



Contents lists available at ScienceDirect

Applied and Computational Harmonic Analysis

www.elsevier.com/locate/acha

Case Studies

Parsimonious representation of nonlinear dynamical systems through manifold learning: A chemotaxis case study

Carmeline J. Dsilva^a, Ronen Talmon^b, Ronald R. Coifman^c,
Ioannis G. Kevrekidis^{a,d,*}^a Department of Chemical and Biological Engineering, Princeton University, Princeton, NJ 08540, USA^b Department of Electrical Engineering, Technion – Israel Institute of Technology, Haifa 3200003, Israel^c Department of Mathematics, Yale University, New Haven, CT 06520, USA^d Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08540, USA

ARTICLE INFO

Article history:

Received 26 May 2015

Accepted 19 June 2015

Available online xxxx

Communicated by Charles K. Chui

Keywords:

Diffusion maps

Repeated eigendirections

Chemotaxis

ABSTRACT

Nonlinear manifold learning algorithms, such as diffusion maps, have been fruitfully applied in recent years to the analysis of large and complex data sets. However, such algorithms still encounter challenges when faced with real data. One such challenge is the existence of “repeated eigendirections,” which obscures the detection of the true dimensionality of the underlying manifold and arises when several embedding coordinates parametrize the same direction in the intrinsic geometry of the data set. We propose an algorithm, based on local linear regression, to automatically detect coordinates corresponding to repeated eigendirections. We construct a more parsimonious embedding using only the eigenvectors corresponding to unique eigendirections, and we show that this reduced diffusion maps embedding induces a metric which is equivalent to the standard diffusion distance. We first demonstrate the utility and flexibility of our approach on synthetic data sets. We then apply our algorithm to data collected from a stochastic model of cellular chemotaxis, where our approach for factoring out repeated eigendirections allows us to detect changes in dynamical behavior and the underlying intrinsic system dimensionality directly from data.

© 2015 Elsevier Inc. All rights reserved.

1. Introduction

In recent years, data mining algorithms have proven useful for many disciplines and applications [1–8]. For multiscale dynamical systems in particular, data-driven methodologies are essential for assisting in model reduction when simple macroscopic models cannot be analytically obtained due to the system complexity [9–12]. For simulations or experimental observations of such systems, the detailed, microscale evolving system state is very high-dimensional, and the reduction to useful macroscale dynamics is not obvious. In

* Corresponding author at: Department of Chemical and Biological Engineering and Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08540, USA.

E-mail address: yannis@princeton.edu (I.G. Kevrekidis).

such cases, data obtained from observations and/or simulations of the dynamical system combined with data mining methodologies can lead to a low-dimensional description which not only provides insight into the underlying dynamics, but also serves as a first step in constructing macroscale models consistent with the observed microscale behavior.

Several manifold learning algorithms obtain parametrizations of the data through the spectral analysis of a Laplace operator [13–16]. The data is then embedded in a new low-dimensional coordinate system given by the eigenvectors of this operator. The premise is that these coordinates, obtained in a data-driven manner, are the right macroscopic “observables”, i.e., the variables which parametrize the macroscale dynamical behavior of the system, thus enabling the construction of a reduced macroscale model.

Such manifold learning algorithms were initially applied to synthetic data sets, to illustrate their geometric properties and flexibility [14,17]. More recently they have been applied to experimental as well as simulation data, enabled by advances in data representation (observers) and metrics [18–23]. A nontrivial shortcoming of these methods is that, from a geometric perspective, not all eigenvectors are guaranteed to parametrize unique directions within the data; some eigenvectors involve “repeated eigendirections” which are associated with the same coordinate in the intrinsic geometry of the data. Identifying these eigenvectors is critical for obtaining a parametrization of the system which captures the true dimensionality of the macroscale dynamics. This task is often done manually, and few methods have been proposed to automate the identification of the unique eigendirections [24].

In this paper, we propose an algorithm to automatically identify such unique eigendirections using local linear regression [25]. We first demonstrate and validate our algorithm on synthetic examples, where a closed-form solution for the eigenvectors and eigenvalues of the Laplace operator is known. We then consider a simulation data set from a stochastic dynamical system modeling cellular chemotaxis [26]. Recent advances in observers and metrics, coupled with the proposed approach for identifying the unique eigendirections, provide a data analysis pipeline which successfully analyzes this simulation in a purely data-driven manner. We will show that this pipeline allows us to detect changes in the dimensionality of the underlying macroscale dynamics, which are related to changes in the regime/mode of the system behavior.

2. Manifold learning based on Laplace operators

Let $z(1), \dots, z(m) \in \mathbb{R}^n$ denote m state observations sampled from an evolving autonomous dynamical system. We assume the n -dimensional observations $z(i)$ lie on a d -dimensional manifold \mathcal{M}_d , where $d < n$. We therefore consider the problem of parametrizing the continuous d -dimensional manifold \mathcal{M}_d embedded in \mathbb{R}^n from data. For linear hyperplanes, the principal axes parametrize the manifold. However, in the case when the manifold is nonlinear, a set of appropriate coordinates is not readily apparent. By using a particular manifold learning technique based on the construction of a Laplace operator, namely, *diffusion maps*, we will show how to extract a d -dimensional parametrization of the observations which is consistent with the geometry of the manifold [13,14,16].

2.1. Eigenfunctions of the Laplace–Beltrami operator

In recent years it has been noted that the eigenfunctions of the *continuous* Laplace–Beltrami operator provide “good” coordinates for a manifold [27]. We start by considering such a continuous setting, where rigorous analysis is possible for specific examples. Using such an example, we will demonstrate that the existence of repeated eigendirections is inherent to the manifold learning setup based on Laplace operators, and that the identification of the unique eigendirections requires some additional effort. The fact that these challenges arise in the continuous setting implies that, even in the limit of infinite data, such repeated eigendirections still pose a problem for data analysis.

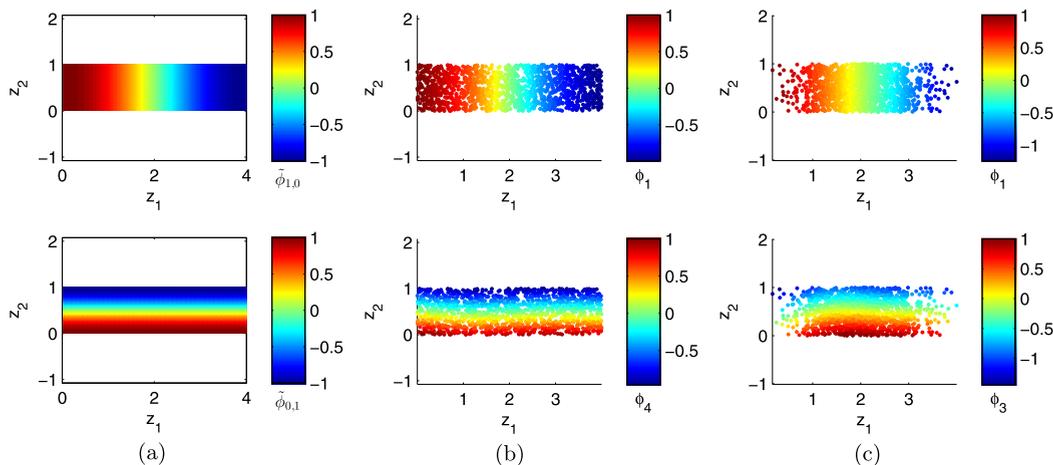


Fig. 1. (a) Two-dimensional continuous strip colored by the eigenfunctions $\tilde{\phi}_{1,0} = \cos(\pi z_1/L_1)$, and $\tilde{\phi}_{0,1} = \cos(\pi z_2/L_2)$. (b) Data, uniformly sampled from a two-dimensional strip, colored by the first and fourth (non-trivial) eigenvectors of the discrete Laplacian. (c) Data, sampled from a Gaussian distribution in z_1 and sampled uniformly in z_2 , colored by the first and third (non-trivial) eigenvectors of the discrete Laplacian. Note that in all cases we uncover parametrizations which are one-to-one with z_1 and z_2 on the domain.

To review and illustrate why these eigenfunctions provide appropriate coordinates, consider a two-dimensional strip with edge lengths L_1 and L_2 . The eigenvalues of the Laplace–Beltrami operator with Neumann boundary conditions are given by

$$\tilde{\mu}_{k_1,k_2} = \left(\frac{k_1\pi}{L_1}\right)^2 + \left(\frac{k_2\pi}{L_2}\right)^2 \tag{1}$$

for $k_1, k_2 = 0, 1, 2, \dots$, and the corresponding eigenfunctions are

$$\tilde{\phi}_{k_1,k_2} = \cos\left(\frac{k_1\pi z_1}{L_1}\right) \cos\left(\frac{k_2\pi z_2}{L_2}\right) \tag{2}$$

where z_1 and z_2 denote the two coordinates of the strip [16]. We note that the eigenfunctions $\tilde{\phi}_{1,0} = \cos(\pi z_1/L_1)$ and $\tilde{\phi}_{0,1} = \cos(\pi z_2/L_2)$ are one-to-one with the z_1 and z_2 coordinates, respectively, on the domain, and therefore provide a parametrization of the underlying manifold (see Fig. 1a). Furthermore, the corresponding eigenvalues $\tilde{\mu}_{1,0}$ and $\tilde{\mu}_{0,1}$ provide an estimate of the relative magnitude of L_1 versus L_2 : as the ratio between L_1 and L_2 increases, the gap between $\tilde{\mu}_{1,0}$ and $\tilde{\mu}_{0,1}$ also increases (this will be discussed further in Section 3.2). The analytic form of the eigenfunctions in (2) illustrates the two issues we address in this paper. First, z_1 and z_2 are not necessarily decoupled in subsequent eigenfunctions, and a proper parametrization of the manifold is not necessarily given by the d eigenfunctions associated with the smallest d eigenvalues. Second, eigenfunctions with $k_1 + k_2 \geq 2$ do not describe any additional directions intrinsic to the strip geometry; we will refer to these as “repeated eigendirections,” and we will refer to the eigenfunctions with $k_1 + k_2 = 1$ as “unique eigendirections.” Clearly, this problem of repeated eigendirections also arises for curved manifolds of more than two variables. We note that, although the eigenfunctions can only be written analytically for very special cases, it has been observed empirically that they often provide appropriate coordinates to usefully parametrize more complicated, nonlinear manifolds.

2.2. Discrete approximation of the Laplace–Beltrami operator: diffusion maps

In most applications, we are not given a description of the continuous manifold on which the data lie. Instead, we are given data *sampled* from the underlying manifold, and the parametrization of the manifold

needs to be uncovered from the data. This can be accomplished by constructing a matrix which approximates the Laplace–Beltrami operator. It was shown in [15] that, in the limit of infinite data, this discrete Laplacian matrix constructed from data converges pointwise to the continuous Laplace–Beltrami operator on the manifold. As a result, the eigenvectors of the discrete Laplacian approximate the eigenfunctions of this continuous operator.

Given observations $z(1), \dots, z(m) \in \mathcal{M}_d$, we first construct the weight matrix $\mathbf{W} \in \mathbb{R}^{m \times m}$, with

$$\mathbf{W}_{ij} = \exp\left(-\frac{\|z(i) - z(j)\|^2}{\epsilon^2}\right), \quad i, j = 1, \dots, m, \tag{3}$$

where $\|\cdot\|$ denotes the appropriate norm for the observations, and ϵ is a characteristic distance between the observations. The kernel’s scale ϵ can be chosen using several heuristics [22,28]; we often take ϵ to be the median of the pairwise distances between the data points. The underlying assumption is that, even in cases where the samples embedded in the high dimensional space lie on a highly nonlinear manifold, within ϵ -local neighborhoods, the Euclidean distance “respects” the tangent plane to the manifold and thus locally conveys meaningful relationships. We then construct the diagonal matrix $\mathbf{D} \in \mathbb{R}^{m \times m}$, with $\mathbf{D}_{ii} = \sum_j \mathbf{W}_{ij}$, and form the matrix $\tilde{\mathbf{W}} = \mathbf{D}^{-\alpha} \mathbf{W} \mathbf{D}^{-\alpha}$, where $0 \leq \alpha \leq 1$. Next, we construct the diagonal matrix $\tilde{\mathbf{D}} \in \mathbb{R}^{m \times m}$, with $\tilde{\mathbf{D}}_{ii} = \sum_j \tilde{\mathbf{W}}_{ij}$, and the matrix $\mathbf{A} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{W}}$.

If the data $z(1), \dots, z(m)$ are sampled from \mathcal{M}_d with some density q , then, for $\epsilon \rightarrow 0$ and $m \rightarrow \infty$ (with the appropriate rates), the discrete matrix converges to the following continuous limit operators with Neumann boundary conditions (as discussed in the previous section) [15]

$$\frac{1}{\epsilon^2}(\mathbf{I} - \mathbf{A})\phi \rightarrow \nabla^2 \phi - 2\nabla U \cdot \nabla \phi, \quad \alpha = 0 \tag{4}$$

$$\frac{1}{\epsilon^2}(\mathbf{I} - \mathbf{A})\phi \rightarrow \nabla^2 \phi - \nabla U \cdot \nabla \phi, \quad \alpha = 1/2 \tag{5}$$

$$\frac{1}{\epsilon^2}(\mathbf{I} - \mathbf{A})\phi \rightarrow \nabla^2 \phi, \quad \alpha = 1 \tag{6}$$

where $U = -\log q$. The different limiting operators, depending on the choice of α , imply that nonuniform sampling on the manifold may have different effects via the potential U in (4)–(6). In particular, by setting $\alpha = 1$, one can factor out the density effects in the weight matrix, and the discrete Laplacian matrix approaches the Laplace–Beltrami operator on the manifold \mathcal{M}_d . Therefore, the eigenvectors $\phi_0, \phi_1, \dots, \phi_{m-1}$ of \mathbf{A} approximate the eigenfunctions of the Laplace–Beltrami operator on \mathcal{M}_d , and the eigenvalues $\mu_0, \mu_1, \dots, \mu_{m-1}$ of \mathbf{A} are related to the eigenvalues of the continuous operator by

$$\mu_k = \exp\left(-\frac{\epsilon^2}{4} \tilde{\mu}_{k_1, k_2}\right). \tag{7}$$

As discussed previously, the eigenfunctions provide a parametrization of the manifold, such that $\mu_\tau^\top \phi_j(i)$ yields the j th embedding coordinate for $z(i)$, where $\tau \in \mathbb{N}$ is a parameter (in our examples, we will take $\tau = 0$). The parameter τ corresponds to the number of iterates of the Laplace operator, or in other words, the number of diffusion steps (for more details, see [15]). This mapping $z(i) \mapsto (\mu_1^\top \phi_1(i), \dots, \mu_{m-1}^\top \phi_{m-1}(i))$ is the standard diffusion maps embedding [14,15], where we order the eigenvectors such that $|\mu_0| \geq |\mu_1| \geq \dots \geq |\mu_{m-1}|$. Because the matrix \mathbf{A} is row-stochastic ($\sum_j \mathbf{A}_{ij} = 1$), $\mu_0 = 1$ and ϕ_0 is a trivial constant vector. The distance induced by this mapping is called the standard diffusion distance,

$$D_\tau^2(z(i), z(j)) = \sum_{k=1}^{m-1} \mu_k^{2\tau} (\phi_k(i) - \phi_k(j))^2. \tag{8}$$

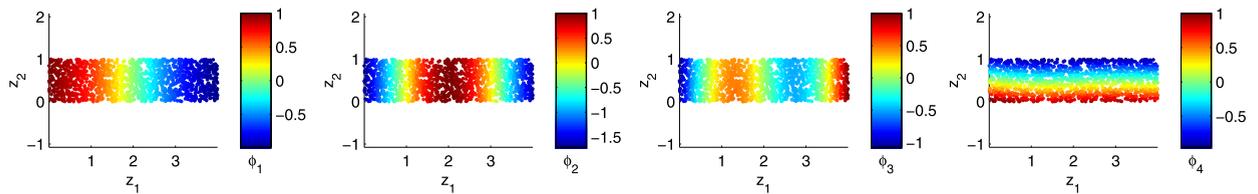


Fig. 2. Data, uniformly sampled from a two-dimensional strip, colored by the first four (non-trivial) eigenvectors from diffusion maps. Note that the first and fourth eigenvectors are one-to-one with z_1 and z_2 , respectively. However, the second and third eigenvectors are higher harmonics of the first eigenvector and do not capture any additional structure within the data set.

Figs. 1b–1c show data sampled from a strip, colored by eigenvectors of \mathbf{A} . In cases of both uniform and nonuniform sampling, the selected eigenvectors are one-to-one with z_1 and z_2 , and thus parametrize the manifold. Although we have considered a very simple example for illustrative purposes, common practice is to use these tools for the analysis of high-dimensional, nonlinear data sets.

From the previous section, we know that the eigenfunctions with $(k_1, k_2) = (1, 0)$ and $(k_1, k_2) = (0, 1)$ provide embedding coordinates for this manifold; these two eigenfunctions describe each a different direction along the data, and are not repeated. From (1), we see that sorting the eigenvectors by the magnitude of the corresponding eigenvalues implies that these two eigenvectors are guaranteed to appear before any coupled or repeated eigendirections. However, these eigenvectors are *not* guaranteed to appear as *the first two* (non-trivial) eigenvectors, as harmonics of the first eigendirection (i.e., $\cos(n\pi z_1/L_1)$ with $n > 1$) could appear before the second. We note that, in contrast to our simple illustrative example, for most data sets of interest, the coordinates for the underlying manifold are unknown and cannot easily be obtained from the coordinates of the original data, so that identifying which eigenvectors correspond to unique eigendirections is nontrivial.

3. Identifying the informative eigenvectors

Common practice is to order the eigenvectors by the magnitude of the corresponding eigenvalues, assuming that the leading eigenvectors provide a parametrization of the underlying manifold. However, as discussed in the previous section, some eigenvectors are higher harmonics of previous eigenvectors and do not describe new directions in the data set [24]. For the case where \mathcal{M}_d is a 2-dimensional strip with edge lengths $L_1 > L_2$, recall that the eigenfunctions $\tilde{\phi}_{1,0} = \cos(\pi z_1/L_1)$ and $\tilde{\phi}_{0,1} = \cos(\pi z_2/L_2)$ provide embedding coordinates for the manifold \mathcal{M}_d . However, these two eigenvectors $\tilde{\phi}_{1,0}$ and $\tilde{\phi}_{0,1}$ are not guaranteed to correspond to the two smallest (non-trivial) eigenvalues (the smallest eigenvalue will always be $\tilde{\mu}_{0,0} = 0$ and correspond to a constant eigenfunction; these eigenvalues are related to the eigenvalues μ_k of the discrete operator via (7)). In fact, if $L_1 > 2L_2$, then $\tilde{\mu}_{2,0} < \tilde{\mu}_{0,1}$, and so the second (non-trivial) eigenvector (when the eigenvectors are ordered by their corresponding eigenvalues) will be a “repeated eigendirection” of the first and still parametrize z_1 (see Fig. 2). Automatic detection of which eigenvectors are harmonics of each other is clearly useful. Utilizing only the eigenvectors ϕ_i which correspond to unique eigendirections yields the most parsimonious representation of the data. We will show how we can automatically detect these eigenvectors to obtain a meaningful representation of the data, and that using only these eigenvectors yields a reduced embedding which is equivalent to the standard diffusion maps embedding. Furthermore, when the data is uniformly sampled from the underlying manifold, the corresponding eigenvalues provide us with an estimate of the relative lengths of the data set along these unique eigendirections.

3.1. Algorithm: local linear regression

Given the eigenvectors $\phi_1, \phi_2, \dots, \phi_{m-1} \in \mathbb{R}^m$, we would like to automatically deduce which ones capture new directions in the data, and which ones are merely repeated eigendirections. This problem was addressed previously in [24] by performing successive iterations of diffusion maps, interspersed with advection along

the first eigendirection at each iteration. However, both the advection procedure and the successive eigendecompositions are expensive and soon become intractable for larger data sets. Here, we propose an alternative approach to address the problem of repeated eigendirections motivated by simple trigonometric arguments (for example, the repeated eigendirection $\cos 2x$ can be written as a function of the unique eigendirection $\cos x$). We therefore attempt to fit a function $f(\phi_1, \dots, \phi_{k-1})$ to ϕ_k ; if the resulting fit is accurate, we assume ϕ_k is a repeated eigendirection of $\phi_1, \dots, \phi_{k-1}$. We use a *local* linear function

$$\phi_k(i) \approx \alpha_k(i) + \beta_k^T(i)\Phi_{k-1}(i) \tag{9}$$

as our functional approximation, where $\Phi_{k-1}(i) = [\phi_1(i) \ \dots \ \phi_{k-1}(i)]^T$, $\alpha_k(i) \in \mathbb{R}$, and $\beta_k(i) \in \mathbb{R}^{k-1}$. The coefficients $\alpha_k(i)$ and $\beta_k(i)$ are not constant because we use a *local* linear fit in the $k - 1$ -dimensional Φ_{k-1} space, and so the coefficients change as a function of the domain.

At each point $\Phi_{k-1}(i)$, we approximate $\phi_k(i)$ by fitting a local linear function using the remaining $m - 1$ data points. We solve the following optimization problem

$$\hat{\alpha}_k(i), \hat{\beta}_k(i) = \underset{\alpha, \beta}{\operatorname{argmin}} \sum_{j \neq i} K(\Phi_{k-1}(i), \Phi_{k-1}(j)) (\phi_k(j) - (\alpha + \beta^T \Phi_{k-1}(j)))^2, \tag{10}$$

where K is a kernel weighting function. We use a Gaussian kernel,

$$K(\Phi_{k-1}(i), \Phi_{k-1}(j)) = \exp\left(-\frac{\|\Phi_{k-1}(i) - \Phi_{k-1}(j)\|^2}{\epsilon_{reg}^2}\right), \tag{11}$$

where ϵ_{reg} is the kernel scale for the regression algorithm. We typically take $\epsilon_{reg} = M/3$, where M is the median of the pairwise distances between $\Phi_{k-1}(i)$, as we empirically found this choice to yield good results. We then define the normalized leave-one-out cross-validation error for this local linear fit as

$$r_k = \sqrt{\frac{\sum_{i=1}^n (\phi_k(i) - (\hat{\alpha}_k(i) + \hat{\beta}_k(i)^T \Phi_{k-1}(i)))^2}{\sum_{i=1}^n (\phi_k(i))^2}}. \tag{12}$$

Note that a small value of r_k implies that ϕ_k can be accurately approximated from $\phi_1, \dots, \phi_{k-1}$. We reason that a small value of r_k implies that ϕ_k is a harmonic of previous modes, i.e., a repeated eigendirection, and conversely, a large value of r_k indicates that ϕ_k parametrizes a new unique eigendirection in the data. We choose to set $r_1 = 1$. The error in (12) can easily be computed (see Section 5.4 of [25]; code is also available at ronen.eew.technion.ac.il/files/2015/01/code.zip).

We propose using *only* those eigenvectors for which r_k is large to embed the data, as this will yield a more parsimonious representation. However, it is important to note that such an embedding also preserves certain features of the standard diffusion map embedding. Let $I = \{i_1, i_2, \dots, i_d\}$ denote the indices of the identified unique eigendirections (i.e., r_{i_1}, \dots, r_{i_d} are large). Under the assumption that $\phi_j = f(\phi_{i_1}, \dots, \phi_{i_d})$ for $j \notin I$, where f is Lipschitz continuous with Lipschitz constant K , one can show that

$$\frac{1}{1 + K^2 \sum_{k \notin I} \mu_k^{2\tau}} D_\tau^2(z(i), z(j)) \leq \sum_{k \in I} \mu_k^{2\tau} (\phi_k(i) - \phi_k(j))^2 \leq D_\tau^2(z(i), z(j)). \tag{13}$$

Therefore (for finite data, or provided the eigenvalues μ_k decay sufficiently fast in the limit of infinite data), the distance induced by the d eigenvectors $\phi_{i_1}, \dots, \phi_{i_d}$ identified as parametrizing unique eigendirections is equivalent to the standard diffusion distance. We will refer to this distance

$$\tilde{D}_\tau^2(z(i), z(j)) = \sum_{k \in I} \mu_k^{2\tau} (\phi_k(i) - \phi_k(j))^2 \tag{14}$$

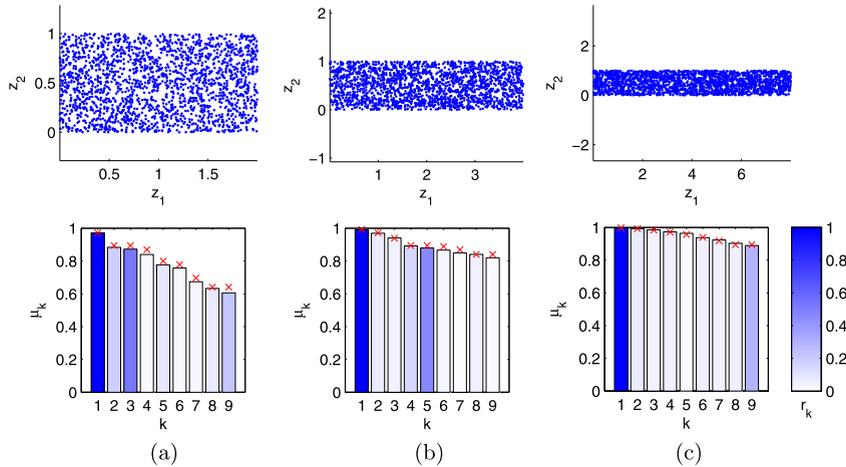


Fig. 3. Data sets (top) and eigenvalue spectra from diffusion maps analysis (bottom) for strips with (a) $L_1 = 2, L_2 = 1$, (b) $L_1 = 4, L_2 = 1$, (c) $L_1 = 8, L_2 = 1$. The empirical eigenvalues are plotted as blue bars, and the analytical eigenvalues are plotted as red x's. The bars are colored by r_k , as defined in (12). From the eigenvalues which are identified as parametrizing unique eigendirections (indicated by the darker blue bars), the estimated length ratio from (15) is (a) 2.2, (b) 4.1, (c) 8.7. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

as the *reduced* diffusion distance, and the embedding obtained from the corresponding eigenvectors as the *reduced* diffusion maps embedding.

3.2. Estimating the “relative lengths”

For a two-dimensional strip, once the eigenvectors which parametrize unique eigendirections have been determined, the corresponding eigenvalues can be used to calculate the relative lengths along these directions. For more general manifolds (at least for the case of uniform sampling), we postulate that the eigenvalues can still provide a measure of the relative significance of the unique eigendirections. Let $\mu_{i_1}, \mu_{i_2}, \dots$ denote the eigenvalues corresponding to eigenvectors which parametrize unique eigendirections (i.e., those eigenvectors where r_{i_j} is large). For the strip example, from (1) and (7), we propose to approximate the relative lengths L_j along the manifold by

$$L_j \propto \frac{1}{\sqrt{-\log \mu_{i_j}}}. \tag{15}$$

These lengths can then be used to evaluate how many components are required to retain most of the information within the data set, similar to the eigenvalues in principal component analysis.

3.3. Illustrative examples

We demonstrate our proposed approach on two synthetic data sets. The first data set is the two-dimensional strip discussed previously, where the eigenvalues and eigenvectors are known analytically. The second data set is a nonlinear manifold which demonstrates the flexibility of our approach.

3.3.1. Strip

We consider three different two-dimensional strip data sets. Each data set contains $m = 2000$ data points uniformly sampled from the strip. Fig. 3 shows the data sets and the diffusion maps eigenspectra. The eigenvalues are colored by the leave-one-out cross-validation error r_k as defined in (12); a small value of r_k indicates that the corresponding eigenvector is a repeated eigendirection, while a large value of r_k indicates that the corresponding eigenvector describes a new direction in the data.

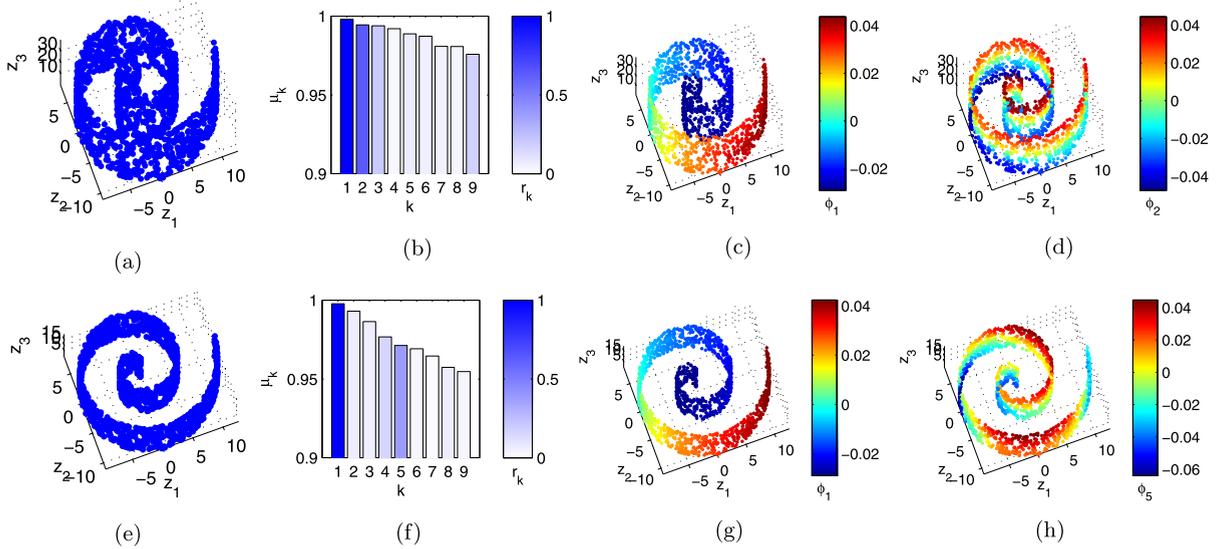


Fig. 4. Swiss roll example. (a) Data set 1; $h = 40$. (b) Eigenvalue spectrum from the diffusion maps analysis of data set 1. (c) Data set 1, colored by the first diffusion maps eigenvector. (d) Data set 1, colored by the second diffusion maps eigenvector. (e) Data set 2; $h = 20$. (f) Eigenvalue spectrum from the diffusion maps analysis of data set 2. (g) Data set 2, colored by the first diffusion maps eigenvector. (h) Data set 2, colored by the fifth diffusion maps eigenvector.

The eigenvalues are consistent with the known analytic eigenvalues of the Laplacian (see (1) and (7)), shown in red. Furthermore, the two unique eigendirections can easily be identified, since their corresponding regression errors r_k are large. As expected, the gap between the two meaningful eigenvalues increases as the strip becomes longer. Using (15), we can accurately estimate the relative lengths of the two unique eigendirections in each data set.

3.3.2. Swiss roll

Our second illustrative example consists of two different Swiss roll data sets. The data are sampled according to

$$(z_1, z_2, z_3) = (\theta \cos \theta, \theta \sin \theta, ht) \tag{16}$$

where θ is sampled such that z_1, z_2 are uniformly sampled along the arclength of the spiral, and t is uniformly sampled from $[0, 1]$. Note that, in this example, z_1, z_2, z_3 are the original data coordinates; however, the data lie on a two-dimensional manifold parametrized by θ and t . The height of the first Swiss roll is $h = 40$, while the height of the second is $h = 20$. Each data set consists of $m = 1500$ points, shown in Figs. 4a and 4e.

Figs. 4b and 4f show the eigenvalue spectra from the analysis of the two data sets. Similar to Fig. 3, the bars are colored by the leave-one-out cross-validation error, where a small value of r_k indicates that the corresponding eigenvector is a repeated eigendirection. From these plots, one can conclude that the first two eigenvectors ϕ_1 and ϕ_2 parametrize the first data set, while ϕ_1 and ϕ_5 parametrize the second. Figs. 4c, 4d, 4g, and 4h show the two data sets, colored by the two eigenvectors identified as parametrizing the unique eigendirections. As expected, these eigenvectors are one-to-one with the arclength along the spiral, and the height of the Swiss roll.

4. Chemotaxis: a case study

Our main motivation for identifying the unique eigendirections is to facilitate the analysis of complex data sets for which the underlying dimensionality and structure is not readily apparent. In data from complex

dynamical systems, noisy microscopic behavior often gives rise to coherent macroscopic dynamics. In general, this mapping from microscale to macroscale is not always obvious. We will show how we can use our proposed methodology to extract a parametrization of data collected at the microscale which is consistent with a certain macroscopic behavior, without any *a priori* knowledge of the appropriate microscopic or macroscopic model. Furthermore, determining the *true* dimensionality of such a parametrization by identifying the unique eigendirections reveals the requisite dimensionality of a reduced model which captures the relevant macroscopic behavior. Knowing appropriate macroscopic variables and the true dimensionality can help inform modeling efforts and aid in positing useful macroscale models.

Our model problem describes the process of cellular chemotaxis [29], where biological cells exhibit coherent macroscopic dynamics regulated by extracellular sensed signals in order to accomplish tasks such as finding food or navigating away from toxins. Several microscopic models have been proposed to describe chemotactic dynamics [26,30]. We will analyze one such model described by a one-dimensional velocity jump process [29]. This specific example has an analytic macroscopic description in which the dynamics of a large ensemble of cells depend on the value of a single system parameter. This description serves as a “ground truth” and will allow us to validate our results which are obtained in an unsupervised manner.

Thus far, we only considered synthetic data sets for which the Euclidean distance between data points served as an informative metric. Here, we will show that our approach, *when utilizing the appropriate statistical observers and affinity metric between pairs of observations*, uncovers a parametrization of the microscopic data which is consistent with the known macroscopic model. Furthermore, we will show that changes in the “mode” of dynamical behavior resulting from changes in the system parameters can be automatically detected.

4.1. Problem description

The microscopic model consists of a collection of N cells (for our simulations, we take $N = 1000$) whose states are defined by their positions and velocities on a line, and the dynamics of each cell are governed by a stochastic process. Let $x_i(t)$ and $v_i(t)$ denote the position and velocity, respectively, of cell i at time t . The velocity of each cell is $\pm s$, where s is a (fixed) speed. We initialize the cells such that

$$\begin{aligned} x_i(0) &= 0 \\ \mathbb{P}\{v_i(0) = +s\} &= p \end{aligned} \quad (17)$$

where $0 < p < 1$ is the probability of a cell initially moving to the right. The velocity of each cell randomly switches between $\pm s$ following an (independent) Poisson process with rate λ (this switching is physically controlled by extracellular signals). For our specific simulations, we set $s^2/\lambda = 1$, which is consistent with the analysis presented in [26]. We note that we have chosen a very specific one-parameter family of initial conditions which lead to simple dynamics and allow us to illustrate our main points; however, our analysis is not restricted to such specific cases and more complex sets of initial conditions could be used. Each data set consists of 10 stochastic simulations, with initial conditions uniformly chosen such that $0.1 \leq p \leq 0.9$. We allow each simulation to evolve for t_{max} time units, and record the states of the N cells every Δt time units; we use data with $t > 0$ for analysis. We will show that the parameter $t_{obs} \equiv t_{max}/N$, which determines the time scale of observation, is of critical importance when we address the limiting asymptotic behavior.

4.2. Observers and metrics

In general, manifold learning techniques have two essential components. One is the appropriate *observers* of the system state. These observers should be informative as to the state of the system, as well as insensitive

to noise in the system. The second is a *distance metric* between the observations that captures a notion of locality: observations which we perceive to be similar should yield a small value for this distance.

For the chemotaxis example, we use histograms of the cell positions as observers. Histograms are invariant to the indexing of the cells, while retaining information about the relative spatial locations of the cells. Histograms are also robust to noise [20]. Instead of the standard Euclidean distance, we use the earth mover’s distance (EMD) [18] as the metric between pairs of histograms. Conceptually, the EMD measures how much “work” it takes to transform one probability density into another. It therefore not only considers where the densities are inconsistent, but also how far apart the inconsistencies are from each other. Although the brute-force computation of the EMD is computationally expensive, there has been a plethora of work in developing efficient algorithms for its approximation [31–33]. For the specific case of spatially one-dimensional data, the EMD is equivalent to the L_1 -norm between the cumulative distribution functions of the data [34], which can be approximated from histograms as

$$\|z(i) - z(j)\|_{EMD} \approx \sum_{l=1}^n \left| \sum_{k=1}^l z_k(i) - \sum_{k=1}^l z_k(j) \right|; \quad (18)$$

the histograms $z(i), z(j)$ are defined on n equally-spaced bins in \mathbb{R} (for our simulations, we will take $n = 32$), and z_k denotes the k th bin. Here, we choose to only use the positions of the cells in the distance computation; utilizing both the positions and velocities produces qualitatively similar results (not shown). We note that although histograms are not essential to the theory of EMD, they make the computation practical.

4.3. Results

4.3.1. Identifying the unique eigendirections

Fig. 5 shows the results of analyzing two sets of chemotaxis simulations. One set of simulations (Figs. 5a–5c) is for a relatively small value of λ , while the other set (Figs. 5d–5f) is for a larger value of λ . For both sets of simulations, the macroscopic variable p , which controls the initial distribution of cell velocities, and t , the time, are well-correlated with the eigenvectors ϕ_1 and ϕ_2 . In particular, p characterizes the fraction of initially right-moving cells (alternatively, the initial *flux* of the right-moving cells). The dominant coordinate ϕ_1 is correlated with p for the small λ case (Fig. 5a), and correlated with t for the large λ case (Fig. 5e), indicating that the relative importance of p and t changes between the two simulations.

As illustrated in the synthetic data sets we discussed previously, the two leading eigenvectors are not guaranteed to correspond to unique eigendirections. Fig. 6 shows the results of analyzing three different simulations which span a larger range of λ values compared to Fig. 5. We see from the eigenspectra and the cross-validation errors r_k that the second unique eigendirection moves down in the spectrum as λ increases. The two-dimensional reduced diffusion maps embeddings obtained from eigenvectors which parametrize unique eigendirections capture the macroscopic variables p and t (we only show the correspondence of the two identified unique eigendirections with p in the middle row of Fig. 6; the correspondence between these eigendirections and t is similarly consistent). In contrast, using the two leading eigenvectors ϕ_1 and ϕ_2 produces uninformative embeddings for large values of λ , as ϕ_2 is a repeated eigendirection which does not parametrize any new directions in the data. Furthermore, although the eigenvalue spectra do not exhibit any large spectral gaps, the gap between the eigenvalues which correspond to new directions in the data (indicated by the dark blue) increases with λ , indicating that the ratio of relative extents along different directions on the underlying manifold becomes smaller, and the corresponding eigenvectors provide us with an informative two-dimensional parametrization of the data. It now becomes clear that looking at the eigenvectors *modulo* repeated eigendirections is essential for extracting an informative parametrization of the data, as well as characterizing the effective dimensionality of the data.

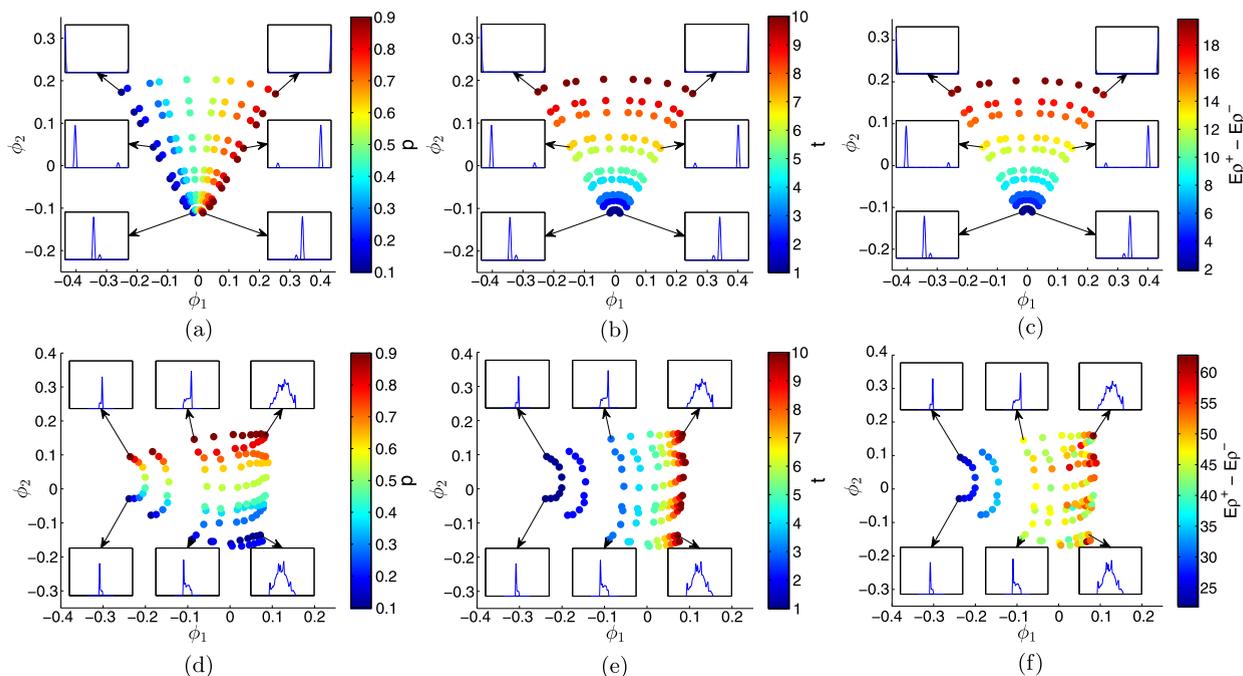


Fig. 5. Embeddings using the first two diffusion maps eigenvectors computed from simulation data of the velocity jump process with $t_{max} = 10$, $\Delta t = 1$, and (a–c) $\lambda = 1$, $s = 1$, (d–f) $\lambda = 400$, $s = 20$. The distances used in the diffusion maps kernel are the earth mover’s distances between the histograms of cell positions. The data are colored by (a, d) p , the initial probability of a cell moving to the right, (b, e) t , time, and (c, f) $E\rho^+ - E\rho^-$, the difference between the average position of the left- and right-moving cells. Representative histograms of the cell positions are shown for selected data points.

4.3.2. Analytic macroscopic description

This particular model system can be usefully described through a known analytic macroscopic equation that governs the overall system behavior. For a large collection of cells ($N \rightarrow \infty$), the system can be described using the evolution of the cell density. Let $\rho(x, t)$ denote this evolving cell density field, and let $\rho^-(x, t)$ and $\rho^+(x, t)$ denote the densities of the left-moving and of the right-moving cells, respectively. It can be shown that, as $N \rightarrow \infty$, the densities satisfy the following set of coupled partial differential equations (PDEs) [29]:

$$\begin{aligned} \frac{\partial \rho^+}{\partial t} + s \frac{\partial \rho^+}{\partial x} &= -\lambda \rho^+ + \lambda \rho^- \\ \frac{\partial \rho^-}{\partial t} - s \frac{\partial \rho^-}{\partial x} &= \lambda \rho^+ - \lambda \rho^-. \end{aligned} \tag{19}$$

Alternatively, (19) can be rewritten as a single, second-order PDE (the telegrapher equation):

$$\frac{\partial^2 \rho}{\partial t^2} + 2\lambda \frac{\partial \rho}{\partial t} = s^2 \frac{\partial^2 \rho}{\partial x^2}. \tag{20}$$

From (20), we see that, for fixed values of λ and s , the macroscopic state of the system (the probability density of the cells) is indeed a function of the initial condition (parametrized by p) and the time t . This implies that for fixed λ and s , the microscopic data in a high-dimensional ambient space (e.g., the positions of all N cells) should lie on a two-dimensional manifold parametrized by p (which controls the initial ratio of ρ^+ and ρ^-) and t (which describes the evolution); uncovering the low-dimensional structure of the microscopic data helps reveal this manifold, leading to a subsequent physically meaningful parametrization of it. This is consistent with the results presented in Figs. 5 and 6, where the two-dimensional embeddings obtained from microscopic simulation data are one-to-one with p and t .

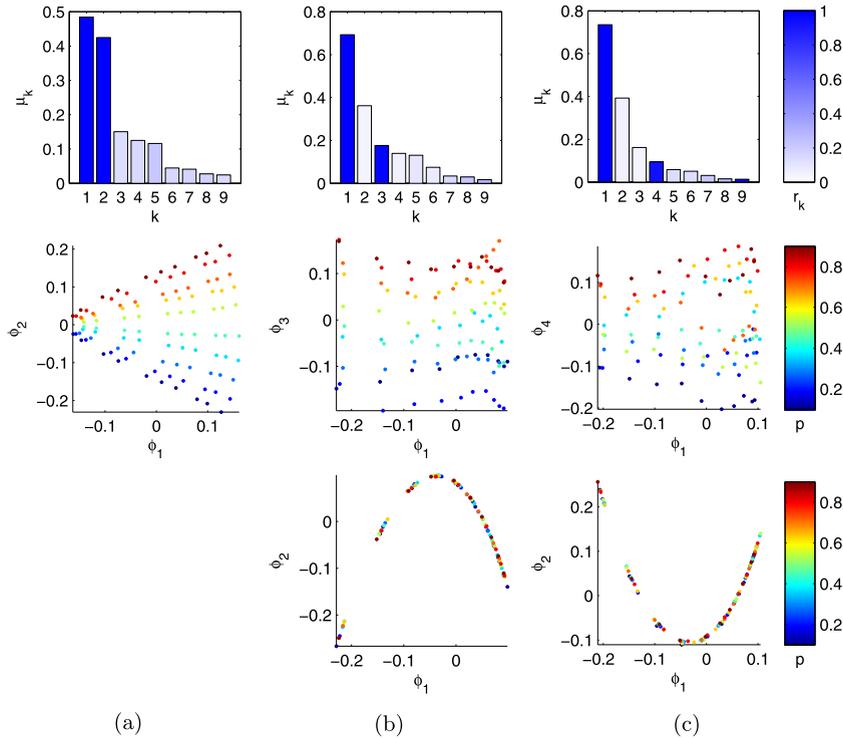


Fig. 6. Analysis of three different chemotaxis simulation data sets. (a) $\lambda = 100$, $s = 10$. (b) $\lambda = 1600$, $s = 40$. (c) $\lambda = 6400$, $s = 80$. We set $t_{max} = 10$ and $\Delta t = 1$, and the distances used in the diffusion maps kernel are the earth mover's distances between the histograms of cell positions. For each data set, the eigenvalue spectrum, colored by r_k , is shown in the top row. From the spectra, we can see that the first two components are informative for (a), that components 1 and 3 are informative for (b), and that components 1 and 4 are informative for (c). The corresponding reduced diffusion maps embeddings are shown in the middle row. For comparison, the standard diffusion maps embeddings using the first two components are also shown in the bottom row in (b) and (c).

Guided by the known theory of the model macroscopic PDE system, we consider two asymptotic regimes for the simulation. When, on the other hand, $\lambda \rightarrow 0$, the right-hand side of (19) tends to 0, and (19) becomes two uncoupled wave equations,

$$\begin{aligned} \frac{\partial \rho^+}{\partial t} + s \frac{\partial \rho^+}{\partial x} &= 0 \\ \frac{\partial \rho^-}{\partial t} - s \frac{\partial \rho^-}{\partial x} &= 0. \end{aligned} \tag{21}$$

When, on the other hand, $\lambda \rightarrow \infty$ with s^2/λ fixed, (20) approaches the diffusion equation,

$$2 \frac{\partial \rho}{\partial t} = D \frac{\partial^2 \rho}{\partial x^2}, \tag{22}$$

where $D = s^2/\lambda$. The above analysis shows that the initial distribution of velocities of the cells (determined by p in the microscopic simulations) plays a very different role depending on the value of λ . When $\lambda \rightarrow 0$, the dynamics are described by the telegrapher equation, and the effect of the initial velocity distribution persists throughout the trajectory. When $\lambda \rightarrow \infty$, the dynamics are described by a single diffusion equation, and the initial conditions for the velocity are not especially important, since the velocity distribution quickly equilibrates and we see purely diffusive behavior.

The relative importance of p and t from the analytic description are consistent with the results in Fig. 5, where p is correlated with the dominant diffusion map coordinate when λ is small, and becomes correlated with the subdominant coordinate when λ is large. Furthermore, in the small λ regime (wave equation),

shown in Figs. 5a–5c, the points corresponding to small times are more tightly clustered than the points corresponding to large times. This is in agreement with the macroscopic model: at small times, the cells are more condensed around $x = 0$, and it is more difficult to distinguish the cells moving to the left from the cells moving to the right. On the other hand, at large times, once the cells move away from the origin, this separation between left-moving and right-moving cells is clear. For the large λ case (diffusion equation), shown in Figs. 5d–5f, we observe that at small times (before the cell velocities have sufficient opportunity to switch and equilibrate), the initial velocity distribution p can be well perceived in the embedding in Fig. 5d, as the initial imbalance in left-right velocities affects the initial displacements. On the other hand, for longer times, we observe that the initial distribution p cannot be well perceived in the embedding in Fig. 5d, as the velocities have equilibrated and their initial distribution is practically forgotten. Overall, in both cases, we obtain, in an unsupervised data-driven manner, a useful low-dimensional embedding of the data. Correlating this embedding with potential candidate macroscopic observables (such as time, densities, fluxes, and initial condition statistics) helps elucidate an accurate picture of the macroscopic variables that govern the system dynamics.

The analytic macroscopic model suggests that when λ is small, the system dynamics can be described by the two densities ρ^+ and ρ^- , and when λ is large, the dynamics can be described by a single density $\rho = \rho^+ + \rho^-$. Figs. 5c and 5f show the data, colored by the difference in the average position of the left- and right-moving cells. Clearly, for small λ , both ρ^+ and ρ^- are required to describe the dynamics, and the difference between the two densities is a useful observable in parametrizing the data manifold (and thus a good macroscopic descriptor, a good “candidate macroscopic variable”). However, for large λ , the two densities rapidly equilibrate, and the difference between the two densities is *not* a useful observable in parametrizing the data manifold.

This analytic macroscopic description allows us to emphasize the importance of using an appropriate distance metric. We also analyzed the two sets of simulation data from Fig. 5 using the standard Euclidean distance between the histograms to compute the distances in (3). We empirically found that there is no appreciable correlation between the embedding coordinates and what we know to be meaningful macroscopic variables, p and t (the correlations between the embedding coordinates and the governing macroscopic variables are all less than 60%, with the correlations for the small λ case being less than 20%) with this similarity measure. In contrast, the correlations between the embedding coordinates and the macroscopic variables when using EMD in the diffusion maps calculation are all greater than 80%. Clearly, using a metric which meaningfully describes the distances between observations is essential for obtaining an informative parametrization of data. We note again that the metrics we used in this work only employed observations of cell positions (and not of cell instantaneous velocities); if the instantaneous velocities are readily measurable, they could be incorporated in a useful and possibly even more informative metric for data analysis.

4.3.3. Detecting changes in dimensionality

Detecting changes in effective dimensionality is essential for the analysis and modeling of (simulated or experimentally observed) dynamical systems. The approximate dimensionality of a data set can be estimated by examining the two eigenvalues $\mu_{i_1} \geq \mu_{i_2}$ corresponding to unique eigendirections. According to (15), we plot

$$\sqrt{\frac{\log \mu_{i_1}}{\log \mu_{i_2}}}; \quad (23)$$

when this ratio becomes small, the data are *effectively* one-dimensional, as the extent of data along the second direction is very small compared with the first.

Fig. 7 shows the eigenvalue ratio in (23) as a function of λ and the time scale of observation $t_{obs} = t_{max}/N$. From this eigenvalue ratio, we can easily detect changes in the estimated dimensionality as we

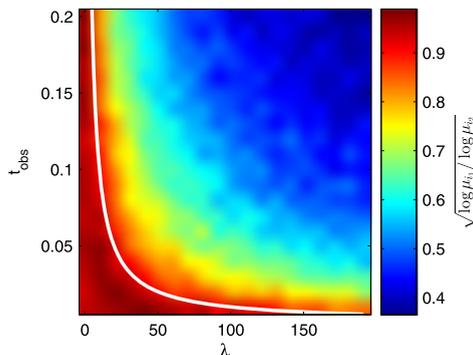


Fig. 7. Detecting the change in dimensionality. The appropriate ratio the two eigenvalues (see (23)) which correspond to unique eigendirections is plotted as a function of t_{obs} and λ . Each data point is the average eigenvalue ratio over three replicate data sets, and in each simulation, $\Delta t = t_{max}/10$. The distances used in the diffusion maps kernel are the earth mover's distances between the histograms of cell positions. The curve $t_{obs} = 1/\lambda$ is shown in white.

vary the relevant parameters. For small λ and/or short t_{obs} , the data are effectively two-dimensional, as both p and t are important to the observed dynamics. However, when λ and/or t_{obs} is large, the system rapidly equilibrates and the data are effectively one-dimensional, parametrized by time. For this specific example, these changes are consistent with the analytically-predicted shift in the dynamical behavior from the telegrapher equation to the diffusion equation. From (20), we can see that t_{obs} , the time scale of observation, should be larger than the time scale $1/\lambda$ for the diffusion equation limit in (22) to hold; we therefore expect a transition in the dynamical behavior at

$$t_{obs} \approx \frac{1}{\lambda}. \quad (24)$$

This curve is also plotted in Fig. 7 and is consistent with the transition predicted by the diffusion maps eigenvalues.

5. Conclusions

This paper addresses repeated eigendirections, a critical issue in applying diffusion maps to the analysis of complex dynamic data. We show that, using local linear regression, we can automate the detection of eigenvectors corresponding to unique directions in the geometry of the data. From this detection, we can then obtain a *reduced diffusion maps embedding* which is both parsimonious and induces a metric which is equivalent to the standard diffusion distance. We showed that this algorithm enables us to more fruitfully analyze data from a complex dynamical system, enabling the extraction of good reduced coordinates and the detection of changes in the dimensionality of the macroscopic dynamics. We are confident that the proposed methodology will be helpful in the analysis of high-dimensional data sets for which the dimensionality of the underlying manifold is unknown.

Acknowledgments

C.J.D. was supported by the Department of Energy Computational Science Graduate Fellowship (CSGF), grant number DE-FG02-97ER25308, and the National Science Foundation Graduate Research Fellowship, grant number DGE 1148900. R.T. was supported by the European Union's Seventh Framework Programme (FP7) under Marie Curie Grant 630657 and by the Horev Fellowship. I.G.K. was supported by the National Science Foundation (CS&E program), grant number 1310173, and by the US AFOSR (Dr. F. Darema, Grant FA9550-12-1-0332). R.R.C. was supported by the National Science Foundation (CS&E program), grant number 1309858.

References

- [1] S. Gepshtein, Y. Keller, Image completion by diffusion maps and spectral relaxation, *IEEE Trans. Image Process.* 22 (8) (2013) 2983–2994.
- [2] Á. Fernández, N. Rabin, R.R. Coifman, J. Eckstein, Diffusion methods for aligning medical datasets: location prediction in CT scan images, *Med. Image Anal.* 18 (2) (2014) 425–432.
- [3] A. Singer, Z. Zhao, Y. Shkolnisky, R. Hadani, Viewing angle classification of cryo-electron microscopy images using eigenvectors, *SIAM J. Imaging Sci.* 4 (2) (2011) 723–759.
- [4] L. Yuan, C. Pan, S. Ji, M. McCutchan, Z.-H. Zhou, S.J. Newfeld, S. Kumar, J. Ye, Automated annotation of developmental stages of *Drosophila* embryos in images containing spatial patterns of expression, *Bioinformatics* 30 (2) (2014) 266–273.
- [5] W. Zhao, R. Chellappa, P.J. Phillips, A. Rosenfeld, Face recognition: a literature survey, *ACM Comput. Surv.* 35 (4) (2003) 399–458.
- [6] C. Trapnell, D. Cacchiarelli, J. Grimsby, P. Pokharel, S. Li, M. Morse, N.J. Lennon, K.J. Livak, T.S. Mikkelsen, J.L. Rinn, The dynamics and regulators of cell fate decisions are revealed by pseudotemporal ordering of single cells, *Nat. Biotechnol.* 32 (4) (2014) 381–386.
- [7] I. Kemelmacher-Shlizerman, E. Shechtman, R. Garg, S.M. Seitz, Exploring photobios, *ACM Trans. Graph.* 30 (4) (2011) 61.
- [8] L. Sifre, S. Mallat, E.N.S. Di, Rotation, scaling and deformation invariant scattering for texture discrimination, in: *Proc. CVPR*, 2013.
- [9] R. Talmon, R.R. Coifman, Intrinsic modeling of stochastic dynamical systems using empirical geometry, *Appl. Comput. Harmon. Anal.* 35 (1) (2014) 138–160.
- [10] T. Berry, J.R. Cressman, Z. Greguric-Ferencek, T. Sauer, Time-scale separation from diffusion-mapped delay coordinates, *SIAM J. Appl. Dyn. Syst.* 12 (2) (2013) 618–649.
- [11] A. Singer, R. Erban, I.G. Kevrekidis, R.R. Coifman, Detecting intrinsic slow variables in stochastic dynamical systems by anisotropic diffusion maps, *Proc. Natl. Acad. Sci. USA* 106 (38) (2009) 16090–16095.
- [12] A.L. Ferguson, A.Z. Panagiotopoulos, P.G. Debenedetti, I.G. Kevrekidis, Systematic determination of order parameters for chain dynamics using diffusion maps, *Proc. Natl. Acad. Sci. USA* 107 (31) (2010) 13597–13602.
- [13] M. Belkin, P. Niyogi, Laplacian eigenmaps for dimensionality reduction and data representation, *Neural Comput.* 15 (2003) 1373–1396.
- [14] R.R. Coifman, S. Lafon, A.B. Lee, M. Maggioni, B. Nadler, F. Warner, S.W. Zucker, Geometric diffusions as a tool for harmonic analysis and structure definition of data: diffusion maps, *Proc. Natl. Acad. Sci. USA* 102 (21) (2005) 7426–7431.
- [15] R.R. Coifman, S. Lafon, Geometric harmonics: a novel tool for multiscale out-of-sample extension of empirical functions, *Appl. Comput. Harmon. Anal.* 21 (1) (2006) 31–52.
- [16] A. Singer, R.R. Coifman, Non-linear independent component analysis with diffusion maps, *Appl. Comput. Harmon. Anal.* 25 (2) (2008) 226–239.
- [17] B. Nadler, S. Lafon, R.R. Coifman, I.G. Kevrekidis, Diffusion maps, spectral clustering and reaction coordinates of dynamical systems, *Appl. Comput. Harmon. Anal.* 21 (1) (2006) 113–127.
- [18] Y. Rubner, C. Tomasi, L.J. Guibas, The earth mover’s distance as a metric for image retrieval, *Int. J. Comput. Vis.* 40 (2) (2000) 99–121.
- [19] S. Mallat, Group invariant scattering, *Comm. Pure Appl. Math.* 65 (10) (2012) 1331–1398.
- [20] R. Talmon, R.R. Coifman, Empirical intrinsic geometry for nonlinear modeling and time series filtering, *Proc. Natl. Acad. Sci. USA* 110 (31) (2013) 12535–12540.
- [21] Z. Zhao, A. Singer, Rotationally invariant image representation for viewing direction classification in cryo-EM, *J. Struct. Biol.* 186 (1) (2014) 153–166.
- [22] M.A. Rohrdanz, W. Zheng, M. Maggioni, C. Clementi, Determination of reaction coordinates via locally scaled diffusion map, *J. Chem. Phys.* 134 (2011) 124116.
- [23] R. Talmon, S. Mallat, H. Zaveri, R.R. Coifman, Manifold learning for latent variable inference in dynamical systems, *IEEE Trans. Signal Process.* 63 (15) (2015) 3843–3856.
- [24] S. Gerber, T. Tasdizen, R. Whitaker, Robust non-linear dimensionality reduction using successive 1-dimensional Laplacian eigenmaps, in: *Proceedings of the 24th International Conference on Machine Learning*, ACM, 2007, pp. 281–288.
- [25] L. Wasserman, *All of Nonparametric Statistics*, Springer Science & Business, Media, 2006.
- [26] H.G. Othmer, S.R. Dunbar, W. Alt, Models of dispersal in biological systems, *J. Math. Biol.* 26 (3) (1988) 263–298.
- [27] P.W. Jones, M. Maggioni, R. Schul, Manifold parametrizations by eigenfunctions of the Laplacian and heat kernels, *Proc. Natl. Acad. Sci. USA* 105 (6) (2008) 1803–1808.
- [28] R.R. Coifman, Y. Shkolnisky, F.J. Sigworth, A. Singer, Graph Laplacian tomography from unknown random projections, *IEEE Trans. Image Process.* 17 (10) (2008) 1891–1899.
- [29] H.G. Othmer, T. Hillen, The diffusion limit of transport equations derived from velocity-jump processes, *SIAM J. Appl. Math.* 61 (3) (2000) 751–775.
- [30] E.A. Codling, M.J. Plank, S. Benhamou, Random walk models in biology, *J. R. Soc. Interface* 5 (25) (2008) 813–834.
- [31] O. Pele, M. Werman, A linear time histogram metric for improved sift matching, in: *ECCV*, 2008.
- [32] O. Pele, M. Werman, Fast and robust earth mover’s distances, in: *ICCV*, 2009.
- [33] W. Leeb, R. Coifman, Lipschitz norms, mixed Lipschitz norms and their duals on spaces with semigroups, with applications to Earth mover’s distance, tech. rep., Yale University, 2014.
- [34] Y. Rubner, C. Tomasi, *Perceptual Metrics for Image Database Navigation*, vol. 1, Springer, 2000.