1), *Koopman operator* \mathcal{K} [4], [5] is defined as an infinite-dimensional linear operator that acts on *observables* $g : \mathcal{M} \rightarrow \mathbb{R}$ (or \mathbb{C}), i.e.,

$$\mathcal{K}g(\mathbf{x}) = g(\mathbf{f}(\mathbf{x})), \quad (2)$$

with which the analysis of nonlinear dynamics (1) can be lifted to a linear (but infinite-dimensional) regime. Since \mathcal{K} is linear, let us consider a set of eigenfunctions $\{\varphi_1, \varphi_2, \dots\}$ of \mathcal{K} with eigenvalues $\{\lambda_1, \lambda_2, \dots\}$, i.e., $\mathcal{K}\varphi_i = \lambda_i\varphi_i$ for $i \in \mathbb{N}$, where $\varphi : \mathcal{M} \rightarrow \mathbb{C}$ and $\lambda \in \mathbb{C}$. Further, suppose that g can be expressed as a linear combination of those infinite number of eigenfunctions, i.e., $g(\mathbf{x}) = \sum_{i=1}^{\infty} \varphi_i(\mathbf{x})c_i$ with a set of coefficients $\{c_1, c_2, \dots\}$. By repeatedly applying \mathcal{K} to both sides of this equation, we obtain the following modal decomposition:

$$g(\mathbf{x}_t) = \sum_{i=1}^{\infty} \lambda_i^t \varphi_i(\mathbf{x}_0)c_i. \quad (3)$$

Here, the value of g is decomposed into a sum of *Koopman modes* $w_i = \varphi_i(\mathbf{x}_0)c_i$, each of which evolves over time with its frequency and decay rate respectively given by $\angle\lambda_i$ and $|\lambda_i|$, since λ_i is a complex value. The Koopman modes and their eigenvalues can be investigated to understand the dominant characteristics of complex phenomena that follow nonlinear dynamics. The above discussion can also be applied straightforwardly to continuous-time dynamical systems [4], [5].

Modal decomposition based on \mathcal{K} , often referred to as Koopman spectral analysis, has been receiving attention in nonlinear physics and applied mathematics. In addition, it is a useful tool for engineering tasks including machine learning and pattern recognition; the spectra (eigenvalues) of \mathcal{K} can be used as features of dynamical systems, the eigenfunctions are a useful representation of time-series for various tasks, such as regression and visualization, and \mathcal{K} itself can be used for prediction and optimal control. Several methods have been proposed to compute modal decomposition based on \mathcal{K} , such as generalized Laplace analysis [5], [16], the Ulam–Galerkin method [17], and DMD [6], [7], [13]. DMD, which is reviewed in more detail in the next subsection, has received significant attention and been utilized in various data analysis scenarios (e.g., [6]–[9]).

Note that the Koopman operator and modal decomposition based on it can be extended to random dynamical systems actuated by process noise [4], [14], [18]. In addition, Proctor *et al.* [19], [20] discussed Koopman analysis of systems with control signals. In this paper, we primarily target autonomous deterministic dynamics (e.g., Eq. (1)) for the sake of presentation clarity.

2.2 Dynamic mode decomposition and Koopman invariant subspace

Let us review DMD, an algorithm for Koopman spectral analysis (further details are in the supplementary). Consider a set of observables $\{g_1, \dots, g_n\}$ and let $\mathbf{g} = [g_1 \ \dots \ g_n]^\top$ be a vector-valued observable. In addition, define two matrices $\mathbf{Y}_0, \mathbf{Y}_1 \in \mathbb{R}^{n \times m}$ generated by \mathbf{x}_0, \mathbf{f} and \mathbf{g} , i.e.,

$$\mathbf{Y}_0 = [\mathbf{g}(\mathbf{x}_0) \ \dots \ \mathbf{g}(\mathbf{x}_{m-1})] \quad \text{and} \quad \mathbf{Y}_1 = [\mathbf{g}(\mathbf{f}(\mathbf{x}_0)) \ \dots \ \mathbf{g}(\mathbf{f}(\mathbf{x}_{m-1}))], \quad (4)$$

where $m + 1$ is the number of snapshots in the dataset. The core functionality of DMD algorithms is computing the eigendecomposition of matrix $\mathbf{A} = \mathbf{Y}_1 \mathbf{Y}_0^\dagger$ [13], [21], where \mathbf{Y}_0^\dagger is the Moore–Penrose pseudoinverse of \mathbf{Y}_0 . The eigenvectors of \mathbf{A} are referred to as *dynamic modes*, and they coincide with the Koopman modes if the corresponding eigenfunctions of \mathcal{K} are in $\text{span}\{g_1, \dots, g_n\}$ [21]. Alternatively (but nearly equivalently), the condition under which DMD works as a numerical realization of Koopman spectral analysis can be described as follows.

Rather than calculating the infinite-dimensional \mathcal{K} directly, we can consider the restriction of \mathcal{K} to a finite-dimensional subspace. Assume the observables are elements of $L^2(\mathcal{M}, \mu)$. The *Koopman invariant subspace* is defined as $\mathcal{G} \subset L^2(\mathcal{M}, \mu)$ s.t. $\forall g \in \mathcal{G}, \mathcal{K}g \in \mathcal{G}$. If \mathcal{G} is spanned by a finite number of functions, then the restriction of \mathcal{K} to \mathcal{G} , which we denote K , becomes a *finite-dimensional* linear operator. In the sequel, we assume the existence of such \mathcal{G} . If $\{g_1, \dots, g_n\}$ spans \mathcal{G} , then DMD’s matrix $\mathbf{A} = \mathbf{Y}_1 \mathbf{Y}_0^\dagger$ coincides with $\mathbf{K} \in \mathbb{R}^{n \times n}$ asymptotically, wherein \mathbf{K} is the realization of K with regard to the frame (or basis) $\{g_1, \dots, g_n\}$. For modal decomposition (3), the (vector-valued) Koopman modes are given by \mathbf{w} and the values of the eigenfunctions are obtained by $\varphi = \mathbf{z}^\text{H} \mathbf{g}$, where \mathbf{w} and \mathbf{z} are the right- and left-eigenvectors of \mathbf{K} normalized such that $\mathbf{w}_i^\text{H} \mathbf{z}_j = \delta_{i,j}$ [14], [21], and \mathbf{z}^H denotes the conjugate transpose of \mathbf{z} .

Here, an important problem in the practice of DMD arises, i.e., we often have no access to \mathbf{g} that spans a Koopman invariant subspace \mathcal{G} . In this case, for nonlinear dynamics, we must manually prepare adequate observables. Several researchers have addressed this issue; Williams *et al.* [14] leveraged a dictionary of predefined basis functions to transform original data, and Kawahara [15] defined Koopman spectral analysis in a reproducing kernel Hilbert space. Brunton *et al.* [22] proposed the use of observables selected in a data-driven manner [23] from a function dictionary. Note that, for these methods, we must select an appropriate function dictionary or kernel function according to the target dynamics. However, if we have no *a priori* knowledge about them, which is often the case, such existing methods do not have to be applied successfully to nonlinear dynamics.

3 Learning Koopman invariant subspaces

3.1 Minimizing residual sum of squares of linear least-squares regression

In this paper, we propose a method to learn a set of observables $\{g_1, \dots, g_n\}$ that spans a Koopman invariant subspace \mathcal{G} , given a sequence of measurements as the dataset. In the following, we summarize desirable properties for such observables, upon which the proposed method is constructed.

Theorem 1. *Consider a set of square-integrable observables $\{g_1, \dots, g_n\}$, and define a vector-valued observable $\mathbf{g} = [g_1 \ \dots \ g_n]^\top$. In addition, define a linear operator G whose matrix form is given as $\mathbf{G} = (\int_{\mathcal{M}} (\mathbf{g} \circ \mathbf{f}) \mathbf{g}^\text{H} d\mu) (\int_{\mathcal{M}} \mathbf{g} \mathbf{g}^\text{H} d\mu)^\dagger$. Then, $\forall \mathbf{x} \in \mathcal{M}, \mathbf{g}(\mathbf{f}(\mathbf{x})) = G\mathbf{g}(\mathbf{x})$ if and only if $\{g_1, \dots, g_n\}$ spans a Koopman invariant subspace.*

Proof. If $\forall \mathbf{x} \in \mathcal{M}, \mathbf{g}(\mathbf{f}(\mathbf{x})) = G\mathbf{g}(\mathbf{x})$, then for any $\hat{\mathbf{g}} = \sum_{i=1}^n a_i g_i \in \text{span}\{g_1, \dots, g_n\}$,

$$\mathcal{K}\hat{\mathbf{g}} = \sum_{i=1}^n a_i g_i(\mathbf{f}(\mathbf{x})) = \sum_{j=1}^n \left(\sum_{i=1}^n a_i G_{i,j} \right) g_j(\mathbf{x}) \in \text{span}\{g_1, \dots, g_n\},$$

where $G_{i,j}$ denotes the (i, j) -element of \mathbf{G} ; thus, $\text{span}\{g_1, \dots, g_n\}$ is a Koopman invariant subspace. On the other hand, if $\{g_1, \dots, g_n\}$ spans a Koopman invariant subspace, there exists a linear operator K such that $\forall \mathbf{x} \in \mathcal{M}, \mathbf{g}(\mathbf{f}(\mathbf{x})) = K\mathbf{g}(\mathbf{x})$; thus, $\int_{\mathcal{M}} (\mathbf{g} \circ \mathbf{f}) \mathbf{g}^\text{H} d\mu = \int_{\mathcal{M}} K\mathbf{g} \mathbf{g}^\text{H} d\mu$. Therefore, an instance of the matrix form of K is obtained in the form of \mathbf{G} . \square

According to Theorem 1, we should obtain \mathbf{g} that makes $\mathbf{g} \circ \mathbf{f} - G\mathbf{g}$ zero. However, such problems cannot be solved with finite data because \mathbf{g} is a function. Thus, we give the corresponding empirical

risk minimization problem based on the assumption of ergodicity of \mathbf{f} and the convergence property of the empirical matrix as follows.

Assumption 1. For dynamical system (1), the time-average and space-average of a function $g : \mathcal{M} \rightarrow \mathbb{R}$ (or \mathbb{C}) coincide in $m \rightarrow \infty$ for almost all $\mathbf{x}_0 \in \mathcal{M}$, i.e.,

$$\lim_{m \rightarrow \infty} \frac{1}{m} \sum_{j=0}^{m-1} g(\mathbf{x}_j) = \int_{\mathcal{M}} g(\mathbf{x}) d\mu(\mathbf{x}), \quad \text{for almost all } \mathbf{x}_0 \in \mathcal{M}.$$

Theorem 2. Define \mathbf{Y}_0 and \mathbf{Y}_1 by Eq. (4) and suppose that Assumption 1 holds. If all modes are sufficiently excited in the data (i.e., $\text{rank}(\mathbf{Y}_0) = n$), then matrix $\mathbf{A} = \mathbf{Y}_1 \mathbf{Y}_0^\dagger$ almost surely converges to the matrix form of linear operator G in $m \rightarrow \infty$.

Proof. From Assumption 1, $\frac{1}{m} \mathbf{Y}_1 \mathbf{Y}_0^H$ and $\frac{1}{m} \mathbf{Y}_0 \mathbf{Y}_0^H$ respectively converge to $\int_{\mathcal{M}} (g \circ \mathbf{f}) g^H d\mu$ and $\int_{\mathcal{M}} g g^H d\mu$ for almost all $\mathbf{x}_0 \in \mathcal{M}$. In addition, since the rank of $\mathbf{Y}_0 \mathbf{Y}_0^H$ is always n , $(\frac{1}{m} \mathbf{Y}_0 \mathbf{Y}_0^H)^\dagger$ converges to $(\int_{\mathcal{M}} g g^H d\mu)^\dagger$ in $m \rightarrow \infty$ [24]. Consequently, in $m \rightarrow \infty$, $\mathbf{A} = (\frac{1}{m} \mathbf{Y}_1 \mathbf{Y}_0^H) (\frac{1}{m} \mathbf{Y}_0 \mathbf{Y}_0^H)^\dagger$ almost surely converges to \mathbf{G} , which is the matrix form of linear operator G . \square

Since $\mathbf{A} = \mathbf{Y}_1 \mathbf{Y}_0^\dagger$ is the minimum-norm solution of the linear least-squares regression from the columns of \mathbf{Y}_0 to those of \mathbf{Y}_1 , we constitute the learning problem to estimate a set of function that transforms the original data into a form in which the linear least-squares regression fits well. In particular, we minimize RSS, which measures the discrepancy between the data and the estimated regression model (i.e., linear least-squares in this case). We define the *RSS loss* as follows:

$$\mathcal{L}_{\text{RSS}}(\mathbf{g}; (\mathbf{x}_0, \dots, \mathbf{x}_m)) = \left\| \mathbf{Y}_1 - (\mathbf{Y}_1 \mathbf{Y}_0^\dagger) \mathbf{Y}_0 \right\|_{\text{F}}^2, \quad (5)$$

which becomes zero when \mathbf{g} spans a Koopman invariant subspace. If we implement a smooth parametric model on \mathbf{g} , the local minima of \mathcal{L}_{RSS} can be found using gradient descent. We adopt \mathbf{g} that achieves a local minimum of \mathcal{L}_{RSS} as a set of observables that spans (approximately) a Koopman invariant subspace.

3.2 Linear delay embedder for state space reconstruction

In the previous subsection, we have presented an important part of the principle of LKIS, i.e., minimization of the RSS of linear least-squares regression. Note that, to define RSS loss (5), we need access to a sequence of the original states, i.e., $(\mathbf{x}_0, \dots, \mathbf{x}_m) \in \mathcal{M}^{m+1}$, as a dataset. In practice, however, we cannot necessarily observe full states \mathbf{x} due to limited memory and sensor capabilities. In this case, only transformed (and possibly degenerated) measurements are available, which we denote $\mathbf{y} = \psi(\mathbf{x})$ with a measurement function $\psi : \mathcal{M} \rightarrow \mathbb{R}^r$. To define RSS loss (5) given only degenerated measurements, we must reconstruct the original states \mathbf{x} from the actual observations \mathbf{y} .

Here, we utilize delay-coordinate embedding, which has been widely used for state space reconstruction in the analysis of nonlinear dynamics. Consider a univariate time-series $(\dots, y_{t-1}, y_t, y_{t+1}, \dots)$, which is a sequence of degenerated measurements $y_t = \psi(\mathbf{x}_t)$. According to the well-known Taken's theorem [25], [26], a faithful representation of \mathbf{x}_t that preserves the structure of the state space can be obtained by $\tilde{\mathbf{x}}_t = [y_t \quad y_{t-\tau} \quad \dots \quad y_{t-(d-1)\tau}]^\top$ with some lag parameter τ and embedding dimension d if d is greater than $2 \dim(\mathbf{x})$. For a multivariate time-series, embedding with non-uniform lags provides better reconstruction [27]. For example, when we have a two-dimensional time-series $\mathbf{y}_t = [y_{1,t} \quad y_{2,t}]^\top$, an embedding with non-uniform lags is similar to $\tilde{\mathbf{x}}_t = [y_{1,t} \quad y_{1,t-\tau_{11}} \quad \dots \quad y_{1,t-\tau_{1d_1}} \quad y_{2,t} \quad y_{2,t-\tau_{21}} \quad \dots \quad y_{2,t-\tau_{2d_2}}]^\top$ with each value of τ and d . Several methods have been proposed for selection of τ and d [27]–[29]; however, appropriate values may depend on the given application (attractor inspection, prediction, etc.).

In this paper, we propose to surrogate the parameter selection of the delay-coordinate embedding by learning a *linear delay embedder* from data. Formally, we learn embedder ϕ such that

$$\tilde{\mathbf{x}}_t = \phi(\mathbf{y}_t^{(k)}) = \mathbf{W}_\phi [y_t^\top \quad y_{t-1}^\top \quad \dots \quad y_{t-k+1}^\top]^\top, \quad \mathbf{W}_\phi \in \mathbb{R}^{p \times kr}, \quad (6)$$

where $p = \dim(\tilde{\mathbf{x}})$, $r = \dim(\mathbf{y})$, and k is a hyperparameter of maximum lag. We estimate weight \mathbf{W}_ϕ as well as the parameters of \mathbf{g} by minimizing RSS loss (5), which is now defined using $\tilde{\mathbf{x}}$ instead of \mathbf{x} . Learning ϕ from data yields an embedding that is suitable for learning a Koopman invariant subspace. Moreover, we can impose L1 regularization on weight \mathbf{W}_ϕ to make it highly interpretable if necessary according to the given application.

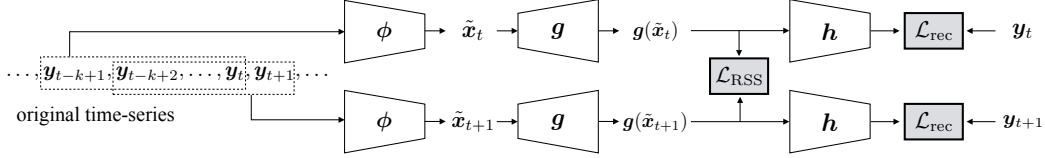


Figure 1: An instance of LKIS framework, in which g and h are implemented by MLPs.

3.3 Reconstruction of original measurements

Simple minimization of \mathcal{L}_{RSS} may yield trivial g , such as constant values. We should impose some constraints to prevent such trivial solutions. In the proposed framework, modal decomposition is first obtained in terms of learned observables g ; thus, the values of g must be back-projected to the space of the original measurements y to obtain a physically meaningful representation of the dynamic modes. Therefore, we modify the loss function by employing an additional term such that the original measurements y can be reconstructed from the values of g by a *reconstructor* h , i.e., $y \approx h(g(\tilde{x}))$. Such term is given as follows:

$$\mathcal{L}_{\text{rec}}(\mathbf{h}, \mathbf{g}; (\tilde{x}_0, \dots, \tilde{x}_m)) = \sum_{j=0}^m \|\mathbf{y}_j - \mathbf{h}(\mathbf{g}(\tilde{x}_j))\|^2, \quad (7)$$

and, if h is a smooth parametric model, this term can also be reduced using gradient descent. Finally, the objective function to be minimized becomes

$$\mathcal{L}(\phi, \mathbf{g}, \mathbf{h}; (\mathbf{y}_0, \dots, \mathbf{y}_m)) = \mathcal{L}_{\text{RSS}}(\mathbf{g}, \phi; (\tilde{x}_{k-1}, \dots, \tilde{x}_m)) + \alpha \mathcal{L}_{\text{rec}}(\mathbf{h}, \mathbf{g}; (\tilde{x}_{k-1}, \dots, \tilde{x}_m)), \quad (8)$$

where α is a parameter that controls the balance between \mathcal{L}_{RSS} and \mathcal{L}_{rec} .

3.4 Implementation using neural networks

In Sections 3.1–3.3, we introduced the main concepts for the LKIS framework, i.e., RSS loss minimization, learning the linear delay embedder, and reconstruction of the original measurements. Here, we demonstrate an implementation of the LKIS framework using neural networks.

Figure 1 shows a schematic diagram of the implementation of the framework. We model g and h using multi-layer perceptrons (MLPs) with a parametric ReLU activation function [30]. Here, the sizes of the hidden layer of MLPs are defined by the arithmetic means of the sizes of the input and output layers of the MLPs. Thus, the remaining tunable hyperparameters are k (maximum delay of ϕ), p (dimensionality of \tilde{x}), and n (dimensionality of g). To obtain g with dimensionality much greater than that of the original measurements, we found that it was useful to set $k > 1$ even when full-state measurements (e.g., $y = x$) were available.

After estimating the parameters of ϕ , g , and h , DMD can be performed normally by using the values of the learned g , defining the data matrices in Eq. (4), and computing the eigendecomposition of $\mathbf{A} = \mathbf{Y}_1 \mathbf{Y}_0^\dagger$; the dynamic modes are obtained by w , and the values of the eigenfunctions are obtained by $\varphi = z^H g$, where w and z are the right- and left-eigenvectors of \mathbf{A} . See Section 2.2 for details.

In the numerical experiments described in Sections 5 and 6, we performed optimization using first-order gradient descent. To stabilize optimization, batch normalization [31] was imposed on the inputs of hidden layers. Note that, since RSS loss function (5) is *not* decomposable with regard to data points, convergence of stochastic gradient descent (SGD) cannot be shown straightforwardly. However, we empirically found that the non-decomposable RSS loss was often reduced successfully, even with mini-batch SGD. Let us show an example; the full-batch RSS loss (denoted $\mathcal{L}_{\text{RSS}}^*$) under the updates of the mini-batch SGD are plotted in the rightmost panel of Figure 4. Here, $\mathcal{L}_{\text{RSS}}^*$ decreases rapidly and remains small. For SGD on non-decomposable losses, Kar *et al.* [32] provided guarantees for some cases; however, examining the behavior of more general non-decomposable losses under mini-batch updates remains an open problem.

4 Related work

The proposed framework is motivated by the operator-theoretic view of nonlinear dynamical systems. In contrast, learning a generative (state-space) model for nonlinear dynamical systems directly has been actively studied in machine learning and optimal control communities, on which we mention a

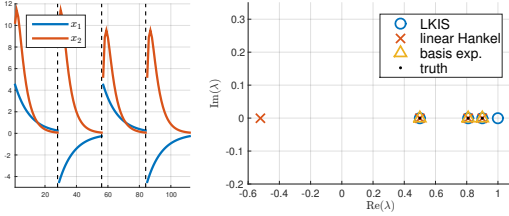


Figure 2: (left) Data generated from system (9) and (right) the estimated Koopman eigenvalues. While linear Hankel DMD produces an inconsistent eigenvalue, LKIS-DMD successfully identifies λ , μ , λ^2 , and $\lambda^0\mu^0 = 1$.

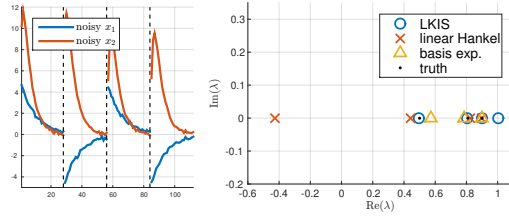


Figure 3: (left) Data generated from system (9) and white Gaussian observation noise and (right) the estimated Koopman eigenvalues. LKIS-DMD successfully identifies the eigenvalues even with the observation noise.

few examples. A classical but popular method for learning nonlinear dynamical systems is using an expectation-maximization algorithm with Bayesian filtering/smoothing (see, e.g., [33]). Recently, using approximate Bayesian inference with the variational autoencoder (VAE) technique [34] to learn generative dynamical models has been actively researched. Chung *et al.* [35] proposed a recurrent neural network with random latent variables, Gao *et al.* [36] utilized VAE-based inference for neural population models, and Johnson *et al.* [37] and Krishnan *et al.* [38] developed inference methods for structured models based on inference with a VAE. In addition, Karl *et al.* [39] proposed a method to obtain a more consistent estimation of nonlinear state space models. Moreover, Watter *et al.* [40] proposed a similar approach in the context of optimal control. Since generative models are intrinsically aware of process and observation noises, incorporating methodologies developed in such studies to the operator-theoretic perspective is an important open challenge to explicitly deal with uncertainty.

We would like to mention some studies closely related to our method. After the first submission of this manuscript (in May 2017), several similar approaches to learning data transform for Koopman analysis have been proposed [41]–[45]. The relationships and relative advantages of these methods should be elaborated in the future.

5 Numerical examples

In this section, we provide numerical examples of DMD based on the LKIS framework (LKIS-DMD) implemented using neural networks. We conducted experiments on three typical nonlinear dynamical systems: a fixed-point attractor, a limit-cycle attractor, and a system with multiple basins of attraction. We show the results of comparisons with other recent DMD algorithms, i.e., Hankel DMD [46], [47], extended DMD [14], and DMD with reproducing kernels [15]. The detailed setups of the experiments discussed in this section and the next section are described in the supplementary.

Fixed-point attractor Consider a two-dimensional nonlinear map on $\mathbf{x}_t = [x_{1,t} \ x_{2,t}]^\top$:

$$x_{1,t+1} = \lambda x_{1,t}, \quad x_{2,t+1} = \mu x_{2,t} + (\lambda^2 - \mu)x_{1,t}^2, \quad (9)$$

which has a stable equilibrium at the origin if $\lambda, \mu < 1$. The Koopman eigenvalues of system (9) include λ and μ , and the corresponding eigenfunctions are $\varphi_\lambda(\mathbf{x}) = x_1$ and $\varphi_\mu(\mathbf{x}) = x_2 - x_1^2$, respectively. $\lambda^i\mu^j$ is also an eigenvalue with corresponding eigenfunction $\varphi_\lambda^i\varphi_\mu^j$. A minimal Koopman invariant subspace of system (9) is $\text{span}\{x_1, x_2, x_1^2\}$, and the eigenvalues of the Koopman operator restricted to such subspace include λ , μ and λ^2 . We generated a dataset using system (9) with $\lambda = 0.9$ and $\mu = 0.5$ and applied LKIS-DMD ($n = 4$), linear Hankel DMD [46], [47] (delay 2), and DMD with basis expansion by $\{x_1, x_2, x_1^2\}$, which corresponds to extended DMD [14] with a right and minimal observable dictionary. The estimated Koopman eigenvalues are shown in Figure 2, wherein LKIS-DMD successfully identifies the eigenvalues of the target invariant subspace. In Figure 3, we show eigenvalues estimated using data contaminated with white Gaussian observation noise ($\sigma = 0.1$). The eigenvalues estimated by LKIS-DMD coincide with the true values even with the observation noise, whereas the results of DMD with basis expansion (i.e., extended DMD) are directly affected by the observation noise.

Limit-cycle attractor We generated data from the limit cycle of the FitzHugh–Nagumo equation

$$\dot{x}_1 = x_1^3/3 + x_1 - x_2 + I, \quad \dot{x}_2 = c(x_1 - bx_2 + a), \quad (10)$$

where $a = 0.7$, $b = 0.8$, $c = 0.08$, and $I = 0.8$. Since trajectories in a limit-cycle are periodic, the (discrete-time) Koopman eigenvalues should lie near the unit circle. Figure 4 shows the eigenvalues

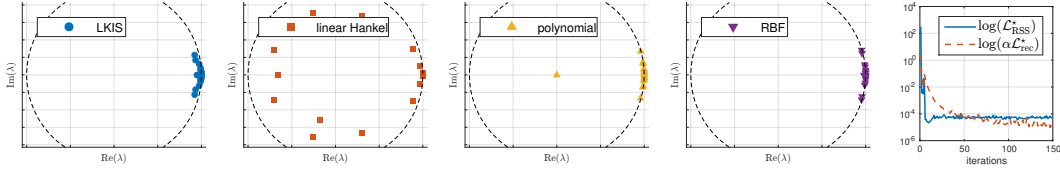


Figure 4: The left four panels show the estimated Koopman eigenvalues on the limit-cycle of the FitzHugh-Nagumo equation by LKIS-DMD, linear Hankel DMD, and kernel DMDs with polynomial and RBF kernels. The hyperparameters of each DMD are set to produce 16 eigenvalues. The rightmost plot shows the full-batch (size 2,000) loss under mini-batch (size 200) SGD updates along iterations. Non-decomposable part $\mathcal{L}_{\text{RSS}}^*$ decreases rapidly and remains small, even by SGD.

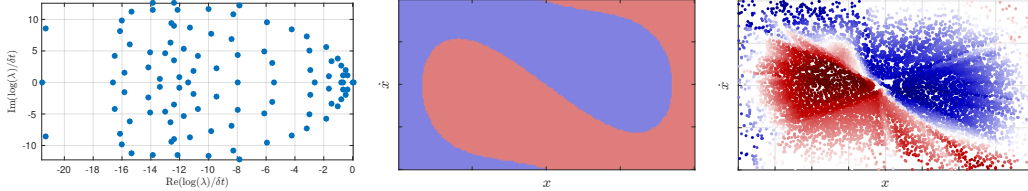


Figure 5: (left) The continuous-time Koopman eigenvalues estimated by LKIS-DMD on the Duffing equation. (center) The true basins of attraction of the Duffing equation, wherein points in the blue region evolve toward $(1, 0)$ and points in the red region evolve toward $(-1, 0)$. Note that the stable manifold of the saddle point is not drawn precisely. (right) The values of the Koopman eigenfunction with a nearly zero eigenvalue computed by LKIS-DMD, whose level sets should correspond to the basins of attraction. There is rough agreement between the true boundary of the basins of attraction and the numerically computed boundary. The right two plots are best viewed in color.

estimated by LKIS-DMD ($n = 16$), linear Hankel DMD [46], [47] (delay 8), and DMDs with reproducing kernels [15] (polynomial kernel of degree 4 and RBF kernel of width 1). The eigenvalues produced by LKIS-DMD agree well with those produced by kernel DMDs, whereas linear Hankel DMD produces eigenvalues that would correspond to rapidly decaying modes.

Multiple basins of attraction Consider the unforced Duffing equation

$$\ddot{x} = -\delta\dot{x} - x(\beta + \alpha x^2), \quad \mathbf{x} = [x \quad \dot{x}]^\top, \quad (11)$$

where $\alpha = 1$, $\beta = -1$, and $\delta = 0.5$. States \mathbf{x} following (11) evolve toward $[1 \ 0]^\top$ or $[-1 \ 0]^\top$ depending on which basin of attraction the initial value belongs to unless the initial state is on the stable manifold of the saddle. Generally, a Koopman eigenfunction whose continuous-time eigenvalue is zero takes a constant value in each basin of attraction [14]; thus, the contour plot of such an eigenfunction shows the boundary of the basins of attraction. We generated 1,000 episodes of time-series starting at different initial values uniformly sampled from $[-2, 2]^2$. The left plot in Figure 5 shows the continuous-time Koopman eigenvalues estimated by LKIS-DMD ($n = 100$), all of which correspond to decaying modes (i.e., negative real parts) and agree with the property of the data. The center plot in Figure 5 shows the true basins of attraction of (11), and the right plot shows the estimated values of the eigenfunction corresponding to the eigenvalue of the smallest magnitude. The surface of the estimated eigenfunction agrees qualitatively with the true boundary of the basins of attractions, which indicates that LKIS-DMD successfully identifies the Koopman eigenfunction.

6 Applications

The numerical experiments in the previous section demonstrated the feasibility of the proposed method as a fully data-driven method for Koopman spectral analysis. Here, we introduce practical applications of LKIS-DMD.

Chaotic time-series prediction Prediction of a chaotic time-series has received significant interest in nonlinear physics. We would like to perform the prediction of a chaotic time-series using DMD, since DMD can be naturally utilized for prediction as follows. Since $\mathbf{g}(\mathbf{x}_t)$ is decomposed as $\sum_{i=1}^n \varphi_i(\mathbf{x}_t) \mathbf{c}_i$ and φ is obtained by $\varphi_i(\mathbf{x}_t) = \mathbf{z}_i^H \mathbf{g}(\mathbf{x}_t)$ where \mathbf{z}_i is a left-eigenvalue of \mathbf{K} , the next step of \mathbf{g} can be described in terms of the current step, i.e., $\mathbf{g}(\mathbf{x}_{t+1}) = \sum_{i=1}^n \lambda_i(\mathbf{z}_i^H \mathbf{g}(\mathbf{x}_t)) \mathbf{c}_i$. In

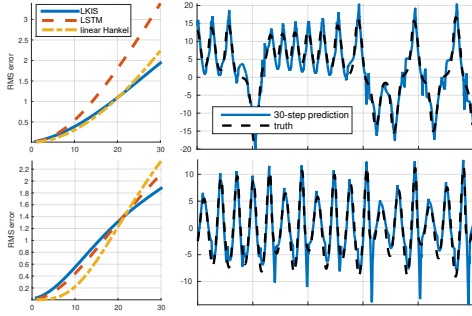


Figure 6: The left plot shows RMS errors from 1- to 30-step predictions, and the right plot shows a part of the 30-step prediction obtained by LKIS-DMD on (*upper*) the Lorenz- x series and (*lower*) the Rossler- x series.

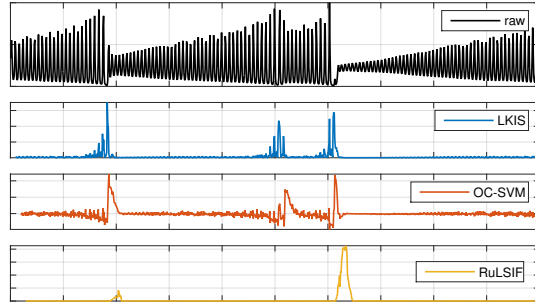


Figure 7: The top plot shows the raw time-series obtained by a far-infrared laser [50]. The other plots show the results of unstable phenomena detection, wherein the peaks should correspond to the occurrences of unstable phenomena.

addition, in the case of LKIS-DMD, the values of \mathbf{g} must be back-projected to \mathbf{y} using the learned \mathbf{h} . We generated two types of univariate time-series by extracting the $\{x\}$ series of the Lorenz attractor [48] and the Rossler attractor [49]. We simulated 25,000 steps for each attractor and used the first 10,000 steps for training, the next 5,000 steps for validation, and the last 10,000 steps for testing prediction accuracy. We examined the prediction accuracy of LKIS-DMD, a simple LSTM network, and linear Hankel DMD [46], [47], all of whose hyperparameters were tuned using the validation set. The prediction accuracy of every method and an example of the predicted series on the test set by LKIS-DMD are shown in Figure 6. As can be seen, the proposed LKIS-DMD achieves the smallest root-mean-square (RMS) errors in the 30-step prediction.

Unstable phenomena detection One of the most popular applications of DMD is the investigation of the global characteristics of dynamics by inspecting the spatial distribution of the dynamic modes. In addition to the spatial distribution, we can investigate the temporal profiles of mode activations by examining the values of corresponding eigenfunctions. For example, assume there is an eigenfunction $\varphi_{\lambda \ll 1}$ that corresponds to a discrete-time eigenvalue λ whose magnitude is considerably smaller than one. Such a small eigenvalue indicates a rapidly decaying (i.e., unstable) mode; thus, we can detect occurrences of unstable phenomena by observing the values of $\varphi_{\lambda \ll 1}$. We applied LKIS-DMD ($n = 10$) to a time-series generated by a far-infrared laser, which was obtained from the Santa Fe Time Series Competition Data [50]. We investigated the values of eigenfunction $\varphi_{\lambda \ll 1}$ corresponding to the eigenvalue of the smallest magnitude. The original time-series and values of $\varphi_{\lambda \ll 1}$ obtained by LKIS-DMD are shown in Figure 7. As can be seen, the activations of $\varphi_{\lambda \ll 1}$ coincide with sudden decays of the pulsation amplitudes. For comparison, we applied the novelty/change-point detection technique using one-class support vector machine (OC-SVM) [51] and direct density-ratio estimation by relative unconstrained least-squares importance fitting (RuLSIF) [52]. We computed AUC, defining the sudden decays of the amplitudes as the points to be detected, which were 0.924, 0.799, and 0.803 for LKIS, OC-SVM, and RuLSIF, respectively.

7 Conclusion

In this paper, we have proposed a framework for learning Koopman invariant subspaces, which is a fully data-driven numerical algorithm for Koopman spectral analysis. In contrast to existing approaches, the proposed method learns (approximately) a Koopman invariant subspace entirely from the available data based on the minimization of RSS loss. We have shown empirical results for several typical nonlinear dynamics and application examples.

We have also introduced an implementation using multi-layer perceptrons; however, one possible drawback of such an implementation is the local optima of the objective function, which makes it difficult to assess the adequacy of the obtained results. Rather than using neural networks, the observables to be learned could be modeled by a sparse combination of basis functions as in [23] but still utilizing optimization based on RSS loss. Another possible future research direction could be incorporating approximate Bayesian inference methods, such as VAE [34]. The proposed framework is based on a discriminative viewpoint, but inference methodologies for generative models could be used to modify the proposed framework to explicitly consider uncertainty in data.

Acknowledgments

This work was supported by JSPS KAKENHI Grant No. JP15J09172, JP26280086, JP16H01548, and JP26289320.

References

- [1] M. W. Hirsch, S. Smale, and R. L. Devaney, *Differential equations, dynamical systems, and an introduction to chaos*, 3rd. Academic Press, 2013.
- [2] A. Lasota and M. C. Mackey, *Chaos, fractals, and noise: Stochastic aspects of dynamics*, 2nd. Springer, 1994.
- [3] B. O. Koopman, “Hamiltonian systems and transformation in Hilbert space,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 17, no. 5, pp. 315–318, 1931.
- [4] I. Mezić, “Spectral properties of dynamical systems, model reduction and decompositions,” *Nonlinear Dynamics*, vol. 41, no. 1-3, pp. 309–325, 2005.
- [5] M. Budišić, R. Mohr, and I. Mezić, “Applied Koopmanism,” *Chaos*, vol. 22, p. 047 510, 2012.
- [6] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, “Spectral analysis of nonlinear flows,” *Journal of Fluid Mechanics*, vol. 641, pp. 115–127, 2009.
- [7] P. J. Schmid, “Dynamic mode decomposition of numerical and experimental data,” *Journal of Fluid Mechanics*, vol. 656, pp. 5–28, 2010.
- [8] J. L. Proctor and P. A. Eckhoff, “Discovering dynamic patterns from infectious disease data using dynamic mode decomposition,” *International Health*, vol. 7, no. 2, pp. 139–145, 2015.
- [9] B. W. Brunton, L. A. Johnson, J. G. Ojemann, and J. N. Kutz, “Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition,” *Journal of Neuroscience Methods*, vol. 258, pp. 1–15, 2016.
- [10] E. Berger, M. Sastuba, D. Vogt, B. Jung, and H. B. Amor, “Estimation of perturbations in robotic behavior using dynamic mode decomposition,” *Advanced Robotics*, vol. 29, no. 5, pp. 331–343, 2015.
- [11] J. N. Kutz, X. Fu, and S. L. Brunton, “Multiresolution dynamic mode decomposition,” *SIAM Journal on Applied Dynamical Systems*, vol. 15, no. 2, pp. 713–735, 2016.
- [12] A. Mauroy and J. Goncalves, “Linear identification of nonlinear systems: A lifting technique based on the Koopman operator,” in *Proceedings of the 2016 IEEE 55th Conference on Decision and Control*, 2016, pp. 6500–6505.
- [13] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor, *Dynamic mode decomposition: Data-driven modeling of complex systems*. SIAM, 2016.
- [14] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley, “A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition,” *Journal of Nonlinear Science*, vol. 25, no. 6, pp. 1307–1346, 2015.
- [15] Y. Kawahara, “Dynamic mode decomposition with reproducing kernels for Koopman spectral analysis,” in *Advances in Neural Information Processing Systems*, vol. 29, 2016, pp. 911–919.
- [16] I. Mezić, “Analysis of fluid flows via spectral properties of the Koopman operator,” *Annual Review of Fluid Mechanics*, vol. 45, pp. 357–378, 2013.
- [17] G. Froyland, G. A. Gottwald, and A. Hammerlindl, “A computational method to extract macroscopic variables and their dynamics in multiscale systems,” *SIAM Journal on Applied Dynamical Systems*, vol. 13, no. 4, pp. 1816–1846, 2014.
- [18] N. Takeishi, Y. Kawahara, and T. Yairi, “Subspace dynamic mode decomposition for stochastic Koopman analysis,” *Physical Review E*, vol. 96, no. 3, 033310, p. 033 310, 3 2017.
- [19] J. L. Proctor, S. L. Brunton, and J. N. Kutz, “Dynamic mode decomposition with control,” *SIAM Journal on Applied Dynamical Systems*, vol. 15, no. 1, pp. 142–161, 2016.
- [20] —, “Generalizing Koopman theory to allow for inputs and control,” arXiv:1602.07647, 2016.
- [21] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz, “On dynamic mode decomposition: Theory and applications,” *Journal of Computational Dynamics*, vol. 1, no. 2, pp. 391–421, 2014.

- [22] S. L. Brunton, B. W. Brunton, J. L. Proctor, and J. N. Kutz, “Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control,” *PLoS ONE*, vol. 11, no. 2, e0150171, 2016.
- [23] S. L. Brunton, J. L. Proctor, and J. N. Kutz, “Discovering governing equations from data by sparse identification of nonlinear dynamical systems,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 113, no. 15, pp. 3932–3937, 2016.
- [24] V. Rakočević, “On continuity of the Moore–Penrose and Drazin inverses,” *Matematički Vesnik*, vol. 49, no. 3-4, pp. 163–172, 1997.
- [25] F. Takens, “Detecting strange attractors in turbulence,” in *Dynamical Systems and Turbulence, Warwick 1980*, ser. Lecture Notes in Mathematics, vol. 898, 1981, pp. 366–381.
- [26] T. Sauer, J. A. Yorke, and M. Casdagli, “Embedology,” *Journal of Statistical Physics*, vol. 65, no. 3-4, pp. 579–616, 1991.
- [27] S. P. Garcia and J. S. Almeida, “Multivariate phase space reconstruction by nearest neighbor embedding with different time delays,” *Physical Review E*, vol. 72, no. 2, 027205, p. 027 205, 2005.
- [28] Y. Hirata, H. Suzuki, and K. Aihara, “Reconstructing state spaces from multivariate data using variable delays,” *Physical Review E*, vol. 74, no. 2, 026202, p. 026 202, 2006.
- [29] I. Vlachos and D. Kugiumtzis, “Nonuniform state-space reconstruction and coupling detection,” *Physical Review E*, vol. 82, no. 1, 016207, p. 016 207, 2010.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the 2015 IEEE International Conference on Computer Vision*, 2015, pp. 1026–1034.
- [31] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 37, 2015, pp. 448–456.
- [32] P. Kar, H. Narasimhan, and P. Jain, “Online and stochastic gradient methods for non-decomposable loss functions,” in *Advances in Neural Information Processing Systems*, vol. 27, 2014, pp. 694–702.
- [33] Z. Ghahramani and S. T. Roweis, “Learning nonlinear dynamical systems using an EM algorithm,” in *Advances in Neural Information Processing Systems*, vol. 11, 1999, pp. 431–437.
- [34] D. P. Kingma and M. Welling, “Stochastic gradient VB and the variational auto-encoder,” in *Proceedings of the 2nd International Conference on Learning Representations*, 2014.
- [35] J. Chung, K. Kastner, L. Dinh, K. Goel, A. C. Courville, and Y. Bengio, “A recurrent latent variable model for sequential data,” in *Advances in Neural Information Processing Systems*, vol. 28, 2015, pp. 2980–2988.
- [36] Y. Gao, E. W. Archer, L. Paninski, and J. P. Cunningham, “Linear dynamical neural population models through nonlinear embeddings,” in *Advances in Neural Information Processing Systems*, vol. 29, 2016, pp. 163–171.
- [37] M. Johnson, D. K. Duvenaud, A. Wiltchko, R. P. Adams, and S. R. Datta, “Composing graphical models with neural networks for structured representations and fast inference,” in *Advances in Neural Information Processing Systems*, vol. 29, 2016, pp. 2946–2954.
- [38] R. G. Krishnan, U. Shalit, and D. Sontag, “Structured inference networks for nonlinear state space models,” in *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 2017, pp. 2101–2109.
- [39] M. Karl, M. Soelch, J. Bayer, and P. van der Smagt, “Deep variational Bayes filters: Unsupervised learning of state space models from raw data,” in *Proceedings of the 5th International Conference on Learning Representations*, 2017.
- [40] M. Watter, J. Springenberg, J. Boedecker, and M. Riedmiller, “Embed to control: A locally linear latent dynamics model for control from raw images,” in *Advances in Neural Information Processing Systems*, vol. 28, 2015, pp. 2746–2754.
- [41] Q. Li, F. Dietrich, E. M. Bollt, and I. G. Kevrekidis, “Extended dynamic mode decomposition with dictionary learning: A data-driven adaptive spectral decomposition of the Koopman operator,” *Chaos*, vol. 27, p. 103 111, 2017.
- [42] E. Yeung, S. Kundu, and N. Hodas, “Learning deep neural network representations for Koopman operators of nonlinear dynamical systems,” arXiv:1708.06850, 2017.

- [43] A. Mardt, L. Pasquali, H. Wu, and F. Noé, “VAMPnets: Deep learning of molecular kinetics,” arXiv:1710.06012, 2017.
- [44] S. E. Otto and C. W. Rowley, “Linearly-recurrent autoencoder networks for learning dynamics,” arXiv:1712.01378, 2017.
- [45] B. Lusch, J. N. Kutz, and S. L. Brunton, “Deep learning for universal linear embeddings of nonlinear dynamics,” arXiv:1712.09707, 2017.
- [46] H. Arbabi and I. Mezić, “Ergodic theory, dynamic mode decomposition and computation of spectral properties of the Koopman operator,” *SIAM Journal on Applied Dynamical Systems*, vol. 16, no. 4, 2096–2126, 2017.
- [47] Y. Susuki and I. Mezić, “A Prony approximation of Koopman mode decomposition,” in *Proceedings of the 2015 IEEE 54th Conference on Decision and Control*, 2015, pp. 7022–7027.
- [48] E. N. Lorenz, “Deterministic nonperiodic flow,” *Journal of the Atmospheric Sciences*, vol. 20, no. 2, pp. 130–141, 1963.
- [49] O. E. RöSSLer, “An equation for continuous chaos,” *Physical Letters*, vol. 57A, no. 5, pp. 397–398, 1976.
- [50] A. S. Weigend and N. A. Gershenfeld, Eds., *Time series prediction: Forecasting the future and understanding the past*, ser. Santa Fe Institute Series. Westview Press, 1993.
- [51] S. Canu and A. Smola, “Kernel methods and the exponential family,” *Neurocomputing*, vol. 69, no. 7-9, pp. 714–720, 2006.
- [52] S. Liu, M. Yamada, N. Collier, and M. Sugiyama, “Change-point detection in time-series data by relative density-ratio estimation,” *Neural Networks*, vol. 43, pp. 72–83, 2013.