

MIT Open Access Articles

*Diffusion Maps Kalman Filter for a
Class of Systems with Gradient Flows*

The MIT Faculty has made this article openly available. ***Please share***
how this access benefits you. Your story matters.

As Published: 10.1109/TSP.2020.2987750

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/135213>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Diffusion Maps Kalman Filter for a Class of Systems with Gradient Flows

Tal Shnitzer, *Student Member, IEEE*, Ronen Talmon, *Member, IEEE*, Jean-Jacques Slotine

Abstract—In this paper, we propose a non-parametric method for state estimation of high-dimensional nonlinear stochastic dynamical systems, which evolve according to gradient flows with isotropic diffusion. We combine diffusion maps, a manifold learning technique, with a linear Kalman filter and with concepts from Koopman operator theory. More concretely, using diffusion maps, we construct data-driven virtual state coordinates, which linearize the system model. Based on these coordinates, we devise a data-driven framework for state estimation using the Kalman filter. We demonstrate the strengths of our method with respect to both parametric and non-parametric algorithms in three tracking problems. In particular, applying the approach to actual recordings of hippocampal neural activity in rodents directly yields a representation of the position of the animals. We show that the proposed method outperforms competing non-parametric algorithms in the examined stochastic problem formulations. Additionally, we obtain results comparable to classical parametric algorithms, which, in contrast to our method, are equipped with model knowledge.

Index Terms—Intrinsic modeling, manifold learning, Kalman filter, diffusion maps, non-parametric filtering.

I. INTRODUCTION

IN many real applications, the system model is not accessible and some estimation of it is required. State estimation and characterization of stochastic, possibly nonlinear, dynamical systems are therefore widely studied problems. Traditionally, such problems are addressed using classical algorithms, which rely on predefined parametric models. On the one hand, parametric models need to be sufficiently simple to allow accurate parameter estimation from measurements. On the other hand, too simple models often fail to accommodate the complexity of real systems. This facilitates the development of non-parametric methods [1]–[5]. Particularly in this paper, we take a non-parametric approach and propose a new method to derive the system model in a data-driven manner.

To demonstrate the primary idea, consider a classical nonlinear Simultaneous Localization and Mapping (SLAM) problem, where the goal is to track the 2D position $\mathbf{x} = (x_1, x_2)^T$ of a moving object. Typically, in the bearing-only version of this problem, the accessible system measurements are given by the azimuth of an object, i.e. by $\phi = \arctan(x_1/x_2)$. This nonlinear model, relating the state (position) of the system \mathbf{x} to the measurements ϕ , poses a challenge for processing and analysis, since traditional linear methods cannot be applied.

In [6], this nonlinearity is resolved by constructing virtual measurements, y , and measurement mapping, \mathbf{H} , based on the knowledge of the system properties and model, which allow for the formulation of an equivalent linear problem and the application of a linear time-varying Kalman filter [7]. Briefly, since the system state is given by $\mathbf{x} = (\sin \phi, \cos \phi)$, linearization is achieved by defining the virtual measurements $y = \mathbf{H}\mathbf{x} + v$, where v is noise and:

$$\mathbf{H} = \begin{bmatrix} \cos \phi & -\sin \phi \end{bmatrix} \quad (1)$$

such that $\mathbf{H}\mathbf{x} = 0$. The resulting measurement equation is linear and can be constructed from the given measurements ϕ . In this work, analogously to the virtual measurements that linearize the system, we propose a computational method to construct the data-driven non-parametric counterparts of a virtual state, linearizing the system dynamics and measurement model. However, in contrast to [6], our construction of this virtual state is data-driven and does not require any explicit knowledge of the system properties or measurement modality, e.g. the knowledge that the measurements represent the azimuth in a 2D tracking problem.

One way to devise such a computational method, which has recently drawn significant research attention, is to address the problem of data-driven system analysis and state estimation from an operator-theoretic point of view. In this approach, the dynamical system is described by two dual operators, the Perron-Frobenius operator, which represents probability density evolution, and the Koopman Operator, which is defined on some linear functional space of infinite dimension and describes the time-evolution of observables [8], [9].

In the context of empirical dynamical systems analysis, the main challenge is to approximate these operators from the system measurements. Several methods for estimating the Koopman Operator have been proposed in recent years [9]–[14]. For example, the Extended Dynamic Mode Decomposition (EDMD) [9], [10] approximates the Koopman eigenfunctions and modes based on two sets of points, related through system dynamics, and a set of dictionary elements. However, the optimal choice of the dictionary in EDMD depends on the data [10]. This framework for estimating the Koopman eigenfunctions and modes was later employed in [3] as part of a non-parametric Kalman filtering framework, where a linear Kalman filter is constructed based on the approximation of the Koopman Operator, along with its eigenvectors, eigenvalues and Koopman modes using EDMD. The Kalman filter propagates the system in the space spanned by the Koopman eigenvectors and then the resulting estimates are

Tal Shnitzer and Ronen Talmon are with the Department of Electrical Engineering, Technion – Israel Institute of Technology, Technion City, Haifa, Israel 32000 (e-mail: shnitzer@campus.technion.ac.il; ronen@ee.technion.ac.il).

Jean-Jacques E. Slotine is with the Nonlinear Systems Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, U.S.A. (e-mail: jjs@mit.edu).

projected back to the state space. Other work on the analysis of nonlinear stochastic dynamical systems based on Koopman operator theory includes [11], [12], [15]. The work in [15] offers a formal definition and rigorous mathematical analysis for the generalization of the Koopman operator to stochastic dynamical systems. In addition, a new framework for approximating the eigenfunctions and eigenvalues of this stochastic Koopman operator is presented. A different approach is proposed in [11] and [12], where the authors characterize the long term behavior of a system (asymptotic dynamics) based on time-averages of functions. In [11], invariant-measures are defined and it is shown that these measures coincide with the eigenfunctions of the Koopman Operator and can be simply calculated by the Fourier transform. In [12], this framework is extended to address dynamical systems which are not measure-preserving using Laplace averages. In both [11] and [12], several trajectories of the system are needed for the analysis.

Different related non-parametric frameworks for state estimation in stochastic dynamical systems based on geometry and manifold learning were presented in [5], [16]–[18]. In [5] and [16] a probabilistic approach is proposed, in which the problem is projected onto coordinates constructed by diffusion maps [19], a manifold learning technique. In these diffusion maps coordinates, the probability density function of the system state can be propagated in time without prior knowledge of the system dynamics, yet, the underlying system state is assumed to be accessible. The work in [17] proposes to construct an ensemble Kalman filter based on delay embedding coordinates (Takens embedding [20]), used for dynamics estimation. The propagation in time is estimated based on nearest neighbors of the current time-lag frame. In [18], a framework for estimating the eigenfunctions of the Koopman Operator generator based on diffusion maps is presented. This work discusses the relationship between diffusion maps and the Koopman operator and proposes to estimate the Koopman generator eigenfunctions using the eigenfunctions of the Laplace-Beltrami operator approximated by diffusion maps.

The relationship between diffusion maps and the Koopman Operator is further discussed in [8]. There, it is shown that EDMD can also provide data-driven dimensionality reduction. Moreover, for systems in which the system state is described by a Markov process, the eigenfunctions of the backward Fokker-Planck operator can be approximated using EDMD, similarly to diffusion maps. The main benefit of such a manifold learning approach using EDMD is that both the dynamics and the geometry of the underlying state are taken into account.

In this paper, we present a framework for state estimation of stochastic dynamical systems, with state dynamics of gradient flows and isotropic diffusion, which reveals the system model with minimal prior assumptions, using diffusion maps [19] and the Kalman filter. We assume that we are given a set of noisy measurements from some unknown nonlinear function of a stochastic underlying state and show that a *linear* model describing the system can be revealed, even for highly nonlinear systems. This is obtained in a completely data-driven manner, based on virtual state coordinates constructed by diffusion maps and their inherent dynamics [21]. A Kalman filter is

then formulated based on the recovered system model and utilized for non-parametric state estimation. By constructing the Kalman filter based on the recovered model, we incorporate system dynamics into diffusion maps, combining geometry and dynamics, as in [8], from a new manifold learning standpoint. We further show that our method uncovers an operator describing the system dynamics, which is analogous to the Stochastic Koopman Operator [8], an extension of the Koopman Operator for stochastic systems. Moreover, our devised method is well suited for high-dimensional systems due to the nonlinear dimensionality reduction obtained by diffusion maps.

The framework in this paper is tightly related to [22], in which a linear observer was constructed by exploiting similar properties of the diffusion maps coordinates and their dynamics. We show that the Kalman filter framework introduced in the present work outperforms the observer framework both by reducing the amount of hyper-parameters and by significantly improving the results in certain scenarios.

With respect to previous work, our method encompasses several key differences. First, it does not rely on accessibility to the state of the system (as assumed in [5], [16]). Second, it does not require predefined dictionary elements (as required in [10]). Third, our method addresses stochastic dynamical systems in contrast to the deterministic systems considered in [3], [17]. Finally, it does not require a training set of samples with known states, in contrast to common non-parametric filtering algorithms [1], [2].

The remainder of the paper is organized as follows. In Section II, the general problem setting is presented. In Section III, we first overview the key-points of our method and then describe the derivation of the Diffusion Maps Kalman filter framework in detail. In Section IV, we demonstrate the strengths of our method on two object tracking problems. We compare our method with both parametric and non-parametric methods. Section V concludes the paper with a short summary.

II. PROBLEM FORMULATION

Consider an ergodic stochastic dynamical system with some nonlinear generator $S : \mathcal{M} \rightarrow \mathcal{M}$, where \mathcal{M} is a compact Riemannian manifold of dimension d . The system is defined by:

$$\dot{\theta}_t = S(\theta_t, \mathbf{u}_t) \quad (2)$$

$$\mathbf{z}_t = g(\theta_t) + \mathbf{v}_t \quad (3)$$

where $\theta_t \in \mathcal{M}$ is the system state, $\dot{\theta}_t$ is its time derivative, $\mathbf{u}_t \in \mathbb{R}^d$ is the process noise, $\mathbf{z}_t \in \mathbb{R}^m$ are the system measurements through some unknown nonlinear function g and $\mathbf{v}_t \in \mathbb{R}^m$ is the measurement noise. The evolution in time of such a dynamical system can be described by the Stochastic Koopman Operator [8], which is defined by

$$(U_{st}f)(\theta_t) = \mathbb{E}[f \circ T(\theta_t, \mathbf{u}_t)] \quad (4)$$

where T is the flow induced by S , \circ is the composition operator, $f : \mathcal{M} \rightarrow \mathbb{R}$ are some observables from an infinite dimensional vector space, closed under composition with T , and \mathbb{E} denotes expectation.

One of the notable properties of the Koopman Operator, which has increased its usage in a line of recent work, is that it is linear in the space of observables, even for highly nonlinear dynamical systems: $U_{st}(\alpha_1 f_1 + \alpha_2 f_2) = \mathbb{E}[(\alpha_1 f_1 + \alpha_2 f_2) \circ T] = \alpha_1 \mathbb{E}[f_1 \circ T] + \alpha_2 \mathbb{E}[f_2 \circ T] = \alpha_1 U_{st} f_1 + \alpha_2 U_{st} f_2$. However, the tradeoff is that even for finite dimensional dynamical systems, the Koopman Operator acts on an infinite dimensional space of observables.

In this work, we focus on the case in which the dynamics equation (2) takes the form of a Langevin equation:

$$\dot{\theta}_t = -\nabla U(\theta_t) + \sqrt{\frac{2}{\beta}} \dot{\mathbf{u}}_t \quad (5)$$

where $U(\theta_t)$ is a smooth and bounded potential, $\sqrt{\frac{2}{\beta}}$ is a constant diffusion coefficient, \mathbf{u}_t is Brownian motion and $\dot{\mathbf{u}}_t$ is its time derivative.

The goal in this work is to build a new coordinate system representing the hidden system state θ_t and devise a filtering framework in these new coordinates, given only the measurements z_t , without any prior knowledge on the system equations (2) and (3).

III. DIFFUSION-BASED KALMAN FILTER

In this section we lay the foundation for our proposed framework and present the theoretical results which allow for the construction of a data-driven linear Kalman filter.

A. Overview

We present a method based on diffusion maps that discovers a new coordinate system describing a model of the state of the system with *linear* drift in a completely data-driven manner. Based on this linear drift, our method constructs a linear operator, analogous to the Stochastic Koopman Operator, from measurements, without prior model knowledge. By exploiting the linearity of this operator, we will formulate a linear Kalman filter framework, which allows for estimation of trajectories of the underlying system state based on noisy nonlinear measurements. Now we will briefly overview the key points of our method, which is described in detail in the following subsections.

Given noisy measurements $z_t \in \mathbb{R}^m$, we use diffusion maps to represent the system state θ_t by a new set of k coordinates, denoted by Φ_t , where $k < m$. We will show that in this new coordinate system, the dynamical system can be described by the following *linear* equations:

$$\dot{\Phi}_t = F\Phi_t + Q_t^{1/2} \dot{\omega}_t \quad (6)$$

$$z_t = H\Phi_t + R_t^{1/2} v_t \quad (7)$$

where F is a linear operator describing the linear drift of the dynamics of the new coordinates Φ_t , H is a linear lift operator from the new coordinates Φ_t to the measurements z_t , $\dot{\omega}_t$ is a standard normally distributed noise process, v_t is measurement noise, and Q_t and R_t are matrices determining the covariance of the driving and measurement noise processes, respectively.

We will further show that the linear operators F and H , in (6) and (7) respectively, can be constructed using

diffusion maps in a data-driven manner, providing a coarse approximation of the system. Yet, this construction can be further improved since diffusion maps ignore the inherent time-dependencies between consecutive samples. Therefore, we formulate a linear Kalman filter using the constructed coordinates Φ_t and the recovered system operators F and H , thereby improving the state estimate by incorporating the dynamics into the diffusion maps coordinates. This leads to a data-driven linear filtering framework, which can be applied to nonlinear systems with an unknown model, revealing a new representation of the system $\hat{\Phi}_t$, which is tightly related to the system state θ_t as will be described in Subsection III-B. In addition, our framework allows for the estimation of specific realizations of system trajectories based on measurements, in contrast to most existing work on the Stochastic Koopman Operator, which only represent the average time-evolution.

The remainder of this section is described as follows. In Subsection III-B and Subsection III-C, we reiterate the derivations presented in [22] for state and model recovery using diffusion maps [19]. In Subsection III-D, we present our proposed Kalman filter framework. In Subsection III-E, we discuss the observability and detectability of the proposed framework and in Subsection III-F, we elaborate on the relation between our framework and the Stochastic Koopman Operator.

B. Recovering the State

Given noiseless measurements, $z_t = g(\theta_t)$, of the hidden system state, θ_t , the following kernel is defined

$$k_\epsilon(s, t) = \exp\left(-\frac{d^2(z_s, z_t)}{\epsilon^2}\right) \quad (8)$$

where $\epsilon > 0$ is a kernel scale, traditionally set to the median of the distances between the measurements, and $d(z_s, z_t)$ is a distance function between z_s and z_t . In our case, we calculated this distance using a modified Mahalanobis distance, first presented in [23]:

$$d(z_s, z_t) = \sqrt{\frac{1}{2} (z_s - z_t) (C_s^\dagger + C_t^\dagger) (z_s - z_t)^T} \quad (9)$$

where C_s^\dagger and C_t^\dagger are the pseudo-inverse of the measurement covariance matrices at times s and t , respectively. In [23], it was shown that this modified Mahalanobis distance between the measurements approximates the Euclidean distance between the hidden states.

The constructed kernel is then normalized according to

$$p_\epsilon(s, t) = \frac{k_\epsilon(s, t)}{d_\epsilon(s)} \quad (10)$$

where $d_\epsilon(s) = \int k_\epsilon(s, t) p_{eq}(\theta_t) d\theta_t$ and $p_{eq}(\theta_t) = e^{-U(\theta_t)}$ is the equilibrium density of the hidden state θ_t .

We can then define the operator P_ϵ by

$$(P_\epsilon f)(\theta_s) = \int p_\epsilon(s, t) f(\theta_t) p_{eq}(\theta_t) d\theta_t$$

where f is a real function defined on the hidden state θ_t . Based on P_ϵ we define

$$L_\epsilon = \frac{P_\epsilon - I}{\epsilon} \quad (11)$$

where I is the identity operator. In [23] it was shown that for hidden states θ_t with dynamics as in (5), the operator L_ϵ converges to the backward Fokker-Planck operator \mathcal{L} defined on the manifold \mathcal{M} , as $\epsilon \rightarrow 0$:

$$\mathcal{L}q = \frac{1}{\beta} \Delta q - \nabla q \cdot \nabla U \quad (12)$$

where q denotes functions in a subspace of observables defined on the system state, which describe averages of functions: $q(\theta_t) = \mathbb{E}[h(\theta_t) \mid \theta_0 = \mathbf{a}_0]$, where \mathbf{a}_0 is some initial state, and h is smooth.

The operator L_ϵ has a discrete set of decreasing eigenvalues, $\{-\lambda^{(\ell)}\}_{\ell \in \mathbb{N}}$, $0 = -\lambda^{(0)} > -\lambda^{(1)} \geq -\lambda^{(2)} \geq \dots$, and eigenfunctions, $\{\phi^{(\ell)}\}_{\ell \in \mathbb{N}}$ [19]. In many stochastic systems, these eigenvalues exhibit a spectral gap, implying on only a few dominant eigenvalues (which are close to 0) [21]. In such systems, we can represent the hidden state using only the principal eigenfunctions, corresponding to the largest (non-trivial $\ell \neq 0$) eigenvalues. The diffusion maps coordinates are then obtained by calculating the eigenvalue decomposition of L_ϵ and using the k eigenfunctions corresponding to the k largest eigenvalues:

$$z_t \mapsto \Phi(\theta_t) = [\phi^{(1)}(\theta_t), \phi^{(2)}(\theta_t), \dots, \phi^{(k)}(\theta_t)] \quad (13)$$

It was shown in [23] that when using a kernel based on the modified Mahalanobis distance (9), the eigenfunctions of L_ϵ converge to the eigenfunctions of the backward Fokker-Planck operator \mathcal{L} , defined on the *hidden state* θ_t , as $\epsilon \rightarrow 0$.

The diffusion maps coordinates in (13) are tightly related to the system state. Consider the adjoint of \mathcal{L} , which is known as the forward Fokker-Planck operator. The forward Fokker-Planck operator exhibits two important properties. First, it describes the evolution of the transition probability density, $p(\theta, t \mid \theta_0)$. Second, it is Hermitian, and therefore, its eigenfunctions form a basis for the space of real functions of the system state (with the equilibrium density as a measure $p_{eq}(\theta) = \exp^{-U(\theta)}$). In addition, since the backward and forward operators are adjoint, their eigenfunctions can be normalized to be bi-orthonormal, and the eigenfunctions of the backward operator can be used as a bi-orthonormal basis. Due to this connection between the eigenfunctions of the backward and forward operators and since the eigenfunctions of the forward operator are tightly related to the system state, the eigenfunctions of the backward Fokker-Planck operator can be used to represent the system state as well. Therefore, from this point on, the diffusion maps coordinates, denoted by $\phi^{(\ell)}$, which approximate these eigenfunctions, can be considered as a new set of coordinates representing the hidden system state.

Note that the above results are obtained only when we have access to clean measurements $z_t = g(\theta_t)$. We will address this issue in Subsection III-D.

C. Recovering the Model

We will now show that the representation of the system state using the diffusion maps coordinates (13) can be used to construct a data-driven representation of the system model.

The Dynamics of the Diffusion Maps Coordinates: Consider the state dynamics in (5) measured through $z_t = g(\theta_t)$. For such systems, based on Itô's Lemma, the eigenfunctions of the operator L_ϵ , obtained using the diffusion maps algorithm, evolve according to a stochastic differential equation (SDE) of known form [21]:

$$\dot{\phi}^{(\ell)}(\theta) = -\lambda^{(\ell)} \phi^{(\ell)}(\theta) + \sqrt{2} \|\nabla_{\theta} \phi^{(\ell)}(\theta)\|_2 \dot{\omega}^{(\ell)} \quad (14)$$

where $\phi^{(\ell)}(\theta)$ is the ℓ th eigenfunction of the operator L_ϵ , $-\lambda^{(\ell)}$ is the corresponding eigenvalue, $\omega^{(\ell)}$ is the ℓ th coordinate of a multidimensional Brownian motion process, where each coordinate is some linear combination of the Brownian motion coordinates of the underlying process, $u^{(\ell)}$, and $\|\cdot\|_2$ is the L^2 norm. Note that we omitted the time notation of the state (θ_t) to indicate that the eigenfunctions are dependent only on the state.

This SDE depicts that the obtained diffusion maps coordinates evolve according to a linear drift, $-\lambda^{(\ell)} \phi^{(\ell)}(\theta)$, and an additional diffusion component. We note that the linear drift component is fully known since we have, using diffusion maps, both the eigenfunctions and the corresponding eigenvalues of the operator L_ϵ , which approximates the backward Fokker-Planck operator. This linear drift will be used as the new *linear* state dynamics.

Constructing the Lift Function: As stated in Subsection III-B, the eigenfunctions of the backward Fokker-Planck operator, approximated by diffusion maps, form a basis for the space of real functions defined on the system state. Therefore, every real function of the system state can be written as an expansion in these eigenfunctions. Specifically, we can represent the measurement function in the following manner:

$$z_t^{(i)} = g^{(i)}(\theta_t) = \sum_{\ell=0}^{\infty} \alpha_{i,\ell} \phi^{(\ell)}(\theta_t) \quad (15)$$

where $\alpha_{i,\ell} = \langle z_t^{(i)}, \phi^{(\ell)} \rangle_{p_{eq}} = \int_{\mathcal{M}} z_t^{(i)} \phi^{(\ell)}(\theta_t) p_{eq}(\theta_t) d\theta_t$. When the eigenvalues of the backward Fokker-Planck operator exhibit a spectral gap, most of the energy is captured by the k eigenfunctions, corresponding to the k largest eigenvalues. In such cases we can approximate the mapping in (15) using only these k eigenfunctions:

$$z_t^{(i)} \approx \sum_{\ell=0}^k \alpha_{i,\ell} \phi^{(\ell)}(\theta_t) \quad (16)$$

We can now write expression (16) in matrix form:

$$z_t \approx \alpha \Phi(\theta_t) \quad (17)$$

where $\Phi(\theta_t)$ is defined in (13) and α is an $m \times k$ matrix in which $(\alpha)_{i,\ell} = \alpha_{i,\ell}$.

These results imply that through the representation of the measurement function using the diffusion maps eigenfunctions, we obtain a linear mapping between the eigenfunctions, $\Phi(\theta_t)$, and the system measurements, z_t . Thus, we set the linear lift function simply to be α .

To conclude this subsection, we note that all of the theoretical derivations above are true for state equations of the form (5). However, it is not essential that specifically the state will

exhibit such dynamics, but rather that the dynamics of some underlying system parameter are governed by the Langevin equation. In such systems, the diffusion maps coordinates, constructed using the modified Mahalanobis distance (9), exhibit useful properties and can be used as a foundation for the proposed time-series filtering framework, as described in the remainder of this paper. Since many natural phenomena are governed by dynamics that can be modeled using Langevin equation (5), our framework is applicable to a wide range of problems.

D. Diffusion Maps Kalman Filter

Based on the Euler-Maruyama method, the dynamics of the diffusion maps coordinates can be discretized to

$$\begin{aligned} \phi^{(\ell)}(\boldsymbol{\theta}_{n+1}) &= (1 - \lambda^{(\ell)}\Delta t) \phi^{(\ell)}(\boldsymbol{\theta}_n) \\ &+ \sqrt{2} \|\nabla_{\boldsymbol{\theta}} \phi^{(\ell)}(\boldsymbol{\theta}_n)\|_2 \Delta\omega_n^{(\ell)} \end{aligned} \quad (18)$$

where Δt is the time step, $\Delta\omega_n^{(\ell)}$ is a normally distributed noise process and $\nabla_{\boldsymbol{\theta}}$ denotes the gradient according to the system state $\boldsymbol{\theta}_n$. To approximate the linear drift in (18) from discrete measurements, discrete approximations of $\phi^{(\ell)}(\boldsymbol{\theta}_n)$ and $\lambda^{(\ell)}$ are required. These approximations can be obtained using the discrete counterpart of diffusion maps as follows.

Given discrete-time measurements $\{\mathbf{z}_n\}_{n=1}^N$, the $N \times N$ kernel matrix is calculated similarly to (8) by:

$$K(i, j) = \exp\left(-\frac{d^2(\mathbf{z}_i, \mathbf{z}_j)}{\epsilon^2}\right) \quad (19)$$

where $\epsilon > 0$ is the kernel scale and $d(\cdot, \cdot)$ is the modified Mahalanobis distance (9).

The kernel matrix is then normalized to be row-stochastic

$$P(i, j) = \frac{K(i, j)}{D(i)} \quad (20)$$

where $D(i) = \sum_{j=1}^N K(i, j)$.

From the eigenvalue-decomposition of P we obtain a set of eigenvectors and eigenvalues, $\{\boldsymbol{\psi}^{(\ell)}\}_{\ell=0}^k$, $\{\mu^{(\ell)}\}_{\ell=0}^k$. It was shown in [19], that in this discrete setting, the diffusion maps eigenvectors, $\{\boldsymbol{\psi}^{(\ell)}\}_{\ell=0}^k$, approximate the continuous diffusion maps eigenfunctions, $\{\phi^{(\ell)}\}_{\ell=0}^k$, discussed in the previous subsections. Moreover, it was shown in [23], that the eigenvalues of the discrete diffusion maps algorithm can be used to approximate the eigenvalues of the continuous operator according to $-\lambda^{(\ell)} = \frac{2}{\epsilon} \log \mu^{(\ell)}$.

We rewrite the state and measurement equations in the discrete setting using the diffusion maps coordinates:

$$\boldsymbol{\Psi}_{n+1} = (I + \Lambda) \boldsymbol{\Psi}_n + Q_n^{1/2} \Delta\omega_n \quad (21)$$

$$\mathbf{z}_n = \boldsymbol{\alpha} \boldsymbol{\Psi}_n + R_n^{1/2} \mathbf{v}_n \quad (22)$$

where $\boldsymbol{\Psi}_n = [\psi_n^{(1)}, \psi_n^{(2)}, \dots, \psi_n^{(k)}]$, Λ is a diagonal matrix with $\{-\lambda^{(\ell)}\Delta t\}_{\ell=1}^k$ as its diagonal elements, I is the identity matrix, $Q_n^{1/2}$ is a matrix containing the coefficients of the second term in (18), $\Delta\omega_n$ is a vector of standard normally distributed noise processes, R_n is the covariance of the measurement noise and $\boldsymbol{\alpha}$ is the lift function from the diffusion maps coordinates to the measurements. For

simplicity of notation, we denote $\boldsymbol{\Psi}_n = \boldsymbol{\Psi}(\boldsymbol{\theta}_n)$, omitting the dependence on $\boldsymbol{\theta}_n$. Importantly, note that these system equations are approximately linear, even for highly nonlinear systems. Therefore, using the diffusion maps coordinates, we obtain a virtual system state which linearizes the problem, in an entirely data-driven manner, given only the measurements. In the discrete setting, the lift function is approximated by $\alpha_{i,\ell} = \sum_{n=1}^N z_n^{(i)} \psi^{(\ell)}(\boldsymbol{\theta}_n)$ [22].

The term $\boldsymbol{\alpha} \boldsymbol{\Psi}_n$ in (22) approximates the nonlinear measurement function $g(\boldsymbol{\theta}_n)$. Therefore, the noise term in (22), denoted by $R_n^{1/2} \mathbf{v}_n$, is approximately the measurement noise from (3). However, since the diffusion maps coordinates are constructed based on the noisy measurements, \mathbf{z}_n , the approximation of the measurement function may be inaccurate, leading to non-Gaussian components in $R_n^{1/2} \mathbf{v}_n$. In regard to this issue, we note that in many cases the leading diffusion maps eigenvectors (corresponding to the largest eigenvalues) represent the slow components in the data, and therefore, are only mildly affected by the noise. This property facilitates the use of our devised Kalman model in many applications, as demonstrated in the experimental study in Section IV, even though the model is inaccurate due to the measurement noise. We further note that the noise term in (22) can be used to represent deviations from the true model as described in [24]. However, significant measurement noise may still lead to some deterioration in the performance of the proposed method, as demonstrated in Subsection IV-A.

Due to the linearity of the derived system in (21) and (22) (except for elements in Q_n as discussed in a subsequent paragraph), we can construct a linear Kalman filter based on the diffusion maps coordinates. Using the Kalman filter framework we harness the linear dynamics of the diffusion maps eigenfunctions to improve the approximation of the eigenvectors of the discrete diffusion maps algorithm. Therefore, our framework incorporates the inherent time-dependencies of the system samples into a manifold learning technique.

Our proposed framework is highly related to the method presented in [22], where a linear observer was constructed based on the system equations described in (21) and (22), by exploiting the linearity of the data-driven representation in a similar manner. The Kalman filtering framework improves the linear observer framework in [22] in two main aspects. First, the constructed observer scheme is a deterministic framework which discards the stochastic term of the dynamics, whereas the Kalman filter takes it naturally into account. Second, the Kalman filter provides an adaptive optimal update of the fixed model parameters in [22]. We empirically show in Section IV, that the adaptive parameter update and the stochastic framework significantly improve the state recovery and robustness to noise, and outperforms the observer framework as well as other competing methods.

We note that Q_n is a nonlinear function of the state and induces dependencies between different time-samples of the system noise process, which does not fit the Kalman filter framework. However, in many applications, the leading diffusion maps eigenvectors, which are used as a low-dimensional representation of the system, are slowly varying functions of

the system state. Therefore, the gradient of $\psi^{(\ell)}$ according to θ_n in (18) may be sufficiently small (or approximately constant), allowing for the proper use of the Kalman filter. These properties are demonstrated in simulations in Section IV.

An alternative approach to address the dependence of the driving noise on the state would be using a particle filter instead of the Kalman filter. Particle filters are designed to support a wider range of driving noise distributions [25], [26]. However, our empirical study showed that the Kalman filter significantly outperformed the particle filter, defined in a corresponding manner using the derived model, since the particle filter is sensitive to errors in the estimation of the gradient of $\psi^{(\ell)}$. Improving our methods by appropriately approximating the driving noise remains a subject for future work.

The implemented Kalman filter framework is described as follows:

$$\begin{aligned}\hat{\Psi}_n &= F_n \hat{\Psi}_{n-1} + \kappa_n \left(z_n - H_n F_n \hat{\Psi}_{n-1} \right) \\ P_n &= (I - \kappa_n H_n) (F_n P_{n-1} F_n^T + Q_n) \\ \kappa_n &= (F_n P_{n-1} F_n^T + Q_n) H_n^T \\ &\quad (H_n F_n P_{n-1} F_n^T H_n^T + H_n Q_n H_n^T + R_n)^{-1}\end{aligned}\quad (23)$$

where $\hat{\Psi}_n$ is the state estimate at time n , z_n is the measurement at time n , I is the identity matrix, $F_n = I + \Lambda$ represents the dynamics of the state, $H_n = \alpha$ is the lift function between the measurements and the state calculated in a data driven manner as described after equation (22), Q_n is the covariance matrix of the state driving noise, and R_n is the covariance matrix of the measurement noise. We refer to the proposed Kalman filter framework as the Diffusion Maps Kalman (DMK).

In the experiments presented in Section IV, the covariance matrices, R_n and Q_n , were estimated from the data according to $R_n(k, k) = \text{var}(z^{(k)})$ and $Q_n(k, k) = \text{var}(\lambda^{(k)} \psi^{(k)})$, respectively, where $\psi^{(k)}$ is the k th diffusion maps eigenvector corresponding to eigenvalue $\mu^{(k)}$ and $-\lambda^{(k)} = 2/\epsilon \log \mu^{(k)}$. This empirical choice led to good results in all applications. However, it is possibly an underestimation of the covariance and could be further improved. For example the adaptive estimation for R_n and Q_n described in [27] could be used for unknown system dynamics and measurement function. In addition, this issue will be addressed in future work, where a method for properly estimating the variance of the eigenfunctions in (14) will be devised.

E. Observability and Detectability of the New System Equations

In this subsection, we discuss the conditions under which the proposed discrete linear system in (21) and (22) is observable and detectable.

To address the observability and detectability conditions, we first recall that in the proposed system, the state transition matrix is given by $F = I + \Lambda$, which is a full rank diagonal matrix, and the observation matrix is given by $H = \alpha$, where α is defined in Subsection III-C, after equation (15).

Proposition 1: The system in (21) and (22) is observable if $\forall \ell = 1, \dots, k, \exists p = 1, \dots, m$ such that $\langle z^{(p)}, \psi^{(\ell)} \rangle \neq 0$, where $\langle z^{(p)}, \psi^{(\ell)} \rangle = \sum_{n=1}^N z_n^{(p)} \psi^{(\ell)}(\theta_n)$.

In other words, for observability it is required that the k diffusion maps eigenvectors, used in the construction of the system, are not orthogonal to all measurement coordinates.

Proof: In the discrete and linear Kalman filter setting, a system is observable if there are no vectors $v^{(\ell)} \neq \mathbf{0}$ such that $Fv^{(\ell)} = \gamma^{(\ell)}v^{(\ell)}$ and $Hv^{(\ell)} = \mathbf{0}$, where $v^{(\ell)}$ and $\gamma^{(\ell)}$ denote the ℓ th eigenvector and corresponding eigenvalue of F , respectively [28].

Since F is a diagonal matrix in (21), its eigenvectors, $v^{(\ell)}$, contain 1 at index ℓ and zeros elsewhere. Therefore, $Hv^{(\ell)}$ is the ℓ th column of H , which based on Subsection III-C, corresponds to the following vector:

$$Hv^{(\ell)} = \begin{bmatrix} \langle z^{(1)}, \psi^{(\ell)} \rangle \\ \vdots \\ \langle z^{(m)}, \psi^{(\ell)} \rangle \end{bmatrix} \quad (24)$$

where $z^{(i)}$ denotes the i th measurement coordinate, $\psi^{(\ell)}$, $\ell = 1, \dots, k$, is the ℓ th diffusion maps eigenvector and $\langle \cdot, \cdot \rangle$ denotes the inner product between two vectors. From this derivation, for observability we require that $Hv^{(\ell)} \neq \mathbf{0}$, $\forall \ell = 1, \dots, k$, i.e. that the diffusion maps coordinates are not orthogonal to all measurement coordinates. ■

Proposition 2: If $\exists \ell = 1, \dots, k$, such that $\langle z^{(p)}, \psi^{(\ell)} \rangle = 0$, $\forall p = 1, \dots, m$, then the system is detectable if $1 - \lambda^{(\ell)} \Delta t < 0$, where $-\lambda^{(\ell)}$ is the ℓ th eigenvalue of the continuous operator, approximated using diffusion maps, $\psi^{(\ell)}$ is the corresponding eigenvector and Δt denotes the time step.

Proof: A discrete linear system is detectable, if there is no $v^{(\ell)} \neq \mathbf{0}$ and $\gamma^{(\ell)}$ such that $Fv^{(\ell)} = \gamma^{(\ell)}v^{(\ell)}$, $Hv^{(\ell)} = \mathbf{0}$ with $\gamma^{(\ell)} + (\gamma^{(\ell)})^* \geq 0$, where $()^*$ denotes the complex conjugate [28].

In the constructed system, the eigenvalues, $\gamma^{(\ell)}$, of F are real and equal to the diagonal elements of F , which take the form of $1 - \lambda^{(\ell)} \Delta t$. Therefore, the constructed system is detectable if $1 - \lambda^{(\ell)} \Delta t < 0$, when $\langle z^{(p)}, \psi^{(\ell)} \rangle = 0$, $\forall p = 1, \dots, m$, i.e. $\gamma^{(\ell)} < 0$ for $\psi^{(\ell)}$, for which $Hv^{(\ell)} = \mathbf{0}$. ■

Note that by using only diffusion maps coordinates that are not orthogonal to all system measurements in the construction of the system, both propositions hold. Therefore, an informed choice of the diffusion maps coordinates leads to an observable and detectable system. According to common practice, we take the leading k diffusion maps eigenvectors, corresponding to the largest eigenvalues, since these coordinates typically represent the system and are only mildly affected by the noise. Moreover, the choice of the new system state dimensionality, k , can be motivated by the two propositions, since we expect that eigenvectors which are orthogonal to the measurements, e.g. due to noise, will correspond to smaller eigenvalues, i.e. will have larger indexes ℓ . Empirically, following this procedure led to observable and detectable systems in the experimental study in Section IV.

F. Koopman Operator

Our presented framework is tightly related to the Koopman Operator. Specifically, we show the analogy of the revealed dynamics in the present work to the Stochastic Koopman Operator [8], $(U_{stf})(\theta_n) = \mathbb{E}[f \circ T(\theta_n, \omega_n)]$. Given measurements from some stochastic nonlinear system with state dynamics of the form of (5), we project the problem onto the eigenvector space obtained by diffusion maps. By taking these eigenvectors as observables we obtain a space in which the evolution of the observables is represented by a known linear operator:

$$\begin{aligned} (\tilde{U}_{st}\psi^{(\ell)})(\theta_n) &= \mathbb{E}\left[\psi^{(\ell)} \circ T(\theta_n, \mathbf{u}_n) \mid \Psi_n\right] \\ &= \mathbb{E}\left[\left(1 - \lambda^{(\ell)}\Delta t\right)\psi^{(\ell)}(\theta_n) \right. \\ &\quad \left. + \sqrt{2}\|\nabla_{\theta}\psi^{(\ell)}(\theta_n)\|_2\Delta\omega_n^{(\ell)} \mid \Psi_n\right] \\ &= \left(1 - \lambda^{(\ell)}\Delta t\right)\psi^{(\ell)}(\theta_n) \end{aligned} \quad (25)$$

where \tilde{U}_{st} is analogous to the Stochastic Koopman Operator. Note that in contrast to the standard Stochastic Koopman Operator, \tilde{U}_{st} is conditioned on Ψ_n .

The use of diffusion maps for the approximation of the Koopman Operator is discussed in detail in [18]. In particular a Galerkin method for approximating the eigenfunctions of the Koopman generator using the diffusion maps eigenfunctions and eigenvalues is presented. Furthermore, it was shown that for ergodic systems with pure-point spectra, the eigenfunctions of the Koopman generator can be robustly estimated from finite data using diffusion maps. Importantly, our method is completely different than the method presented in [18], since we approximate the Stochastic Koopman Operator, whereas in [18] the generator is approximated.

In our proposed framework, we combine the constructed linear operator and observables, obtained by diffusion maps, with a Kalman filter. This leads to two main benefits. First, instead of representing the average time evolution of the observables in stochastic systems, we obtain an estimation of individual realizations of trajectories from the measurements. In addition, the Kalman filter compensates for the noise and deviations from the measurements. Second, due to the use of diffusion maps, we obtain a data-driven dimensionality reduction and approximate the Koopman Operator based on a finite set of orthonormal functions, spanning the state space of the system [21].

A related work, combining the Koopman Operator and a Kalman filter is presented in [3]. There, the authors define the Koopman Observer Form (KOF) for noiseless systems and the Koopman Kalman Filter (KKF) for systems with measurement noise. They construct a set of linear update equations based on the eigenvectors and modes of the Koopman Operator which provides a linear filtering framework for nonlinear systems, where the Koopman Operator of a given data-set is approximated using EDMD [10]. The EDMD algorithm requires a dictionary of basis functions which affects the resulting estimations [10]. Conversely, in the proposed work,

we obtain the linear dynamics and observables based on the data, from the diffusion maps algorithm without a pre-defined dictionary. Another difference is the problem setting, which includes in the present work stochastic system dynamics rather than measurement noise only. We compare our proposed framework to the one presented in [3] in Subsection IV-A.

Our framework can also be partially related to the work presented in [4]. There, linear update equations are learned from the data for chaotic systems using concepts from Koopman theory. In order to represent the chaotic dynamics using a finite approximation, a nonlinear forcing term representing the deviation from linearity is added. In our work, we rely on the dynamics of the diffusion maps eigenfunctions, which can be expressed as a sum a linear drift component and a nonlinear stochastic component. The stochastic component represents deviations from the simple linear dynamics and can be used for error analysis as well. In the proposed Kalman filter framework we consider the stochastic component as the system noise.

IV. EXPERIMENTAL RESULTS

In this section we present two simulated examples of object tracking and a real tracking application based on neuronal spiking activity. All three tracking problems are nonlinear with unknown system dynamics and measurement functions, where each example depicts a different measurement modality. We compare our proposed Diffusion Maps Kalman (DMK) to several competing algorithms, which are detailed in the following. We show that our DMK framework leads to improved state estimates compared with *non-parametric algorithms*. In addition, it obtains results which are comparable with *parametric methods*, which, in contrast to DMK, are provided with the system model.

A. Nonlinear Object Tracking

We first present a model with Gaussian noise, where the location of an object in a 2-dimensional space is measured through its radius and azimuth angle. The underlying process, describing the Cartesian position of the object at each time point is given by the following discrete time Langevin equations:

$$\begin{aligned} \Delta\theta_{n+1}^{(1)} &= -\frac{1}{2}\left(\theta_n^{(1)} - 1\right)^3 + \left(\theta_n^{(1)} - 1\right) + \sqrt{2}u_n^{(1)} \\ \Delta\theta_{n+1}^{(2)} &= -\frac{1}{2}\left(\theta_n^{(2)} - 6\right)^3 + \left(\theta_n^{(2)} - 6\right) + \sqrt{2}u_n^{(2)} \end{aligned} \quad (26)$$

where $u_n^{(1)}$ and $u_n^{(2)}$ denote standard Gaussian noises and the drift terms in these equations describe double-well potentials. An example of the resulting 2-dimensional process is presented in Figure 1

The object location is measured in polar coordinates, through the azimuth and radius:

$$\begin{aligned} \phi_n &= \arctan\left(\frac{\theta_n^{(1)}}{\theta_n^{(2)}}\right) \\ r_n &= \sqrt{\left(\theta_n^{(1)}\right)^2 + \left(\theta_n^{(2)}\right)^2} \end{aligned} \quad (27)$$




Fig. 1: Example of the underlying 2-dimensional process.

and the system measurements are created by adding Gaussian noise, $z_n = [\phi_n + v_n^{(\phi)}, r_n + v_n^{(r)}]$, where $v_n^{(\phi)}$ and $v_n^{(r)}$ are Gaussian noise processes with variance σ_ϕ^2 and σ_r^2 , respectively. We created trajectories of 1000 samples with a time-step of $\Delta t = 0.01$ and different Signal-to-Noise (SNR) ratios, by varying σ_ϕ^2 and σ_r^2 .

We applied the diffusion maps algorithm to the measurements, z_n , with ϵ set as the *median* of the distances. The covariance matrices used in the modified Mahalanobis distance (9) were estimated similarly to previous work [29], by calculating the empirical covariance of overlapping windows of 30 time frames, centered at each measurement, i.e. $C_n = \text{cov}(z_{n-N:n+N-1})$ where $N = 15$. Since the covariance matrices are not necessarily full rank, their pseudoinverse was calculated using singular value decomposition (SVD). We then constructed the Diffusion Maps Kalman (DMK) based on the first two largest eigenvalues and corresponding eigenvectors obtained by diffusion maps. The dimensionality of the diffusion maps coordinates was determined based on the existence of a significant spectral gap after the second coordinate. The covariance matrices in the Kalman filter update equations, R_n and Q_n in (23), were estimated from the data, according to $R_n(k, k) = \text{var}(z^{(k)})$ and $Q_n(k, k) = \text{var}(\lambda^{(k)}\psi^{(k)})$, where $-\lambda^{(k)} = 2/\epsilon \log \mu^{(k)}$ and $\mu^{(k)}$ and ψ_k are the k th diffusion maps eigenvalue and eigenvector, respectively.

Using this setting, we evaluated our DMK algorithm, which requires access only to the noisy measurements, z_n , and compared it to a parametric algorithm, the particle filter (PF). For evaluation purposes, the particle filter was provided with the *true system model*, which is considered unknown in our setting.

In addition, we compared our results to three non-parametric algorithms: Gaussian Process filtering (GP) [1], the Kalman filter based algorithm described in [3] (KKF) and the observer framework presented in [22]. Note that there is a fundamental difference between our method and two of the competing non-parametric methods, GP and KKF. Both GP and KKF require a subset of data pairs, $\{\theta_n, z_n\}_{n=1}^N$, in their construction. In our experiments, we used a subset of $N = 100$ and $N = 21$ samples, covering the entire state space, for algorithms GP and KKF, respectively. Moreover, the KKF algorithm requires




Fig. 2: Average and standard deviation of the nRMSE of the clean measurement estimates. The nRMSE values were averaged over 50 realizations of process and noise trajectories.

a choice of a kernel function. Here, we used the same kernel function as in [3]. In contrast, our DMK framework can provide a filtered version of the measurements and a new set of coordinates representing the system characteristics without any information on the underlying state or a specific choice of a kernel function. However, note that if an estimate of a *specific state representation* is required, some alignment between the DMK coordinates and the underlying state may be needed.

Here, we used the mapping defined in (17), α , and obtained an estimate of the clean system measurement, $g(\theta_n)$. Therefore, the comparison between the different algorithms was performed in the measurement domain.

Figure 2 presents plots of the normalized root mean square error (nRMSE) values (in log scale) of the clean measurement estimates, ϕ_n and r_n , obtained by the compared algorithms and the measurement error (denoted by Meas.) as an upper bound. The nRMSE values are presented as a function of the measurement noise level (SNR), where the average and standard deviation of the nRMSE were calculated over 50 different process realizations, for each noise level separately. We note that the extended Kalman filter (EKF), with the true system model, was considered as well. However, the EKF led to results which were similar to the particle filter and was omitted for brevity therefor. Figure 2 depicts that the DMK outperforms all three non-parametric algorithms in all noise levels. In addition, for high SNR, the DMK errors are close to the errors obtained by the parametric particle filter, *which has access to the true system equations*. Figure 3 further demonstrates this and presents an example for the clean measurement estimation obtained by DMK (in blue) and by the particle filter (in gray), compared with the true clean measurements (in dotted black) and the noisy measurements (gray 'x'), with SNR=1. Plot (a) presents coordinate ϕ_n and plot (b) presents coordinate r_n . This figure depicts that the DMK estimation closely follows the particle filter estimation.

Note that for lower SNR values the DMK estimation de-

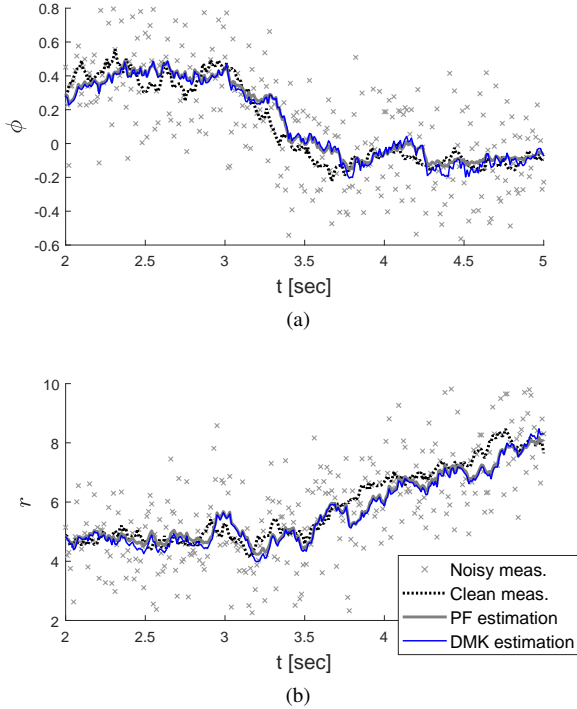


Fig. 3: Example of a trajectory of the filtered measurement coordinates using the DMK algorithm (in blue) compared with the PF estimation (in gray), the true clean measurements (in dotted black) and the noisy measurements (gray 'x').

teriorates in comparison with the particle filter. This is due to errors in the inferred model, caused by the high noise. As described in Subsection III-B and Subsection III-C, the new system model, which is derived based on diffusion maps, is accurate only for the noiseless case. The noise affects the modified Mahalanobis distance and, as a result, the obtained diffusion maps coordinates and eigenvalues contain errors. Since the eigenvalues are used in the model dynamics, these errors affect the estimation quality of our framework. Nevertheless, DMK works quite well in noisy situations, especially in comparison with the non-parametric frameworks.

In order to demonstrate the extent of the estimation deterioration due to noise, we compared the DMK algorithm, which was constructed based on the noisy measurements, and a modified DMK algorithm, where the dynamics (diffusion maps eigenvalues) were obtained by applying diffusion maps to the *clean* measurements and calculating the corresponding “clean dynamics”. The resulting nRMSE for different SNR values are presented in Figure 4, where the average and standard deviation of the nRMSE were calculated based on 50 process and noise realizations for each SNR value. This figure presents the DMK based on the noisy measurements in blue, the DMK with the “clean dynamics” (λ) in cyan and the particle filter, as a baseline, in gray. As expected, for high SNR, both the noisy dynamics and the clean dynamics lead to a similar result, whereas for low SNR, the clean dynamics improve the result. This indicates that our method is indeed affected by the measurement noise. However, note that even with these model errors, our method still outperforms the competing non-

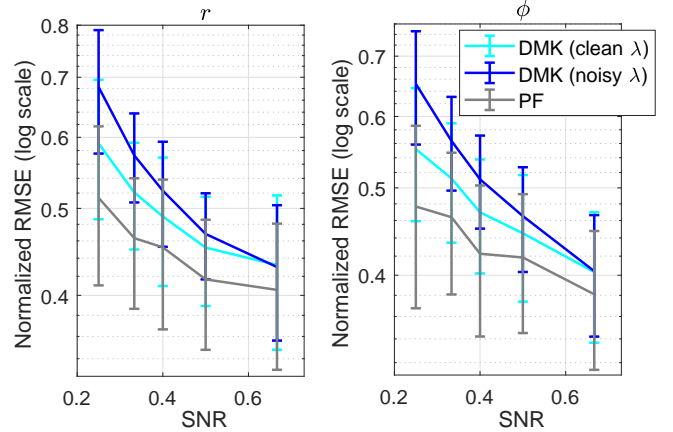


Fig. 4: Average and standard deviation of the nRMSE of the clean measurement estimates, over 50 process and noise realizations. A comparison between the DMK algorithm constructed based on the noisy measurements (in blue) the DMK with the “clean dynamics” (in cyan) and the PF (in gray).

parametric algorithms, as depicted in Figure 2.

To complete the analysis of this example, we compare the convergence rates of DMK based on the noisy measurements, DMK with the “clean dynamics” and the convergence rates of the parametric and non-parametric algorithms.

Figure 5 and Figure 6 present the average asymptotic RMSE (aRMSE), i.e. $\sqrt{\frac{1}{M} \sum_{k=1}^M \left\| \hat{z}_n^{(k)} - \zeta_n^{(k)} \right\|_2^2}$, where M is the number of process and noise realizations, $M = 50$ in both figures, $\hat{z}_n^{(k)}$ denotes an estimation of the measurements (filtered) and $\zeta_n^{(k)}$ denotes the true clean measurements at time n in realization k . Figure 5 presents a comparison between the DMK algorithm based on the noisy measurements (in blue), the DMK algorithm with the “clean dynamics” (in cyan) and the particle filter (in gray), for two SNR values. Plot (a) presents the aRMSE values for measurement coordinate ϕ_n and plot (b) presents the aRMSE values for measurement coordinate r_n . This figure depicts that the convergence rate of DMK is not affected by the noise. In addition, the convergence rate of DMK is either comparable (in plot (a)) or faster (in plot (b)) compared with the convergence rate of the particle filter, for both SNR values.

Figure 6 presents a comparison of the asymptotic RMSE values of the cleaned measurements obtained by the DMK algorithm (in blue), by the two competing non-parametric algorithms, GP and KKF (in light gray and in black, respectively) and by the particle filter (in dark gray), for two different SNR values. Plots (a) and (c) were created using an SNR of 0.67 and plots (b) and (d) were created using an SNR of 0.18. This figure depicts that the convergence rates of all the non-parametric algorithms are similar in both measurement coordinates and for both SNR values.

B. Non-Gaussian Nonlinear Object Tracking

We now present an example which is based on the setting given in [22], [29]. In this setting, the location of a radiating object moving on a 3D sphere is estimated based on




Fig. 5: Asymptotic RMSE, averaged over 50 realizations, of the measurements estimation using the DMK algorithm constructed based on the noisy measurements, colored in blue, the DMK algorithm with the “clean dynamics”, colored in cyan, and the particle filter, colored in gray.

measurements from three sensors, $\mathbf{s}^{(1)}, \mathbf{s}^{(2)}, \mathbf{s}^{(3)}$, which are modeled as “Geiger Counters”. The movement of the object is defined by two underlying Langevin processes, describing the elevation and azimuth angles as follows

$$\Delta\theta_{n+1}^{(1)} = \left(\frac{\pi}{2} \cdot c - c \cdot \theta_n^{(1)}\right) + u_n^{(1)} \quad (28)$$

$$\Delta\theta_{n+1}^{(2)} = \left(\frac{\pi}{5} \cdot c - c \cdot \theta_n^{(2)}\right) + u_n^{(2)} \quad (29)$$

where b is the diffusion coefficient and c is the drift rate parameter. In our experiments, b was set to 0.01 and c was varied between 0.1 and 1 to simulate different trajectory properties.

The 3D location of the object at the n th time step is computed by:

$$\begin{aligned} x_n^{(1)} &= \cos(\theta_n^{(2)}) \sin(\theta_n^{(1)}) \\ x_n^{(2)} &= \sin(\theta_n^{(2)}) \sin(\theta_n^{(1)}) \\ x_n^{(3)} &= \cos(\theta_n^{(1)}) \end{aligned} \quad (30)$$

We mark the 3D position of the object by $\mathbf{x}_n = [x_n^{(1)}, x_n^{(2)}, x_n^{(3)}]$.

The system measurements are given by three Poisson processes, with a rate parameter that is based on the 3D location of the object.

$$y_n^{(j)} \sim Pois(r_n^{(j)}) \quad j = 1, 2, 3 \quad (31)$$

where $r_n^{(j)} = \exp(-\|\mathbf{s}^{(j)} - \mathbf{x}_n\|)$.

Finally, a Poisson noise process with a fixed rate parameter, $v_n^{(j)} \sim Pois(\lambda_v)$, is added to each sensor

$$z_n^{(j)} = y_n^{(j)} + v_n^{(j)} \quad (32)$$

where $z_n^{(j)}$ are the accessible system measurements.

Note that the presented setting is non-Gaussian and therefore the Kalman filter assumptions are not held. However, we show that DMK still provides good results, due to the incorporation of time dependencies, and is either better or comparable to the observer framework presented in [22].

In order to obtain an estimated representation of the system state (the azimuth and elevation angles) from the noisy measurements, $\{z_n\}_{n=1}^N$, we apply DMK.

We simulated 300,000 time samples of the two underlying angles $\theta_n^{(1)}, \theta_n^{(2)}$ with $\Delta t = 0.5$ and constructed the measurements according to (30) and (32). After obtaining the system measurements z_n , we first performed a pre-processing stage, similarly to [22], [29]. This includes constructing histograms for overlapping frames of 60 time-samples of z_n and then calculating the modified Mahalanobis distance (9) between pairs of histograms. The covariance matrices for the modified Mahalanobis distance between the measurements were calculated similarly to Subsection IV-A, based on the empirical covariance of overlapping windows of 20 time frames, centered at each measurement. At this point we are left with $300,000/60 = 5000$ system measurements.

The diffusion maps algorithm, described in Subsection III-B, is applied to the measurements using the calculated Mahalanobis distance, with an empirical choice of ϵ as the median of the Euclidean distances. We obtain a set of eigenvectors and eigenvalues representing intrinsic properties of the system. However, these eigenvectors do not necessarily correspond to the true system state, θ_n , and can represent some linear combination of the state coordinates [30]. Therefore, for evaluation purposes, we perform a linear regression on 3 eigenvectors, corresponding to the largest eigenvalues, and the true system state, based on 100 samples. Note that in this specific example, the linear regression provides a good representation of the underlying state using only a few diffusion maps coordinates. This is due to the choice of the underlying state equation and the lift function which places the data on a sphere. This system was specifically chosen in this application, since the measurements (spike trains) cannot be easily interpretable, in contrast to the example in Subsection IV-A. In Subsection IV-C, a more complicated system is presented, where real neuronal spiking activity is analyzed. For such data, the underlying state equation and system model are completely unknown, and indeed, there, a larger number of coordinates is required to obtain a good representation of the true underlying state (the animal position).

Based on the resulting eigenvectors and eigenvalues we construct the Kalman filter, described in Subsection III-D. The covariance matrices of the measurement noise, R_n , and the state noise, Q_n , were estimated from the data (according to the variance of the histogram measurements and the covariance of the obtained eigenvectors and eigenvalues) as described in Subsection IV-A.

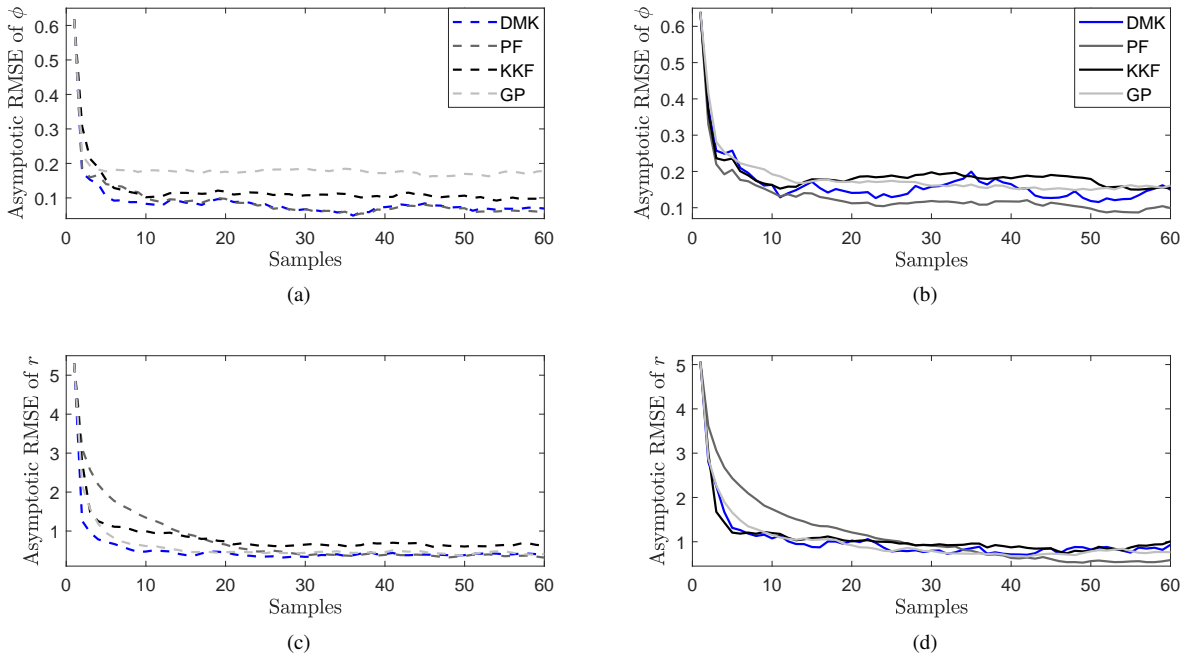


Fig. 6: Asymptotic RMSE, averaged over 50 realizations, of the measurements estimation using the DMK (in blue), the PF (in dark gray), KKF, (in black) and GP (in light gray). The SNR was set to 0.67 in (a) and (c) and to 0.18 in (b) and (d). Plots (a) and (b) present the asymptotic RMSE values for measurement coordinate ψ_n and plots (c) and (d) present the asymptotic RMSE values for measurement coordinate r_n .

We compare our results to the observer framework, described in [22], with a choice of $\gamma = 0.01$ (which led to the best results in this case). In Figure 7, a comparison between the DMK, the observer framework and the diffusion maps coordinates (without additional analysis) is presented. This figure contains 6 identical scatter plots, each presenting the true underlying angles, $\theta_n^{(1)}, \theta_n^{(2)}$. Each plot is colored according to a different coordinate, plots (a) and (d) are colored according to the first and second estimated coordinates of the observer framework, plots (b) and (e) are colored according to the DMK estimation and plots (c) and (f) are colored according to the diffusion maps coordinates (DM). The color gradients in Figure 7 depict that the DMK significantly improves the estimation of the two underlying angles, compared with the diffusion maps coordinates and the observer framework. Moreover, the coordinates obtained by the observer framework suffer from inaccuracies at the boundaries of the data. This is visible for example, in plot (a), when $\theta_n^{(1)} < 0.5$ and is due to the inaccuracy of the linear lift function at the boundaries. These inaccuracies are not apparent in the DMK coordinates which recover the true underlying angles more accurately even at the boundaries of the data.

Figure 8 presents the correlation between the DMK coordinates (after linear regression based on 100 samples) and the true underlying states, $\theta^{(1)}$ and $\theta^{(2)}$, colored in blue. The average and standard deviation over 50 realizations are presented. For comparison, we present the correlations of the observer coordinates, in green, and the diffusion maps coordinates, in black. These plots depict that even though the

Kalman model is inaccurate in this example, the DMK obtains results which are either better or comparable to the observer, in different drift-diffusion rate ratios.

C. Location Tracking based on Rat Hippocampal Neuronal Activity

To demonstrate our framework on real data, we consider a publicly available data set [31]. This data set contains simultaneous recordings of hippocampal regions *CA1* and *CA3* or regions *CA1* and *MEC* in 9 rats during a spatial alternation task in a W-shaped maze. We show that when applying DMK to the neuronal spiking activity, it reveals a meaningful representation which is related to the 2D position of the rat.

A related work [32] addressed this task using a *parametric* framework that recovers a latent state governing the spike rates. Their framework is based on a Poisson process that generates the spiking data and two Gaussian processes, which model the latent state dynamics and the mapping between the latent state and the firing rate. This parametric method assumes a specific model and is especially suited for neuronal spike train data, whereas our method is non-parametric and can be applied to a larger class of problems.

From the available data set [31], we focus on experiments in which the animals were actively moving in the maze. In addition, since some recordings include only a small number of active neurons, we restrict our analysis to experiments which contain more than 20 active neurons and animals which have more than 3 such experiments (for evaluation of error




Fig. 7: Scatter plots of the azimuth and elevation angles, colored by the coordinate estimates. The plots are colored according to the the first and second state estimates of the observer framework (in (a) and (d)), the first and second state estimates of our suggested DMK filter (in (b) and (e)) and the first and second coordinates obtained by diffusion maps (in (c) and (f)).




Fig. 8: Correlation between the DMK coordinates and the true underlying system state coordinates, in blue, compared with the observer coordinates, in green, and the diffusion maps coordinates, in black.

statistics). The remaining data include six animals (abbreviated Bon, Ten, Cor, Eig, Fra and Mil) with 10 to 24 experiments each. The length of each experiment ranged between 7 to 15 minutes.

We apply the proposed method to the recorded neuronal spiking data and construct a new data-driven representation. We show that this new representation is highly related to the true position of the animal.

For this purpose, we first bin the data to obtain a histogram for each neuron depicting the spike rate over time. The histogram bin sizes were chosen empirically, such that at least one neuron is active at any time frame (bin). The bin sizes

ranged between 1 – 2.6 seconds (chosen differently in different animals). The bin sizes, number of experiments per animal and number of active neurons used in our analysis are summarized in Table I.

The diffusion maps algorithm was applied to the obtained histograms, which are treated as measurements in the construction of the affinity matrix described in (19). The covariance matrices for the modified Mahalanobis distance were estimated based on overlapping windows of 15 time frames (histogram bins), and the kernel scale, ϵ , was empirically set to be 3 times the median of the distances, to avoid outliers in the diffusion maps coordinates. The affinity matrix was normalized according to (20) and its first 20 (non-trivial) eigenvectors, corresponding to the largest eigenvalues, were computed. The DMK was then applied to these eigenvectors, resulting in a new representation of the data, $\hat{\Psi}_n = [\psi_n^{(1)}, \dots, \psi_n^{(20)}]$. The DMK covariance matrices were estimated based on the data and the calculated eigenvectors, similarly to Subsection IV-A and Subsection IV-B, where $Q_n(k, k) = \text{var}(\lambda_k \psi_k)$ and R_n was set to be a diagonal matrix with the variance of the histograms (measurements) on its diagonal.

Each experiment in each animal was analyzed and evaluated separately in order to avoid batch effects.

An example for the resulting representation of one experiment is presented in Figure 9. This figure presents 6 identical plots containing the true x-y positions of the animal (marked by filled circles), where each plot is colored by a different DMK coordinate, according to $\ln|\hat{\psi}_i|$. We note that in this example, only 6 eigenvectors were used as input to the DMK algorithm, since these 6 coordinates provided a good representation for the position of the animal. In other experiments, 6 eigenvectors were not always sufficient and

	Bon	Ten	Cor	Eig	Fra	Mil
Histogram bin size (sec)	1	1	2.6	1.6	1	1
Number of experiments	24	14	10	12	18	12
Number of neurons per exp.	33 – 67	42 – 60	20 – 26	29 – 59	40 – 69	25 – 38

TABLE I: Hippocampal neuronal activity - analysis parameters

therefore, for consistency, 20 eigenvectors were used in all experiments in the following analysis.

Figure 9 depicts that the DMK algorithm provides a meaningful representation for the animal location, since different 2D locations are highlighted in different coordinates (colored in yellow). By combining information from all coordinates, the 2D location can be inferred. For example, the middle arm in the W-shaped maze, $(x, y) \approx (160, 110)$, is represented by high values (mostly) in coordinates $\hat{\psi}_2$ and $\hat{\psi}_3$, whereas the lower arm, $(x, y) \approx (160, 80)$, is represented by high values (mostly) in coordinates $\hat{\psi}_2$ and $\hat{\psi}_5$. In addition, Figure 9 demonstrates that the obtained representation covers different regions in the 2D space in a relatively smooth manner.

We emphasize that the DMK coordinates presented in Figure 9 were obtained in a completely data-driven manner and with minimal model assumptions.

In order to quantitatively assess the quality of the new representation obtained by the DMK, we divided the data in each experiment into a train set, consisting of 80% of the data, and a test set, consisting of 20% of the data and performed cross-validation. We performed linear regression between the DMK coordinates representing the data in the train set and the true position of the animal and then used the regression parameters to estimate the animal’s position based on the DMK representation of the test set.

An example for the resulting position estimation based on the test set of one experiment is presented in Figure 10. This figure presents the estimated x position (top plot) and estimated y position (bottom plot) based on the DMK coordinates (in blue) compared with the true position (in black). Figure 10 depicts that after the linear regression, the estimated position based on the DMK closely follows the true position of the animal.

The corresponding 2D position of the animal in this experiment is presented in Figure 11, where the position estimation based on the DMK coordinates is marked by blue ‘x’ and the true 2D position is marked by black circles. This plot depicts that after linear regression, the DMK coordinates are highly related to the 2D position of the animal in most locations. Note that the edge of the lower arm of the W-shaped maze is not represented properly by the DMK coordinates. This is consistent with the result in Figure 9, where none of the coordinates captures this specific location properly.

For comparison, we applied linear regression to two additional data representations: (1) diffusion maps with no additional processing and (2) principal component analysis (PCA) applied to the spike rate histograms.

We calculated the correlation between the estimated position and the true position for each experiment in all animals.

Figure 12 presents the correlation values of the x and y positions, for each animal separately, based on the train set and on the test set. The average and standard deviation of the correlation values were calculated over 5-fold cross-validation, where the data was divided into 5 consecutive segments. This figure depicts that the DMK coordinates provide a meaningful representation that relates to the true position of the animal, since simply by linear regression, the DMK coordinates give rise to a good estimation of the true position. In addition, this representation is significantly better than the representations obtained by the diffusion maps coordinates and by PCA.

Note that the linear regression here was performed solely for a quantitative evaluation of the constructed coordinates. Our main result here is that the data driven DMK coordinates contain meaningful information regarding the location of the animal.

We conclude by noting that in this example, the true system model is completely unknown and is not necessarily compatible with the assumed process form in (5). However, we demonstrated that our method can still be used to obtain a meaningful representation in this application.

V. CONCLUSION

In this work we addressed the analysis of high-dimensional nonlinear stochastic dynamical systems with measurement noise, and presented a non-parametric filtering framework in which a data-driven linear Kalman filter is constructed based on diffusion maps coordinates, utilizing their inherent dynamics and properties. We showed that the presented framework recovers new coordinates that capture meaningful properties of the system given only a set of noisy measurements and with no further knowledge on the system properties. These coordinates may not be directly related to the underlying system state, yet they can be employed for filtering the measurements. To obtain an estimate of a *specific* underlying state, some alignment is required based on a training set, for example, using linear regression.

The theoretical justification of the proposed framework required few assumptions. First, the derivation of the recovered system dynamics was based on the assumption that the underlying system dynamics evolve according to gradient flows with constant diffusion. Second, for compatibility of the recovered system with the Kalman filtering framework, we assumed that the leading diffusion maps coordinates are slowly evolving functions of the system state and are only mildly affected by the measurement noise. Several studies, e.g. [21], [33] and [34], have shown that based on properties of the diffusion maps coordinates, this assumption commonly holds. Third, we specified the conditions under which the




Fig. 9: True x-y positions of the animal in one experiment (filled circles), colored according to 6 DMK coordinates ($\ln|\hat{\psi}_i|$) calculated based on the spike rate histograms




Fig. 10: Example of the estimated x-y positions of the animal in one experiment. The estimations are presented for the test-set, after applying linear regression to the DMK coordinates.




Fig. 11: Estimated 2D position in one experiment. The estimation is presented for the test-set after applying linear regression to the DMK coordinates




Fig. 12: Correlations between the estimated coordinates (after linear regression) and the true x-y positions, on the train-set and the test-set, for a 5-fold cross-validation.

devised system is observable and detectable, and showed that these conditions hold for an informed choice of the diffusion maps coordinates. Although these assumptions are required for the theoretical derivations, the experimental results depict that our framework performs well in comparison to other non-parametric algorithms, even when these assumptions are not fully met. We showed that the proposed method obtains an improved representation of dynamical systems compared with the diffusion maps coordinates and other non-parametric methods.

In the future, we plan to extend our method and address multi-modal data-sets arising from stochastic dynamical systems. We will devise methods for revealing the underlying

common dynamics based on measurements from different sensors.

VI. ACKNOWLEDGEMENTS

We would like to thank the Associate Editor and the anonymous reviewers for their helpful comments and suggestions.

REFERENCES

- [1] M. P. Deisenroth, R. D. Turner, M. F. Huber, U. D. Hanebeck, and C. E. Rasmussen, "Robust filtering and smoothing with gaussian processes," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1865–1871, 2012.
- [2] Y. Wang, M. A. Brubaker, B. Chaib-draa, and R. Urtasun, "Bayesian filtering with online gaussian process latent variable models," in *UAI*, 2014, pp. 849–857.
- [3] A. Surana and A. Banaszuk, "Linear observer synthesis for nonlinear systems using koopman operator framework," *IFAC-PapersOnLine*, vol. 49, no. 18, pp. 716–723, 2016.
- [4] S. L. Brunton, B. W. Brunton, J. L. Proctor, E. Kaiser, and J. N. Kutz, "Chaos as an intermittently forced linear system," *Nature Communications*, vol. 8, 2017.
- [5] T. Berry, D. Giannakis, and J. Harlim, "Nonparametric forecasting of low-dimensional dynamical systems," *Physical Review E*, vol. 91, no. 3, pp. 032915, 2015.
- [6] F. Tan, W. Lohmiller, and J. J. Slotine, "Analytical slam without linearization," in *Robotics Research*, pp. 89–105. Springer, 2018.
- [7] R.E. Kalman, "A new approach to linear filtering and prediction problems," *Trans. ASME J. Basic Eng.*, vol. 82, pp. 34–45, 1960.
- [8] I. Mezić, "Spectral properties of dynamical systems, model reduction and decompositions," *Nonlinear Dynamics*, vol. 41, no. 1, pp. 309–325, 2005.
- [9] M. Budišić, R. Mohr, and I. Mezić, "Applied koopmanism a," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 22, no. 4, pp. 047510, 2012.
- [10] 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, December 2015.
- [11] I. Mezić and A. Banaszuk, "Comparison of systems with complex behavior," *Physica D: Nonlinear Phenomena*, vol. 197, no. 1, pp. 101–133, 2004.
- [12] R. Mohr and I. Mezić, "Construction of eigenfunctions for scalar-type operators via laplace averages with connections to the koopman operator," *arXiv preprint arXiv:1403.6559*, 2014.
- [13] S. M. Ulam, "Problems of modern mathematics. science editions, new york, 1964. originally published as: A collection of mathematical problems," 1960.
- [14] S. Klus, P. Koltai, and C. Schütte, "On the numerical approximation of the perron-frobenius and koopman operator," *arXiv preprint arXiv:1512.05997*, 2015.
- [15] N. Črnjarić-Zic, S. Maćešić, and I. Mezić, "Koopman operator spectrum for random dynamical system," *arXiv preprint arXiv:1711.03146*, 2017.
- [16] T. Berry and J. Harlim, "Nonparametric uncertainty quantification for stochastic gradient flows," *SIAM/ASA Journal on Uncertainty Quantification*, vol. 3, no. 1, pp. 484–508, 2015.
- [17] F. Hamilton, T. Berry, and T. Sauer, "Ensemble kalman filtering without a model," *Physical Review X*, vol. 6, no. 1, pp. 011021, 2016.
- [18] D. Giannakis, "Data-driven spectral decomposition and forecasting of ergodic dynamical systems," *Applied and Computational Harmonic Analysis*, in press, 2017.
- [19] R. Coifman and S. Lafon, "Diffusion maps," *Appl. Comput. Harmon. Anal.*, vol. 21, pp. 5–30, Jul. 2006.
- [20] F. Takens, "Dynamical systems and turbulence (detecting strange attractors in fluid turbulence)," 1981.
- [21] R.R. Coifman, I.G. Kevrekidis, S. Lafon, M. Maggioni, and B. Nadler, "Diffusion maps, reduction coordinates, and low dimensional representation of stochastic systems," *Multiscale Modeling & Simulation*, vol. 7, no. 2, pp. 842–864, 2008.
- [22] T. Shnitzer, R. Talmon, and J. J. E. Slotine, "Manifold learning with contracting observers for data-driven time-series analysis," *IEEE Transactions on Signal Processing*, vol. 65, no. 4, pp. 904–918, 2017.
- [23] A. Singer and R. Coifman, "Non-linear independent component analysis with diffusion maps," *Appl. Comput. Harmon. Anal.*, vol. 25, pp. 226–239, 2008.
- [24] Y. Bulut and O. Bayat, "Kalman filtering with model uncertainties," in *Topics in Modal Analysis I, Volume 5*, pp. 447–455. Springer, 2012.
- [25] M.S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, pp. 174–188, 2003.
- [26] Arnaud Doucet, Simon Godsill, and Christophe Andrieu, "On sequential monte carlo sampling methods for bayesian filtering," *Statistics and computing*, vol. 10, no. 3, pp. 197–208, 2000.
- [27] T. Berry and T. Sauer, "Adaptive ensemble kalman filtering of non-linear systems," *Tellus A: Dynamic Meteorology and Oceanography*, vol. 65, no. 1, pp. 20331, 2013.
- [28] Eduardo D Sontag, *Mathematical control theory: deterministic finite dimensional systems*, vol. 6. Springer Science & Business Media, 2013.
- [29] R. Talmon and R.R. Coifman, "Empirical intrinsic geometry for nonlinear modeling and time series filtering," *Proc. Nat. Acad. Sci.*, vol. 110, no. 31, pp. 12535–12540, 2013.
- [30] C.J. Dsilva, R. Talmon, R.R. Coifman, and I.G. Kevrekidis, "Parsimonious representation of nonlinear dynamical systems through manifold learning: A chemotaxis case study," *Applied and Computational Harmonic Analysis*, 2015, In Press.
- [31] M Karlsson, M Carr, and LM Frank, "Simultaneous extracellular recordings from hippocampal areas ca1 and ca3 (or mec and ca1) from rats performing an alternation task in two w-shaped tracks that are geometrically identically but visually distinct. crcns. org," 2015.
- [32] Anqi Wu, Nicholas G Roy, Stephen Keeley, and Jonathan W Pillow, "Gaussian process based nonlinear latent structure discovery in multi-variate spike train data," in *Advances in Neural Information Processing Systems*, 2017, pp. 3496–3505.
- [33] A. Singer, R. Erban, I. G. Kevrekidis, and R. Coifman, "Detecting intrinsic slow variables in stochastic dynamical systems by anisotropic diffusion maps," *PNAS*, vol. 106, no. 38, pp. 16090–16095, 2009.
- [34] C.J. Dsilva, R. Talmon, C.W. Gear, R.R. Coifman, and I.G. Kevrekidis, "Data-driven reduction for a class of multiscale fast-slow stochastic dynamical systems," *SIAM Journal on Applied Dynamical Systems*, vol. 15, no. 3, pp. 1327–1351, 2016.