

ADM-CLE APPROACH FOR DETECTING SLOW VARIABLES IN CONTINUOUS TIME MARKOV CHAINS AND DYNAMIC DATA*

MIHAI CUCURINGU[†] AND RADEK ERBAN[‡]

Abstract. A method for detecting intrinsic slow variables in stochastic chemical reaction networks is developed and analyzed. It combines anisotropic diffusion maps (ADMs) with approximations based on the chemical Langevin equation (CLE). The resulting approach, called ADM-CLE, has the potential of being more efficient than the ADM method for a large class of chemical reaction systems, because it replaces the computationally most expensive step of ADM (running local short bursts of simulations) by using an approximation based on the CLE. The ADM-CLE approach can be used to estimate the stationary distribution of the detected slow variable, without any a priori knowledge of it. If the conditional distribution of the fast variables can be obtained analytically, then the resulting ADM-CLE approach does not make any use of Monte Carlo simulations to estimate the distributions of both slow and fast variables.

Key words. diffusion maps, stochastic chemical reaction networks, slow variables, stationary distributions

AMS subject classifications. 65C20, 65C40

DOI. 10.1137/15M1017120

1. Introduction. The time evolution of a complex chemical reaction network often occurs at different time scales, and the observer is interested in tracking the evolution of the slowly evolving quantities (i.e., of the so-called *slow variables*), as opposed to recording each and every single reaction that takes place in the system. Whenever a separation of scales exists, one has to simulate a large number of reactions in the system in order to capture the evolution of the slowly evolving variables. With this observation in mind, it becomes crucial to be able to detect and parametrize the underlying slow manifold corresponding to the slow variables intrinsic to the system. In this paper, we introduce an unsupervised method of discovering the underlying hidden slow variables in chemical reaction networks, and of their stationary distributions, using the anisotropic diffusion map (ADM) framework [41].

The ADM is a special class of diffusion maps which has gained tremendous popularity in machine learning and statistical analysis, as a robust nonlinear dimensionality reduction technique, in recent years [37, 2, 10, 6]. Diffusion maps have been successfully used as a manifold learning tool, where it is assumed that the high-dimensional data lies on a lower-dimensional manifold, and one tries to capture the underlying

*Submitted to the journal's Computational Methods in Science and Engineering section April 16, 2015; accepted for publication (in revised form) October 14, 2016; published electronically February 22, 2017.

<http://www.siam.org/journals/sisc/39-1/M101712.html>

Funding: This work was supported by the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement 239870. It is based on work supported in part by award KUK-C1-013-04, made by King Abdullah University of Science and Technology (KAUST). The first author's research was supported by Amit Singer of PACM, by AFOSR MURI grant FA9550-10-1-0569, by award R01GM090200 from the NIGMS, and by award FA9550-09-1-0551 from AFOSR. The second author's research was supported by a University Research Fellowship from the Royal Society, by a Nicholas Kurti Junior Fellowship from Brasenose College, University of Oxford, and by a Philip Leverhulme Prize from the Leverhulme Trust. This prize money was used to support research visits of Mihai Cucuringu in Oxford.

[†]Department of Mathematics, UCLA, Los Angeles, CA 90095-1555 (mihai@math.ucla.edu).

[‡]Mathematical Institute, University of Oxford, Oxford, OX2 6GG, United Kingdom (erban@maths.ox.ac.uk).

geometric structure of the data, a setup where the traditional linear dimensionality reduction techniques (such as principal component analysis) have been shown to fail. In the diffusion maps setup, one constructs or is given a sparse weighted connected graph (usually in the form of a weighted k -nearest-neighbor graph, with each node connected only to its k nearest or most similar neighbors) and uses it to build the associated *combinatorial Laplacian* $\tilde{L} = D - W$, where W denotes the matrix of weights and D denotes a diagonal matrix with D_{ii} equal to the sum of all weights of the node i . Next, one considers the generalized eigenvalue problem $\tilde{L}x = \lambda Dx$, whose solutions are related to the solutions of the eigenvalue problem $Lx = \lambda x$, where $L = D^{-1}W$ is a row-stochastic matrix often dubbed as the *random walk normalized Laplacian*. Whenever the pair (λ, x) is an eigenvalue-eigenvector solution to $Lx = \lambda x$, then so is $(1 - \lambda, x)$ for $\tilde{L}x = \lambda Dx$. The (nonsymmetric) matrix L can also be interpreted as a transition probability matrix of a Markov chain with state space given by the nodes of the graph, and with entries L_{ij} denoting the one-step transition probability from node i to j .

In the diffusion map framework, one exploits a property of the top nontrivial eigenvector of the graph Laplacian of being piecewise constant on subsets of nodes in the domain that correspond to the same state associated to the underlying slow variable. We make this statement precise in section 4, and further use the resulting classification in section 5 to propose an unsupervised method for computing the stationary distribution of the hidden slow variable, without using any prior information on its structure. Since the top eigenvectors of the above Laplacian define the coarsest modes of variation in the data, and have a natural interpretation in terms of diffusion and random walks, they have been used in a very wide range of applications, including but not limited to partitioning [43, 42], clustering and community detection [35, 47, 34], image segmentation [39], ranking [49, 18], and data visualization and learning from data [7, 37].

The main application area studied in this paper involves stochastic models of chemical reaction networks. They are written in terms of stochastic simulation algorithms (SSAs) [22, 23] which have been used to model a number of biological systems, including the phage λ lysis-lysogeny decision circuit [1], circadian rhythms [46], and the cell cycle [28]. The Gillespie SSA [22] is an exact stochastic method that simulates every chemical reaction, sampling from the solution of the corresponding chemical master equation (CME). To characterize the behavior of a chemical system, one needs to simulate a large number of reactions and realizations, which leads to very computationally intensive algorithms. For suitable classes of chemically reacting systems, one can sometimes use exact algorithms which are equivalent to the Gillespie SSA, but are less computationally intensive, such as the Gibson–Bruck SSA [20] and the optimized direct method [5]. However, these methods also stochastically simulate the occurrence of every chemical reaction, which can be a computationally challenging task for systems with a very large number of species. One way to tackle this problem is to use parallel stochastic simulations [29]. In this work, we discuss an alternative approach which does not make use of parallel stochastic simulations, but at the same time, the proposed approach can also benefit from large processing power and parallel computing, as many steps of our proposed algorithms are highly parallelizable.

An alternative approach to treating the molecular populations as discrete random variables is to describe them in terms of their continuously changing concentration, which can be done via the chemical Langevin equation (CLE), a stochastic differential equation that links the discrete SSA with the deterministic reaction rate equations [21]. Although such an approach can be less computationally expensive, it

comes with the disadvantage that, for certain chemical systems, it can lead to negative populations [48]. In addition, note that none of the above approaches takes explicit advantage of the separation of scales, if one exists, something we will exploit in this paper, as detailed in sections 4 and 5.

It is often the case that a modeler is not interested in every single reaction which takes place in the system but only in the slowly evolving quantities. Certain systems possess multiple time scales, meaning that one has to simulate a large number of reactions to reveal the slow dynamics. Several algorithms for chemical networks with fast and slow variables have already been developed in the literature. The authors of [26] proposed simulating the fast reactions using Langevin dynamics and the slow reactions using the Gillespie algorithm. This approach requires both the time scale separation and a sufficiently large system volume; however, the latter constraint can be avoided using probability densities of the fast species conditioned on the slow species and estimating the effective propensity functions of the slow species [3, 4, 14, 38, 45]. An alternative approach to simulating the evolution of the slow variables while avoiding doing so for the fast variables is to estimate the probability distribution of the slow variables [19]. The key point in this approach is to use short bursts of appropriately initialized stochastic simulations to estimate the drift and diffusion coefficients for an approximating Fokker–Planck equation written in terms of the slow variables [16]. The success of this approach has already been demonstrated in a range of applications, including materials science [25], cell motility [15], and social behavior of insects [44].

Reference [8] introduces the conditional stochastic simulation algorithm (CSSA) that allows one to sample efficiently from the distribution of the fast variables conditioned on the slow ones [8] and to estimate the coefficients of the effective stochastic differential equation (SDE) on the fly via a proposed constrained multiscale algorithm (CMA). The CMA can be further modified by estimating the drift and diffusion coefficients in the form given by the CLE for the slow subsystem, which requires the estimation of effective propensity functions of slow reactions [9]. The main question we plan to address in this paper builds on and combines two already existing ideas investigated in [8] and [41] and brings several computational and algorithmic improvements. The above-mentioned CSSA algorithm explicitly makes use of the knowledge of the slow variables (often unavailable in many real applications), a drawback we plan to address as explained later in section 4, where, driven by the top eigenvector of an appropriately constructed Laplacian, we discover the underlying slow variable. In doing so, we make use of the ADM framework [41], which modifies the traditional diffusion map approach to take into account the time dependence of the data, i.e., the time stamp of each of the data points under consideration. By integrating local similarities at different scales, the ADM gives a global description of the data set.

The rest of this paper is organized as follows. In section 2 we provide a mathematical framework for multiscale modeling of stochastic chemical reaction networks and detail the two chemical systems via which we illustrate our approach. In section 3 we introduce the ADM-CLE framework and highlight its differences from the approaches which were previously introduced in the literature. In section 4 we propose a robust mapping from the observable space to the “dynamically meaningful” inaccessible space, which allows us to recover the hidden slow variables. In section 5 we introduce a Markov-based approach for approximating the steady distribution of the slow variable and compare our results with another recently proposed approach. We conclude with a summary and discussion of future work in section 6.

2. Problem formulation. A multiscale modeling framework for stochastic chemical reaction networks can be formulated as follows. We consider a well-mixed system of ℓ chemical species, denoted by X_1, X_2, \dots, X_ℓ , that interact through m reaction channels R_1, R_2, \dots, R_m in a reactor of volume V . We denote the state of the system by $X(t) = [X_1(t), X_2(t), \dots, X_\ell(t)]$, where $X_i(t)$, $i = 1, 2, \dots, \ell$, represents the number of molecules of type X_i in the system at time t . With a slight abuse of notation, we interchangeably use X_i to denote the type i of the molecule. In certain scenarios, one may assume that the reactions can be classified as either fast or slow, depending on the time scale of occurrence [3]. As expected, the fast reactions occur many times on a time scale for which the slow reactions occur with very small probability. As defined in [3], the fast species denoted by \mathbf{F} are those species whose population gets changed by a fast reaction. Slow species (denoted by \mathbf{S}) are not changed by fast reactions. Considering that slow species are not only species from the set $\{X_1, X_2, \dots, X_\ell\}$, but also that their functions are not changed by fast reactions, the components of the fast and slow species can be used as a basis for the state space of the system, whose dimension equals the number of linearly independent species.

For each reaction channel R_j , $j = 1, 2, \dots, m$, there exists a corresponding propensity function $\alpha_j \equiv \alpha_j(\mathbf{x})$, such that $\alpha_j dt$ denotes the probability that, given $\mathbf{X}(t) = \mathbf{x}$, reaction R_j occurs within the infinitesimal time interval $[t, t + dt)$. We denote by ν the stoichiometric matrix of size $m \times \ell$, with entry ν_{ji} denoting the change in the number of molecules of type X_i caused by one occurrence of reaction channel R_j . The continuous time discrete in space Markov chain can be further approximated by the CLE for a multivariate continuous Markov process [21]. Using time step Δt , the Euler–Maruyama discretization of the CLE is given by

$$X_i(t + \Delta t) = X_i(t) + \Delta t \sum_{j=1}^m \nu_{ji} \alpha_j(\mathbf{X}(t)) + \sum_{j=1}^m \nu_{ji} \sqrt{\alpha_j(\mathbf{X}(t))} N_j(t) \sqrt{\Delta t} \quad (2.1)$$

for all $i = 1, 2, \dots, \ell$,

where X_i , with another slight abuse of notation, denotes a real-valued approximation of the number of molecules of the i th chemical species, $i = 1, 2, \dots, \ell$. Here, $N_j(t)$, $j = 1, 2, \dots, m$, denotes the set of m independent normally distributed random variables with zero mean and unit variance.

2.1. Illustrative example CS-I. As the first illustrative example, we consider the following simple two-dimensional chemical system, with the two chemical species denoted by X_1 and X_2 (i.e., $\ell = 2$) which are subject to four reaction channels R_j , $j = 1, 2, 3, 4$ (i.e., $m = 4$), given by



Throughout the rest of this paper, we shall refer to the chemical system (2.2) as CS-I (i.e., “chemical system I”). We label R_j the reaction corresponding to the reaction rate subscript k_j , $j = 1, 2, 3, 4$, and note that each reaction R_j has a propensity function $\alpha_j(t)$ given by [22] associated to it:

$$\alpha_1(t) = k_1 V, \quad \alpha_2(t) = k_2 X_1(t), \quad \alpha_3(t) = k_3 X_2(t), \quad \alpha_4(t) = k_4 X_2(t), \quad (2.3)$$

where V denotes the volume of the reactor. We consider the system with the following dimensionless parameters:

$$k_1 V = 100, \quad k_2 = k_3 = 200, \quad \text{and} \quad k_4 = 1. \quad (2.4)$$




FIG. 2.1. (a) Trajectories of CS-I considered in (2.2) showing the behavior of the slow variable $S = (X_1 + X_2)/2$ in contrast to the behavior of the fast variables X_1 and X_2 , where the system propensity functions and parameters are given by (2.3) and (2.4). (b) Trajectories of CS-II considered in (2.7), showing the slow behavior of the variable $S = X_1 + 2X_2$ in contrast to the fast behavior of variables X_1 and X_2 , where the system parameters are given by (2.8).

We plot in Figure 2.1(a) the time evolution of the two different species in system (2.2), together with the slow variable $S = (X_1 + X_2)/2$, starting from initial conditions $X_1(0) = X_2(0) = 100$. As the figure shows, the system variables X_1 and X_2 are changing very frequently (thus we label them as fast variables), while the newly defined variable S changes very infrequently and can be considered to be a slow variable.


Following [27], for the chemical system in (2.2) comprised only of monomolecular reactions, it is possible to compute analytically the stationary distribution of the slow variable S , since the joint probability distribution of the two variables X_1 and X_2 is a multivariate Poisson distribution

$$(2.5) \quad \mathbb{P}(X_1 = n_1, X_2 = n_2) = \frac{\bar{\lambda}_1^{n_1} \bar{\lambda}_2^{n_2}}{n_1! n_2!} \exp(-\bar{\lambda}_1 - \bar{\lambda}_2)$$

with parameters given by

$$(2.6) \quad \bar{\lambda}_1 = \frac{k_1 V (k_3 + k_4)}{k_2 k_4} = 100.5 \quad \text{and} \quad \bar{\lambda}_2 = \frac{k_1 V}{k_4} = 100.$$

2.2. Illustrative example CS-II. The second example is taken from [8]. We shall refer to it as CS-II from now on. We consider the following system:



involving two molecular species X_1 and X_2 , whose reactions R_1, R_2, \dots, R_6 have the propensity functions given by

$$\begin{aligned} \alpha_1(t) &= k_1 X_2(t), & \alpha_2(t) &= k_2 X_1(t) X_2(t) / V, & \alpha_3(t) &= k_3 V, \\ \alpha_4(t) &= k_4 X_1(t), & \alpha_5(t) &= k_5 \frac{X_1(t)(X_1(t) - 1)}{V}, & \alpha_6(t) &= k_6 X_2(t), \end{aligned}$$

where V denotes the system volume. Figure 2.1(b) shows a simulated trajectory of this chemical system using the Gillespie algorithm for the following dimensionless parameters [8]:

$$(2.8) \quad k_1 = 32, \quad k_2 = 0.04V, \quad k_3 V = 1475, \quad k_4 = 19.75, \quad k_5 = 10V, \quad k_6 = 4000,$$

where we use $V = 8$. Note that in this second example, reactions R_5 and R_6 are occurring on a much faster time scale than the other four reactions R_1 , R_2 , R_3 , and R_4 . A natural choice for the slow variable is $S = X_1 + 2X_2$, which is invariant with respect to all fast reactions [8], as we illustrate in Figure 2.1(b).

2.3. Main problem. Our end goal in this paper is to propose an algorithm that efficiently and accurately estimates the stationary probability density of the hidden slow variable S , without any prior knowledge of it. The approach we propose builds on the anisotropic diffusion map (ADM) framework to implicitly discover the mapping from the observable state space to the dynamically meaningful coordinates of the fast and slow variables, as previously introduced in [41], and on the CLE approximation (2.1).

3. ADM-CLE approach. Let us consider example CS-II, and assume that $s = s(x_1, x_2) = x_1 + 2x_2$ and $f = f(x_1, x_2) = x_1$ are the slowly and rapidly changing variables, respectively. They together define a mapping $g : (x_1, x_2) \mapsto (s, f)$ from the observable state variables x_1 and x_2 in the accessible space \mathcal{O} to the “dynamically meaningful” (but in more complicated examples inaccessible) slow variable s and the fast accessible variable f , both in space \mathcal{H} . In other words, g maps $(x_1, x_2) \mapsto (x_1 + 2x_2, x_1)$ and, conversely, its inverse $h := g^{-1} : (s, f) \mapsto (f, \frac{s-f}{2})$.

The approach introduced in [41] exploits the local point clouds generated by many local bursts of simulations at each point (x_1, x_2) in the observable space \mathcal{O} . Such observable local point clouds are the image under h of similar local point clouds in the inaccessible space \mathcal{H} (at corresponding coordinates (s, f) such that $h(s, f) = (x_1, x_2)$), which, due to the separation of scales between the fast and slow variables f and s , have the appearance of thin elongated ellipses. It is precisely this separation of scales that we leverage into building a sparse anisotropic graph Laplacian L in the observable space, and we use it as an approximation of the isotropic graph Laplacian in the inaccessible space \mathcal{H} . As we shall see, the top nontrivial eigenvector of L will robustly indicate all pairs of original states (x_1, x_2) that correspond to the same slow variable $S = s$ (where $s = x_1 + 2x_2$ for CS-II). In other words, we discover on the fly the structure of the slow variable S , and further integrate this information into a Markov-based method for estimating its stationary distribution $\mathbb{P}(S = s)$, while also computing along the way an analytical expression for the conditional distribution of the fast variable given the slow variable $\mathbb{P}(F = f | S = s)$.

Singer et al. [41] run many local bursts of simulations for a short time step δt starting at (x_1, x_2) . Such trajectories end up at random locations, forming a cloud of points in the observable plane \mathcal{O} , with a bivariate normal distribution with 2×2 covariance matrix Σ . The shape of the resulting point cloud is an ellipse, whose axes reflect the dynamics of the data points. In other words, when there is a separation of scales, the ellipses are thin and elongated, with the ratio between the axis of the ellipse given by the ratio

$$(3.1) \quad \tau = \frac{\widehat{\lambda}_1}{\widehat{\lambda}_2}$$

of the two eigenvalues of Σ . The first eigenvector corresponding to $\widehat{\lambda}_1$ points in the direction of the fast dynamics on the line $x_1 + 2x_2 = s$, while the second one points in the direction of the slow dynamics. In particular, τ is a small parameter, i.e., $0 < \tau \ll 1$. In general, we wish to piece together locally defined components into a globally consistent framework, a nontrivial task when the underlying unobservable

slow variables (or the propensity functions of the system) are complicated nonlinear functions of the observable variables in \mathcal{O} .

The construction of the ADM framework in [41] relates the anisotropic graph Laplacian in the observable space \mathcal{O} with the isotropic graph Laplacian in the inaccessible space \mathcal{H} . In that setup, each of the N data points $\mathbf{x}^{(i)}$, $i = 1, 2, \dots, N$, lives in an ℓ -dimensional data space. For both CS-I and CS-II, the data is two-dimensional, and thus $\ell = 2$. For the former system, we consider each lattice point in the domain $[50, 150] \times [50, 150]$, and hence there are $N = 101^2 = 10,201$ states, while for the latter one we consider the domain $[1, 110] \times [1, 110]$, i.e., $N = 110^2 = 12,100$. Throughout this paper, we will often refer to the N data points $\mathbf{x}^{(i)} = (x_1, x_2)^{(i)}$, $i = 1, 2, \dots, N$, as \mathcal{O} -states of the chemical system. The ADM [41] then generates ensembles of short simulation bursts at each of the N points in the data set, computes the averaged position after statistically averaging over the many simulated trajectories, and obtains an estimate of the local 2×2 covariance matrix $\Sigma_{(i)}$. For each data point $\mathbf{x}^{(i)}$, the inverse of $\Sigma_{(i)}$ is computed and symmetric Σ -dependent squared distance between pairs of data points in the two-dimensional observable space \mathbb{R}^2 (given by (3.3) below) is defined. The ADM framework then uses this dynamic distance measure to approximate the Laplacian on the underlying hidden slow manifold. We provide further details on the ADM framework in section 3.2. We now highlight the first difference between the approach taken in this paper and in [41].

3.1. Replacing short simulation bursts by the CLE approximation. The local bursts of simulations initiated at each data point in order to estimate the local covariances may be computationally expensive to estimate. In this paper, we bypass these short bursts of simulations by using an approximation given by the CLE (2.1), which allows for a theoretical derivation of the local 2×2 covariance matrices. Using (2.1), we obtain

$$\begin{aligned}
 \text{Cov}(X_i(t + \Delta t), X_k(t + \Delta t)) &= \mathbb{E}[X_i(t + \Delta t)X_k(t + \Delta t)] \\
 &\quad - \mathbb{E}[X_i(t + \Delta t)]\mathbb{E}[X_k(t + \Delta t)] \\
 (3.2) \qquad \qquad \qquad &= \Delta t \sum_{j=1}^m \nu_{ji} \nu_{jk} \alpha_j(\mathbf{X}).
 \end{aligned}$$

Computing the eigendecomposition of a local covariance matrix is analogous to performing principal component analysis on the local cloud of points, generated by the short simulation bursts. The advantage of (3.2) over the computational approach used in [41] is that $\Sigma_{(i)}$ can be computed at each data point without running (computationally intensive) short bursts of simulations. The error of the CLE approximation depends on the values of coordinates of the data point $\mathbf{x}^{(i)}$, i.e., on the system volume V [24, 13].

However, for the chemical system CS-II, the domain does include states with small molecule numbers, and the CLE approximation is less well justified. From looking at the estimated distribution of the slow variable $S = X_1 + 2X_2$ shown in Figure 5.4(b), the most probable states contain a lot less than 100 molecules of each chemical species; however, the fact that we are able to robustly recover the distribution even in this scenario is perhaps a good indication that the CLE approximation was justified. To circumvent the assumptions needed for the CLE approximation, one could alternatively consider estimating, for small molecule numbers, the local covariance matrices via direct simulation, as opposed to relying on the CLE approximation. A potential improvement in this direction may also alleviate the issues encountered in the

slow variable detection step via our bin denoising procedure detailed later in section 4, which leads us to truncate at the boundary of the slow variable (perhaps due to unreliable CLE approximation). Combining the CLE approximation for large numbers of molecules with direct simulation for the few states that involve small molecule numbers is an interesting research direction to explore [12].

3.2. Anisotropic diffusion kernels. The next task is the integration of all local principal components into a global framework, with the purpose of identifying the hidden slow variable. We estimate the distance (and hence the similarity measure) between the slow variables in the underlying inaccessible manifold using the anisotropic graph Laplacian [41]. We derive a symmetrized second order approximation of the (unknown) distances in the inaccessible space \mathcal{H} , based on the Jacobian of the unknown mapping from the inaccessible to the observable space. The Σ -dependent distance between two \mathcal{O} -states is given by

$$(3.3) \quad d_{\Sigma}^2 \left((x_1, x_2)^{(i)}, (x_1, x_2)^{(j)} \right) = \frac{1}{2} \left((x_1, x_2)^{(i)} - (x_1, x_2)^{(j)} \right) \left(\Sigma_{(x_1, x_2)^{(i)}}^{-1} + \Sigma_{(x_1, x_2)^{(j)}}^{-1} \right) \times \left((x_1, x_2)^{(i)} - (x_1, x_2)^{(j)} \right)^T$$

and represents a second order approximation¹ of the Euclidean distance in the inaccessible $(s, \tau f)$ -space

$$(3.4) \quad d_{\Sigma}^2[(x_1, x_2)^{(i)}, (x_1, x_2)^{(j)}] \approx (s^{(i)} - s^{(j)})^2 + \tau^2 (f^{(i)} - f^{(j)})^2 \approx (s^{(i)} - s^{(j)})^2,$$

where the last approximation is due to the fact that τ is a small parameter; see (3.1).

Note that it is also possible to extend (3.4) to higher dimensions, as long as there exists a separation of scales between the set of slow variables and the set of fast variables [41]. Using approximation (3.3)–(3.4) of the distance between states of the slow variable, we next construct (an approximation of) the Laplacian on the underlying hidden slow manifold, using the Gaussian kernel as a similarity measure between the slow variable states. We build an $N \times N$ similarity matrix W with entries

$$(3.5) \quad W_{ij} = \exp \left\{ \frac{-d_{\Sigma}^2[(x_1, x_2)^{(i)}, (x_1, x_2)^{(j)}]}{\varepsilon^2} \right\} \approx \exp \left\{ \frac{-(s^{(i)} - s^{(j)})^2}{\varepsilon^2} \right\},$$

$i, j = 1, 2, \dots, N,$

where the single smoothing parameter ε (the kernel scale) has a two-fold interpretation. On one hand, ε denotes the squared radius of the neighborhood used to infer local geometric information; in particular, W_{ij} is $O(1)$ when $s^{(i)}$ and $s^{(j)}$ are in a ball of radius $\sqrt{\varepsilon}$, and thus close on the underlying slow manifold, but it is exponentially small for states that are more than $\sqrt{\varepsilon}$ apart. On the other hand, ε represents the discrete time step at which the underlying random walker jumps from one point to another. We refer the reader to [32] for a detailed survey of random walks on graphs

¹The principal components of Σ are the local directions of the rapidly changing variables at a particular state of the system, whereas components with small eigenvalues correspond to the slow variables. We remark that a singular covariance matrix would correspond to having only changes of the fast variable, and almost no changes in the slow variable, which is unlikely but possible, depending on the system under consideration. We have not run into such problems in the examples considered. Furthermore, we remark that in the initial work of Singer and Coifman [40], which introduced the above second order approximation of the Euclidean distance in the inaccessible space, the authors rely on the pseudoinverse Σ^\dagger .

and their applications. We normalize W using the diagonal matrix D to define the row-stochastic matrix L by

$$(3.6) \quad D_{ii} = \sum_{j=1}^N W_{ij}, \quad L = D^{-1}W.$$

Since L is a row-stochastic matrix, it has eigenvalue $\lambda_0 = 1$ with trivial eigenvector $\Phi_0 = (1, 1, \dots, 1)^T$. The remaining eigenvalues can be ordered as

$$1 = \lambda_0 \geq \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{N-1}.$$

We denote by Φ_i the corresponding eigenvectors, i.e., $L\Phi_i = \lambda_i\Phi_i$. The top d nontrivial eigenvectors of the random walk anisotropic Laplacian L describe the geometry of the underlying d -dimensional manifold [17]; i.e., the i th data point $\mathbf{x}^{(i)}$ is represented by the following diffusion map:

$$(3.7) \quad (\Phi_1(i), \Phi_2(i), \dots, \Phi_d(i)), \quad i = 1, 2, \dots, N,$$

where $\Phi_j(i)$ denotes the i th component of the eigenvector Φ_j . However, note that some of the considered eigenvectors can be higher harmonics of the same principal direction along the manifold; thus in practice one computes the correlation between the computed eigenvectors before selecting the above d eigenvectors chosen to parametrize the underlying manifold. For the two chemical systems considered in this paper, we show in the remainder of this section how the top (i.e., $d = 1$) nontrivial eigenvector of L can be used to successfully recover the underlying slow variable.

Using the stochasticity of L , we can interpret it as a random walk matrix on the weighted graph $G = (V, E)$, where the set of nodes corresponds to the original observable states $(x_1, x_2)^{(i)}$, $i = 1, 2, \dots, N$ (and implicitly to states $s^{(i)}$ of the slow variable), and there is an edge between nodes i and j if and only if $W_{ij} > 0$. The associated combinatorial Laplacian is given by $\tilde{L} = D - W$. Whenever the pair (λ_i, Φ_i) is an eigenvalue-eigenvector solution to $L\Phi_i = \lambda_i\Phi_i$, then so is $(1 - \lambda_i, \Phi_i)$ for the generalized eigenvalue problem $\tilde{L}\Phi_i = \lambda_i D\Phi_i$. We plot in Figures 3.1(a) and 3.2(a) the spectrum of the combinatorial Laplacian $\tilde{L} = D - W$ for the chemical systems CS-I and CS-II. In Figures 3.1(b) and 3.2(b) we color the states of the network with the top nontrivial eigenvector Φ_1 .

Before considering the top eigenvector of L for determining the underlying slow variable and estimating its stationary distribution, we propose using a sparse graph Laplacian which differs from the ADM method in [41], where the Laplacian matrix is associated to a complete weighted graph. However, using a complete graph leads to computing the Σ -dependent squared distance in (3.3) for any pair of nodes; thus an $O(N^2)$ number of computations is used. In light of the approximation (3.4), a pair of points which are far away in the observable space (i.e., for which $d_\Sigma^2((x_1, x_2)^{(i)}, (x_1, x_2)^{(j)})$ is large) denotes a pair of corresponding states of the slow variable which are also far away in the inaccessible space. Thus we do not have to do such computations, because points far away in the unobservable space will have an exponentially small similarity W_{ij} close to 0. The fact that the shape of the local point cloud is an ellipse provides some insight in this direction. Thus we will build a sparse graph of pairwise measurements and no longer compute the Σ -dependent distance between all points of the data set but only between a very small subset of the points. The spectrum of the covariance matrix Σ_i , in particular the ratio τ of its two eigenvalues given by (3.1), guides us in building locally at each point a sparse

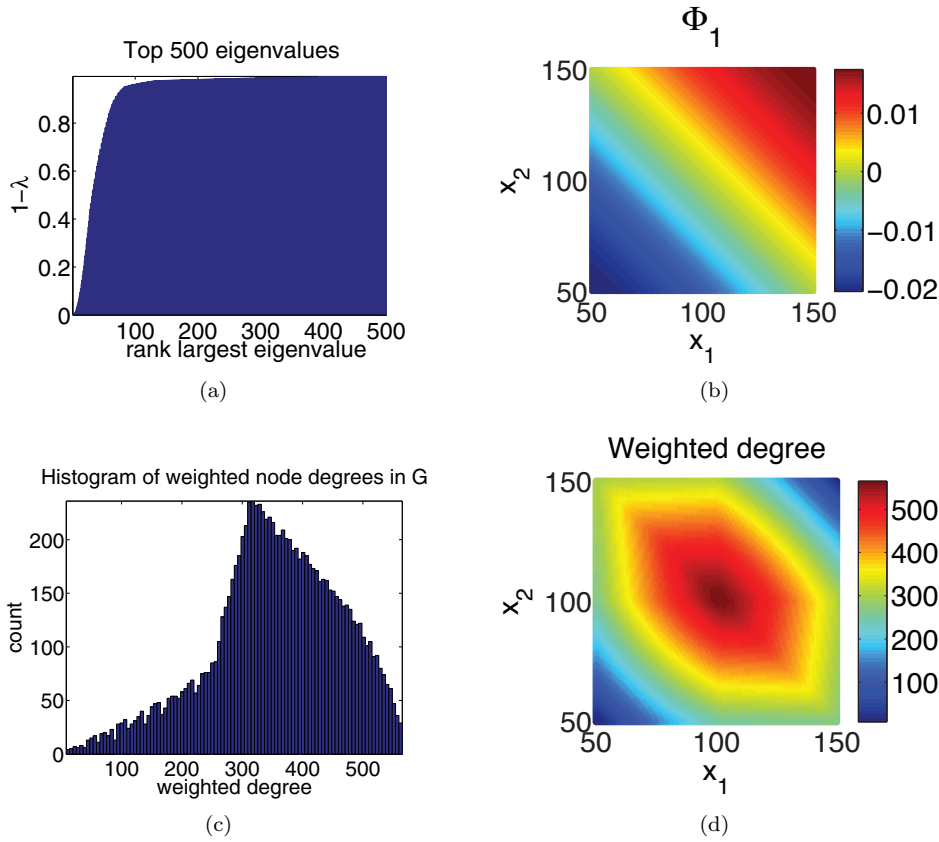


FIG. 3.1. Illustrative example CS-I. (a) The top 500 eigenvalues of the associated combinatorial Laplacian, i.e., $(1 - \lambda_i)$ for $i = 1, 2, \dots, 500$. (b) The coloring of the nodes of G (states of the observable space) according to their corresponding entry in the top eigenvector Φ_1 of L given by (3.6). (c) The weighted degree distribution of the ground state graph G . (d) A scatterplot of the states of the system, colored by their weighted degree.

ellipsoid-like neighborhood graph. However, we cautiously remark that the second order approximation (3.3) of the distance in the inaccessible space is not a monotonically increasing function of the distance in the accessible space and depends on the particular covariance matrices used and potentially the separation of scales. In the examples we have considered, we have indeed observed a monotonically decreasing Gaussian similarity (3.5) as a function of the distance between the points in the accessible space, which motivated our approach to consider ellipsoid-like neighborhood graphs.

For each observable state $(x_1, x_2)^{(i)}$, we build a local adjacency graph, denoted by G_i , in the shape of an ellipse pointing in the direction of the fast dynamics, and whose small axis points in the direction of the slow dynamics. In our computations, we fix the length of the semiminor axis of any ellipse to be equal to 3, while the length of its corresponding semimajor axis is given by 3τ , where τ is given by the ratio of the two eigenvalues of the covariance matrix as in (3.1). Figure 3.3 shows an example of such a local 1-hop neighborhood graph G_i , where the central node $(x_1, x_2)^{(i)}$ is connected to all points contained within the boundaries of an appropriately scaled ellipse centered at $(x_1, x_2)^{(i)}$. Finally, we define the sparse graph G of size $N \times N$ associated to the

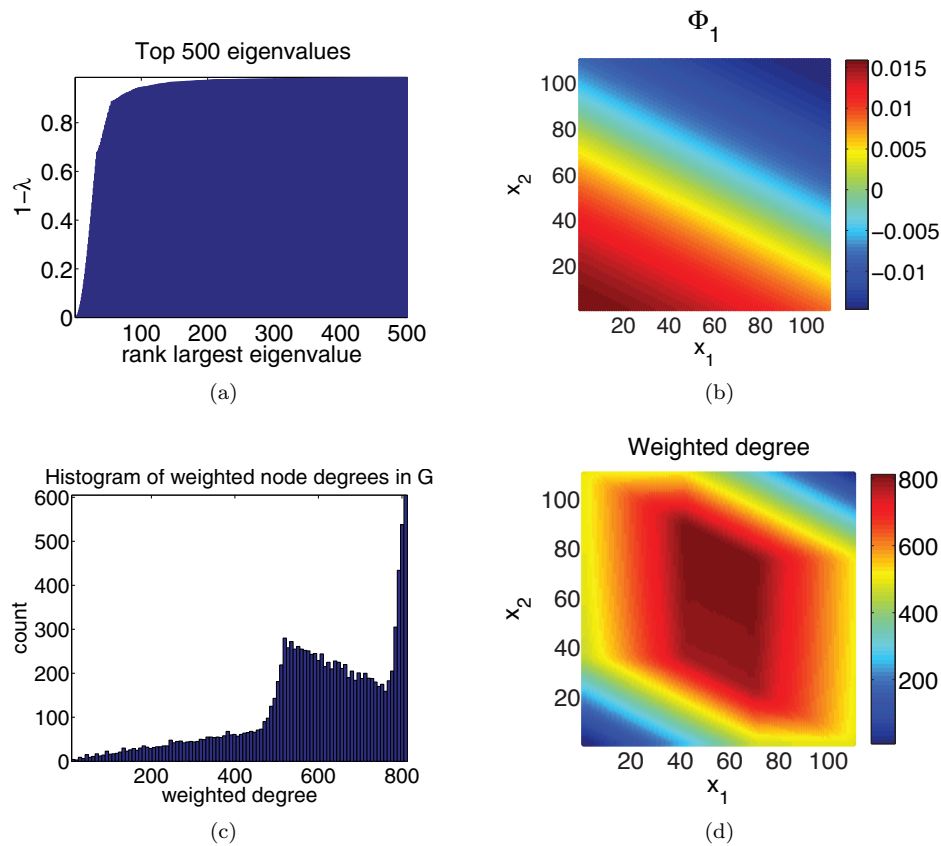


FIG. 3.2. Illustrative example CS-II. (a) The top 500 eigenvalues of the associated combinatorial Laplacian, i.e., $(1 - \lambda_i)$ for $i = 1, 2, \dots, 500$. (b) The coloring of the nodes of G (states of the observable space) according to their corresponding entry in the top eigenvector Φ_1 of L given by (3.6). (c) The weighted degree distribution of the ground state graph G . (d) A scatterplot of the states of the system, colored by their weighted degree.




FIG. 3.3. The local neighborhood graph G_i at a given node $(x_1, x_2)^{(i)}$; the shape is an ellipsoid whose axis ratio is given by the ratio of the eigenvalues τ of the local covariance matrix $d_{\Sigma}^2((x_1, x_2)^{(i)}, (x_1, x_2)^{(j)})$, i.e., by (3.1). The corresponding eigenvectors are used to calculate the orientation of the ellipse.

entire network as the union of all locally defined ellipsoid-like neighborhood graphs $G = \bigcup_{i=1}^N G_i$. Note that the union graph G is still a simple graph, with no self-edges and no multiple edges connecting the same pair of nodes. We compute the distance d_Σ by (3.3) (and thus the similarity W_{ij}) between a pair of nodes $(x_1, x_2)^{(i)}$ and $(x_1, x_2)^{(j)}$ if and only if the corresponding edge (i, j) exists in G .

We plot in Figures 3.1(c) and 3.2(c) the histogram of the weighted degrees of the nodes in the weighted graph W defined in (3.5), while Figures 3.1(d) and 3.2(d) show a scatterplot of the states of the system, where each state i is colored by its weighted degree, i.e., the sum of all its outgoing weighted edges W_{ij} , $j = 1, 2, \dots, n$. Throughout the computational examples in this paper, the smoothing parameter ε which appears in (3.5) was set to $\varepsilon = 0.1$. In contrast to the approach in [41], which computes all $O(N^2)$ pairwise similarities, the sparsity of G (and thus of the associated graph Laplacian L) in our approach only requires the computation of a much smaller number of distances, as low as linear, depending on the discretization of the domain, and makes it computationally feasible to solve problems with thousands or even tens of thousands of nodes.

4. A robust mapping from the observable space \mathcal{O} to the “dynamically meaningful” inaccessible space \mathcal{H} . As a first step towards partitioning the nodes of the original graph G and detecting the associated slow variable, we sort the entries of the top eigenvector Φ_1 , which we then denote by $\bar{\Phi}_1$ with $\bar{\Phi}_1(1) \geq \bar{\Phi}_1(2) \geq \dots \geq \bar{\Phi}_1(N)$. This sorting process defines permutation σ of the original index set $i = 1, 2, \dots, N$ so that $\bar{\Phi}_1(\sigma(i)) = \Phi_1(i)$. We consider the increments between two consecutive (sorted) values

$$(4.1) \quad \delta_i = \bar{\Phi}_1(i) - \bar{\Phi}_1(i+1), \quad i = 1, 2, \dots, N-1.$$

Next, we sort the vector of such increments, denote its entries by $\bar{\delta}_1 \geq \bar{\delta}_2 \geq \dots \geq \bar{\delta}_{N-1}$, and show in Figure 4.1(a) (resp., Figure 4.2(a)) the top 300 (resp., top 420) largest such increments $\bar{\delta}_i$ for illustrative example CS-I (resp., CS-II). Note that this already gives us an idea about the number of distinct slow states in the system, a set which we denote by \mathcal{S} . Ideally, the difference $\Phi_1(i) - \Phi_1(j)$ in the entries of the top eigenvector corresponding to two observable states $(x_1, x_2)^{(i)}$ and $(x_1, x_2)^{(j)}$ that belong to the same slow variable s (i.e., $x_1^{(i)} + 2x_2^{(i)} = x_1^{(j)} + 2x_2^{(j)} = s$ for illustrative example CS-II) should be zero or close to zero, in which case we expect that only approximately $|\mathcal{S}|$ of the $N-1$ increments δ_i are significantly larger than zero, while the remaining majority are zero or close to zero.

In Figures 4.1(b) and 4.2(b) we highlight the correlation between the entries of the top nontrivial eigenvector Φ_1 and the corresponding slow variable S . In Figures 4.1(c) and 4.2(c), we zoom on a subset of states to make the point that the eigenvector Φ_1 is almost constant on the \mathcal{O} -states that correspond to the same value of the slow variable. The plots in Figures 3.1(b) and 3.2(b) show a coloring of the networks generated by the two chemical systems CS-I and CS-II, based on the first nontrivial eigenvector of the associated sparse Laplacian L . Note that the eigenvector looks almost piecewise constant along the lines that point to the evolution of the fast variable, for a given value of the slow variable ($S = (X_1 + X_2)/2$ for CS-I and $S = X_1 + 2X_2$ for CS-II), yet nowhere along the way do we have to input this information into the method. In the next step we use this top eigenvector to identify all nodes of the graph (original states of the chemical system) that correspond to the same value of the underlying slow variable. In other words, all nodes whose corresponding eigenvector entries are between an appropriately chosen interval (that we shall refer to as a *bin*) will be

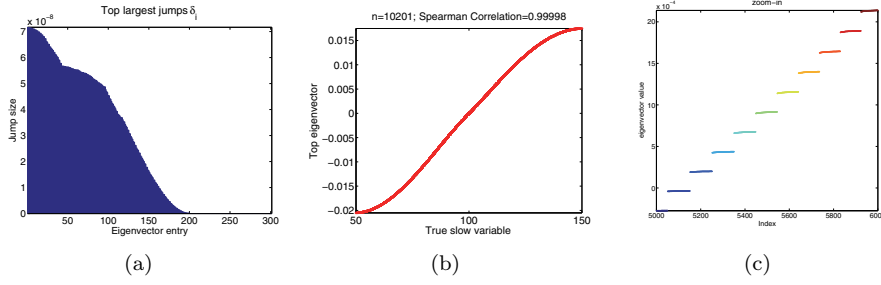


FIG. 4.1. Illustrative example CS-I. (a) Jump sizes (4.1) of the sorted eigenvector Φ_1 of the sparse anisotropic graph Laplacian L . (b) The correlation of Φ_1 with the ground truth slow variable $S = (X_1 + X_2)/2$. On the y-axis we plot the i th entry of the top eigenvector $\Phi_1(i)$ associated to state $(x_1, x_2)^{(i)}$ versus its corresponding true slow variable $s = (x_1 + x_2)/2$. (c) Zoom-in on the sorted top eigenvector Φ_1 (the colors denote the corresponding slow variable) showing that Φ_1 is almost piecewise constant on the bins that correspond to distinct slow variable states. The kernel scale is set to $\epsilon = 0.1$.

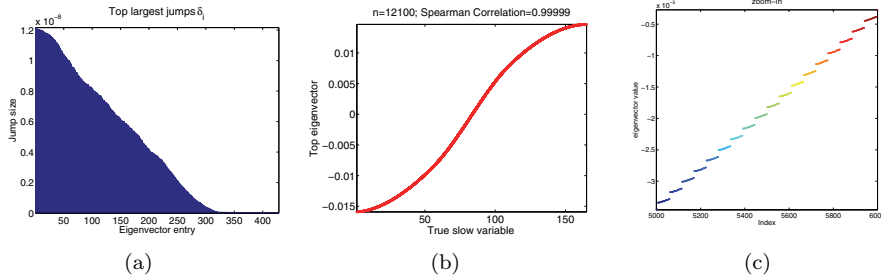


FIG. 4.2. Illustrative example CS-II. (a) Jump sizes (4.1) of the sorted eigenvector Φ_1 of the sparse anisotropic graph Laplacian L . (b) The correlation of Φ_1 with the ground truth slow variable $S = X_1 + 2X_2$. On the y-axis we plot the i th entry of the top eigenvector $\Phi_1(i)$ associated to state $(x_1, x_2)^{(i)}$ versus its corresponding true slow variable $s = x_1 + 2x_2$. (c) Zoom-in on the sorted top eigenvector Φ_1 (the colors denote the corresponding slow variable) showing that Φ_1 is almost piecewise constant on the bins that correspond to distinct slow variable states. The kernel scale is set to $\epsilon = 0.1$.

labeled as belonging to the same slow variable S . In other words, we seek a partition of the observable states in \mathcal{O} , i.e., of the nodes of G , such that all original states $(x_1, x_2)^{(i)}$ with the same value of the corresponding slow variable $s((x_1, x_2)^{(i)})$ end up in the same bin. Our goal is to find a partition $\mathcal{P} = \{\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_k\}$ of \mathcal{O} such that

$$(4.2) \quad \mathcal{P}_j = \{(x_1, x_2)^{(i)} \in \mathcal{O} \mid s(x_1, x_2) = q_j\} \quad \text{and} \quad \bigcup_{j=1}^k \mathcal{P}_j = \mathcal{O},$$

where k denotes the number of distinct values q_j , $j = 1, 2, \dots, k$, of the slow variable S . As an example, in the case of CS-I given by (2.2), the partition $\mathcal{P}_j = \{(1, 99), (2, 98), \dots, (99, 1)\}$ corresponds to all nodes in the graph for which the value of the associated slow variable is constant: $q_j = 50$. The key observation we exploit here is that the top eigenvector of the Laplacian matrix is almost piecewise constant on the bins that partition \mathcal{O} , since the nodes of G that correspond to the same value of the slow variable have a very high pairwise similarity, with W_{ij} very close to 1.

One may also interpret the above problem as a clustering problem, where the




FIG. 4.3. *Illustrative example CS-I. (a) The eigenvector-based slow variable cardinality. The Theta score Θ is the smoothness measure of the bin cardinalities, defined in (4.5). The algorithm perfectly recovers the ground truth partition. (b) The heatmap of the pairwise Jaccard similarity matrix given by (4.4). (c) The correlation between the ordering of the ground truth slow variable and the eigenvector recovered slow variable. (d) The Jaccard index of the pairwise matched bins (from the maximum matching).*

similarity between pairs of points is given by (3.5), and is such that nodes that belong to the same bin have a much higher similarity compared to nodes that belong to two different bins, an effect due to the strong separation of scales. In the case of illustrative example CS-I, the clusters correspond to lines in the two-dimensional plane such that $(x_1 + x_2)/2 = c$ for a constant c . We point the interested reader to the work of [33], where the top eigenvectors of the random walk Laplacian are used for clustering. While in practice one uses the top several eigenvectors as the reduced eigenspace where clustering is subsequently performed, in our case the top eigenvector alone suffices to capture the many different clusters (i.e., bins), a fact we attribute to the strong separation of scales exhibited by the illustrative chemical systems CS-I and CS-II. If several eigenvectors were considered, then one could use a clustering algorithm, such as k -means or spectral clustering [39, 47], to obtain the partitioning (4.2). However, a simpler method has been successfully used for the one-dimensional eigenspace in both examples we considered. It is described as follows. Recall the sorted vector of increments $\bar{\delta}_1 \geq \bar{\delta}_2 \geq \dots \geq \bar{\delta}_{N-1}$ defined in (4.1), and consider the set of the $k-1$ largest such increments $\{\bar{\delta}_1, \bar{\delta}_2, \dots, \bar{\delta}_{k-1}\}$ where $\bar{\delta}_1 \geq \bar{\delta}_2 \geq \dots \geq \bar{\delta}_{k-1}$. Next, from the sorted eigenvector $\bar{\Phi}_1$ we extract the position of the entries whose associated increment (with respect to its right-next neighbor index) belongs to $\{\bar{\delta}_1, \bar{\delta}_2, \dots, \bar{\delta}_{k-1}\}$. In others words, we compute

$$(4.3) \quad b_t = \arg \min_{i=1,2,\dots,N-1} \bar{\Phi}_1(i) - \bar{\Phi}_1(i+1) = \bar{\delta}_t, \quad \text{where } t = 1, 2, \dots, k-1,$$

and $b_0 = 0$ and $b_k = N$. Finally, we compute an estimated partition $\hat{\mathcal{P}}$ of \mathcal{O} by $\hat{\mathcal{P}}_q = \{i \in \mathcal{O} \mid \sigma(i) \in (b_{q-1}, b_q]\}$, where $q = 1, 2, \dots, k$, and σ is the permutation of the original index set $i = 1, 2, \dots, N$, given in the definition of $\bar{\Phi}_1$, i.e. $\bar{\Phi}_1(\sigma(i)) = \Phi_1(i)$. We remark here that k is not a parameter, per se, but represents the number of distinct states of the slow variable in the system, corresponding to the current bounded domain. If not known a priori, one can estimate k from the sorted barplot of the increments between two consecutive entries in the sorted eigenvector, as shown in Figures 4.1(a) and 4.2(a).

To illustrate the correctness of our proposed technique, we compute the Jaccard index between each proposed partition set $\hat{\mathcal{P}}_j$, $j = 1, 2, \dots, k$, and each ground truth

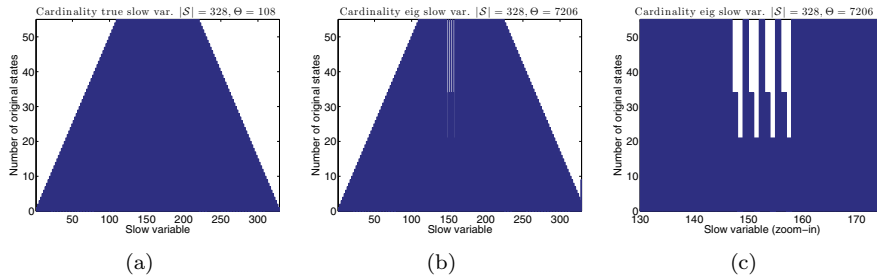


FIG. 4.4. Illustrative example CS-II. (a) The ground truth slow variable cardinality. (b) The eigenvector-based slow variable cardinality. Θ captures the smoothness of the bin cardinalities, as introduced in (4.5). (c) Plot of the cardinalities of a subset of bins, showing the erroneous bin assignments in the eigenvector-based partition. This is a zoomed-in version of panel (b).

partition set \mathcal{P}_i , $i = 1, 2, \dots, |\mathcal{S}|$:

$$(4.4) \quad J_{ij} = \frac{|\mathcal{P}_i \cap \hat{\mathcal{P}}_j|}{|\mathcal{P}_i \cup \hat{\mathcal{P}}_j|}, \quad \text{where } i = 1, 2, \dots, |\mathcal{S}|, \quad j = 1, 2, \dots, k,$$

and we show a heatmap of this matrix in Figure 4.3(b). Since we are interested not only in the partition but also in recovering the ordering of the slow variable, we show in Figure 4.3(c) the correlation between the ground truth ordering of the slow variable and our recovered ordering. Note that we can only recover the ordering up to a global sign, since $-\Psi_1$ is also an eigenvector of L . Finally, we compute the maximum weight matching (using, for example, the Hungarian method [30]) in the bipartite graph with node set $\mathcal{P} \cup \hat{\mathcal{P}}$ and edges across the two sets given by matrix J in (4.4). In Figure 4.3(d) we plot the Jaccard index of the matched partitions. For the first chemical system CS-I, note that the algorithm perfectly recovers the ground truth partition. In Figure 4.4 we present the outcome of the binning algorithm for the illustrative example CS-II, which is no longer satisfactorily by itself and requires further improvement. Though the bin cardinalities in the initial solution visually resemble the ground truth, there are numerous mistakes being made. To illustrate this, for a given partition \mathcal{P} , we compute the following measure of continuity of the recovered bin cardinalities:

$$(4.5) \quad \Theta_{\mathcal{P}} = \Theta(\mathcal{P}_1, \dots, \mathcal{P}_S) = \sum_{i=1}^{S-1} (|\mathcal{P}_i| - |\mathcal{P}_{i+1}|)^2.$$

In other words, $\Theta_{\mathcal{P}}$ captures the squared difference in the cardinalities of two consecutive bins. For the chemical system CS-II, the ground truth yields a score $\Theta = 108$, while for the eigenvector-recovered solution $\Theta = 7206$, thus indicating already that numerous misclassifications are being made, without even computing the Jaccard similarity matrix (4.4) between the two partitions. To this end, we introduce in the next subsection a heuristic denoising technique followed by a truncation of the domain, which altogether lead to a better partitioning of \mathcal{O} into groups of states that correspond to the same slow variable.

4.1. A bin denoising scheme. While the eigenvector-based partition procedure detailed above yields accurate results for CS-I in (2.2), this procedure alone is not sufficient for obtaining a satisfactory partition for the more complex CS-II considered in (2.7), as illustrated by the high corresponding Θ -score ($\Theta = 7206$) shown

Algorithm 1. Bin-merging algorithm.

```

1: Initialize FLAG = TRUE
2: while FLAG is TRUE do
3:   Compute  $\alpha_i := \Theta(\mathcal{P}_1, \dots, \mathcal{P}_i \cup \mathcal{P}_{i+1}, \dots, \mathcal{P}_{|\mathcal{S}|}) \quad \forall i = 1, 2, \dots, |\mathcal{S}| - 1$  using
   definition (4.5)
4:   if  $\min_{i=1,2,\dots,|\mathcal{S}|-1} \alpha_i < \Theta(\mathcal{P}_1, \dots, \mathcal{P}_{|\mathcal{S}|})$  then
5:      $q = \operatorname{argmin}_{i=1,2,\dots,|\mathcal{S}|-1} \alpha_i$ 
6:      $\mathcal{P} := \mathcal{P}_1, \dots, \mathcal{P}_q \cup \mathcal{P}_{q+1}, \dots, \mathcal{P}_{|\mathcal{S}|}$ 
7:      $|\mathcal{S}| = |\mathcal{S}| - 1$ 
8:   else
9:     FLAG = FALSE
10:  end if
11: end while

```

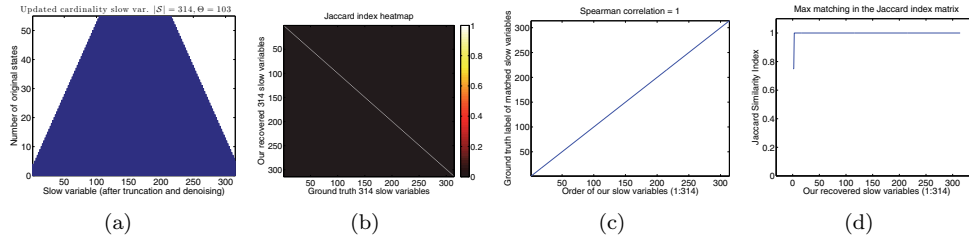


FIG. 4.5. *Illustrative example CS-II.* (a) *The eigenvector-based slow variable cardinality after truncating and bin denoising. The Theta score Θ is the smoothness measure of the bin cardinalities, defined in (4.5).* (b) *The heatmap of the pairwise Jaccard similarity matrix given by (4.4).* (c) *The correlation between the ordering of the ground truth slow variable and the eigenvector recovered slow variable.* (d) *The Jaccard index of the pairwise matched bins (from the maximum matching).*

in Figure 4.4(b). In Figure 4.4(c) we zoom into some of the recovered bins, showing that the eigenvector-based reconstruction splits some of the inner bins, which explains the high associated Θ -score. In other words, states/bins which in the ground truth solution correspond to the same values of the slow state variable, are divided into two adjacent bins, and are mistaken for two distinct states of the slow variable.

To solve this issue, we propose a bin-denoising heuristic that robustly assigns data points to their respective bins. In hindsight, the continuity of the eigenfunctions of the Laplacian should be reflected in the continuity of the histogram of state counts in bins corresponding to adjacent intervals. We detail in Algorithm 1 an iterative heuristic procedure which, at each step, merges two adjacent bins such that the resulting Θ -score is minimized across all possible pairs of adjacent bins that can be merged. We show in Figure 4.5(a) the resulting bin cardinalities after the bin-merging heuristic and after truncating at the boundary of the slow variable. Note that the new denoised partition yields $\Theta = 103$, and the number of bins (states of the slow variable) decreases from $|\mathcal{S}| = 328$ to $|\mathcal{S}| = 314$. Furthermore, in Figure 4.5(b) we compute the Jaccard similarity matrix between the ground truth and the newly obtained partition, showing in Figure 4.5(d) that we almost perfectly recover the structure of the ground truth bins.

Finally, we remark that our approach extends to higher-dimensional subspaces as well; e.g., if $S = \sum_{i=1}^l a_i X_i$, the top eigenvector will take constant, or almost constant,

Algorithm 2. One iteration of the CSSA for computing the conditional distribution $\mathbb{P}(F|S = s)$ of the fast variable F given a value s of the slow variable S .

- 1: Compute the propensity functions $\alpha_i(t)$, for $i = 1, 2, \dots, m$, and their sum $\alpha_0(t) = \sum_{i=1}^m \alpha_i(t)$.
 - 2: Generate r_1 and r_2 , two uniformly distributed random numbers in $(0, 1)$.
 - 3: Compute the next reaction time as $t + \tau$, where $\tau = -\log(r_1)/\alpha_0(t)$.
 - 4: Use r_2 to select reaction R_j which occurs at time $t + \tau$ (each reaction R_i , $i = 1, 2, \dots, m$, occurs with probability α_i/α_0).
 - 5: If the slow variable S changes its current state from s to $s' \neq s$ due to reaction R_j occurring, reset $S = s$ to its previous value while not changing the value of the fast variable F .
 - 6: If any of the variables X_i goes outside the boundary of the considered domain, then revert to the state of the system in step 4 before reaction R_j occurred.
-

values on the states $(x_1, x_2, \dots, x_l)^{(i)}$ corresponding to the slow state $S = s$, such that $\sum_{i=1}^l a_i x_i = s$. In other words, both the above-mentioned methodology that recovers the mapping and the bin-denoising scheme can be thought of in terms of *bins* and *urns* and are agnostic to the dimension of the system.

5. A Markov approach for computing the steady distribution of the slow variable. In this section, we focus on the final step of the ADM-CLE approach, of estimating the stationary distribution of the slow variable, without any prior knowledge of what the slow variable actually is. One of the ingredients needed along the way is an estimation of the conditional distribution $\mathbb{P}(F|S = s)$ of the fast variable F given a value s of the slow variable S , which we compute via two approaches. As the first approach, we consider the conditional stochastic simulation algorithm (CSSA) [8] which is given in Algorithm 2. It samples from the distribution of the fast variable conditioned on the slow variable. The second approach is entirely analytic and free of any stochastic simulations and amounts to analytically solving the CME for each set in the partition $\mathcal{P} = \{\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_k\}$. We then compare our results to the constrained multiscale algorithm (CMA) introduced in [8], which approximates the effective dynamics of the slow variable as an SDE, after estimating the effective drift and diffusion using the CSSA (Algorithm 2).

5.1. A stochastic simulation algorithm for estimating the conditional probability (CSSA). Our next task is to estimate the conditional distribution $\mathbb{P}(F|S = s)$ of the fast variable F given a value s of the slow variable S . One possible approach for doing this relies on the CSSA algorithm to globally integrate the effective dynamics of the slow variable. One iteration of the CSSA is given in Algorithm 2. Ideally, one repeats steps 1–6 of Algorithm 2 and samples values of F until the distribution $\mathbb{P}(F|S = s)$ converges. In practice, we run Algorithm 2 until L_c changes of the slow variable S occur. This computation is done for each value in the range of the slow variable $\mathcal{S} = \{s_1, s_2, \dots, s_{|\mathcal{S}|}\}$. As an example, Table 5.1 illustrates the computation of the transitioning probability from the slow variable state $S = 7$ to $S = 8$, associated to CS-II. We remark that this approach is in line with that of previous work [8], in which the running time of CSSA was implicitly defined by counting up to L_c changes in S . We consider the same range of values for L_c as in [8], i.e., $L_c = \{10, 10^2, 10^3, 10^4, 10^5\}$, as illustrated in Figure 5.5 below. An alternative approach would be to check for convergence in distribution by comparing the cumulative distribution functions at L_c and $2L_c$.

TABLE 5.1

Illustrative example CS-II. The set of all ground states of the system (x_1, x_2) corresponding to the slow variable $S = X_1 + 2X_2 = 7$. We denote by (x'_1, x'_2) the states reachable from (x_1, x_2) in one transition step and by S' the associated corresponding slow variable such that $S' = X'_1 + 2X'_2$. R_j denotes the reaction channel that takes the chemical system from state (x_1, x_2) to (x'_1, x'_2) , with corresponding propensity α_j . We highlight in bold letters the subset of all states via which the system can transition in one step from the slow variable $S = 7$ to $S = 8$.

X_1	X_2	S	X'_1	X'_2	S'	R_j	α_j/α_0	α_j
1	3	7	2	3	8	1	7.8×10^{-3}	96
1	3	7	0	3	6	2	7.8×10^{-5}	0.96
1	3	7	2	3	8	3	1.5×10^{-2}	184.38
1	3	7	0	3	6	4	1.6×10^{-3}	19.75
1	3	7	3	2	7	6	9.7×10^{-1}	12000
3	2	7	4	2	8	1	7.3×10^{-3}	64
3	2	7	2	2	6	2	2.0×10^{-4}	1.92
3	2	7	4	2	8	3	2.1×10^{-2}	184.38
3	2	7	2	2	6	4	6.7×10^{-3}	59.25
3	2	7	1	3	7	5	5.4×10^{-2}	480
3	2	7	5	1	7	6	9.1×10^{-1}	8000
5	1	7	6	1	8	1	5.4×10^{-3}	32
5	1	7	4	1	6	2	3.0×10^{-4}	1.6
5	1	7	6	1	8	3	3.1×10^{-2}	184.38
5	1	7	4	1	6	4	1.6×10^{-2}	98.75
5	1	7	3	2	7	5	2.7×10^{-1}	1600
5	1	7	7	0	7	6	6.7×10^{-1}	4000
7	0	7	8	0	8	3	5.0×10^{-2}	184.38
7	0	7	6	0	6	4	3.7×10^{-2}	138.25
7	0	7	5	1	7	5	9.1×10^{-1}	3360




FIG. 5.1. Transition diagrams for the two chemical systems: (a) CS-I and (b) CS-II.

5.2. An analytical derivation of the conditional distribution. An alternative approach which we follow in this paper relies on an analytical computation of the conditional distribution $\mathbb{P}(F|S = s)$, thus eliminating the need for any expensive stochastic simulations. We illustrate in Figure 5.1 the transition diagrams for the two chemical systems we consider in this paper. For chemical system CS-II, the system can transition from a given state (x_1, x_2) to four adjacent distinct \mathcal{O} -states: to $(x_1 - 2, x_2 + 1)$ via channel R_5 , to $(x_1 + 2, x_2 - 1)$ via channel R_6 , to $(x_1 - 1, x_2)$ via channels R_2 and R_4 , and finally to state $(x_1 + 1, x_2)$ via channels R_1 and R_3 . However, in terms of the underlying slow variable S , the system can transition to only two adjacent states $S = s - 1$ (via channels R_2 and R_4) and $S = s + 1$ (via

channels R_1 and R_3), or it can remain at the current state $S = s$ via channels R_5 and R_6 . Considering the subsystem of fast reactions of CS-II and conditioning on the line $s = x_1 + 2x_2$, the stationary CME takes the form

$$\begin{aligned} 0 &= k_5(x_1 + 2)(x_1 + 1) \mathbb{P}(X_1 = x_1 + 2, X_2 = x_2 - 1) \\ &\quad + k_6(x_2 + 1) \mathbb{P}(X_1 = x_1 - 2, X_2 = x_2 + 1) \\ &\quad - (k_5x_1(x_1 - 1) + k_6x_2) \mathbb{P}(X_1 = x_1, X_2 = x_2). \end{aligned}$$

Thus, the conditional distribution for CS-II is, for $0 \leq x_1 \leq s$, given by

$$\mathbb{P}(F = x_1 | S = s) = \frac{C}{x_1! x_2!} \left(\frac{k_5}{k_6}\right)^{x_2} = \frac{C}{x_1! ((s - x_1)/2)!} \left(\frac{k_5}{k_6}\right)^{(s-x_1)/2}$$

if $(s - x_1)$ is an even number.

Here, C is the normalization constant and $\mathbb{P}(F = x_1 | S = s) = 0$ if $(s - x_1)$ is an odd number.

A similar argument can be used for CS-I. The stationary CME of the fast subsystem of CS-I is written as

$$\begin{aligned} 0 &= k_2(x_1 + 1) \mathbb{P}(X_1 = x_1 + 1, X_2 = x_2 - 1) \\ &\quad + k_3(x_2 + 1) \mathbb{P}(X_1 = x_1 - 1, X_2 = x_2 + 1) \\ &\quad - (k_2x_1 + k_3x_2) \mathbb{P}(X_1 = x_1, X = x_2), \end{aligned}$$

where $s = (x_1 + x_2)/2$. Thus, the conditional distribution for CS-I is, for $0 \leq x_1 \leq 2s$, given by

$$\mathbb{P}(F = x_1 | S = s) = \frac{C}{x_1! x_2!} \left(\frac{k_2}{k_3}\right)^{x_2} = \frac{C}{x_1! (2s - x_1)!} \left(\frac{k_2}{k_3}\right)^{2s-x_1},$$

where C is again the normalization constant.

While in the above two examples we were able to compute analytical derivations, we remark that most often this is the case only for elementary systems. For more complex systems which are not amenable to analytical derivations, we would have to revert to the CSSA approach discussed in section 5.1.

5.3. Aggregated transition rates and a Markov chain on the state of slow variables. In the final step of the ADM-CLE approach, we set up a Markov chain on the state space of slow variables with the end goal of estimating the stationary distribution of the slow variable. As illustrated in Figure 5.1(b), the system CS-II can transition from a given state $S = s$ to two adjacent states $S = s - 1$ (via reaction channels R_2 and R_4) and $S = s + 1$ (via channels R_1 and R_3), or it can remain at the current state $S = s$ via channels R_5 and R_6 . Consider now the set $\mathcal{P}_s = \{\mathbf{x}^{(i)} = (x_1, x_2)^{(i)} | x_1 + 2x_2 = s\}$, illustrated as the middle bin in Figure 5.2. To compute the transition rate between two adjacent bins \mathcal{P}_s and \mathcal{P}_{s+1} , one has to aggregate over possible ways of getting from an observable state in bin \mathcal{P}_s to an observable state in bin \mathcal{P}_{s+1} . We compute $\Theta_1^{(s)}$ to be the aggregated transition rate from state \mathcal{P}_s to state \mathcal{P}_{s+1} , over all possible states $(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$, such that $\mathbf{x}^{(i)} \in \mathcal{P}_s$ and $\mathbf{x}^{(j)} \in \mathcal{P}_{s+1}$, by

$$(5.1) \quad \Theta_1^s = \sum_{\mathbf{x}^{(i)} \in \mathcal{P}_s} \sum_{\mathbf{x}^{(j)} \in \mathcal{P}_{s+1}} \mathbb{P}(F = x_1^{(i)} | S = s) \sum_{k=1}^m \alpha_k(\mathbf{x}^{(i)}) Q(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}, R_k),$$




FIG. 5.2. The Markov chain on the slow variable state space, using the aggregated transition probabilities (5.1)–(5.2) for the chemical system CS-II. Each slow variable state pools together multiple states $(x_1, x_2)^{(i)}$ of the original chemical system, which are placed in the same partition \mathcal{P}_s by our robust mapping procedure introduced in section 4.




FIG. 5.3. Plot of the aggregated transition rates for the illustrative example CS-II: (a) Θ_1 ; (b) Θ_2 ; and (c) $\Theta_1 - \Theta_2$.

where $Q(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}, R_k)$ denotes the indicator functions of whether one can transition from the \mathcal{O} -state $\mathbf{x}^{(i)}$ to \mathcal{O} -state $\mathbf{x}^{(j)}$ via reaction R_k . We define similarly the aggregated transition rate Θ_2^s , that the chemical system transitions from the slow variable state \mathcal{P}_s to \mathcal{P}_{s-1} by

$$(5.2) \quad \Theta_2^s = \sum_{\mathbf{x}^{(i)} \in \mathcal{P}_s} \sum_{\mathbf{x}^{(j)} \in \mathcal{P}_{s-1}} \mathbb{P}(F = x_1^{(i)} | S = s) \sum_{k=1}^m \alpha_k(\mathbf{x}^{(i)}) Q(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}, R_k).$$

We illustrate in Figure 5.3 the aggregated transition rates between the slow state $S = s$ and its adjacent states $S = s - 1$ and $S = s + 1$ for all values of the slow variable S . Note that in the derivations (5.1) and (5.2), we can either rely on the CSSA algorithm to sample from the conditional distribution of fast variables given values for the slow variables, as shown in section 5.1, or use the analytic formulation which is possible to derive for both CS-I and CS-II; see section 5.2.

Finally, we compute the solution to the stationary CME associated to the system in Figure 5.2 which can be written as

$$(5.3) \quad 0 = \Theta_1^{s-1} \pi(s-1) + \Theta_2^{s+1} \pi(s+1) - (\Theta_1^s + \Theta_2^s) \pi(s),$$

where $\pi(s) \approx \mathbb{P}(S = s)$ is the probability that $S = s$ at time t . Assuming that $\pi(s) = 0$ for all $s \notin \mathcal{S}$ and using no-flux boundary conditions, we arrive at a linear system. The eigenvector of the resulting matrix, with associated eigenvalue $\lambda = 0$, yields an approximate solution of the stationary CME, which we plot in blue in Figure 5.4. Our result is visually indistinguishable from the exact solution (plotted as the red solid line).

5.4. A comparison with the CMA. We compare the approach we introduced in section 5.3, relying on the analytical derivations in section 5.2, with the




FIG. 5.4. The final estimated stationary distribution of the slow variable S for the ADM-CLE approach, computed without knowledge of the slow variable, for (a) the chemical system CS-I; and (b) the chemical system CS-II (blue histograms). Red solid lines are exact solutions computed by solving the CME of the full model and using the corresponding definition of the slow variable.

TABLE 5.2

The distance (as measured by the error in (5.4)) between the estimated and the ground truth probability distributions of the slow variable, for the CMA algorithm which runs the CSSA algorithm for each value of the slow variable $S = s$, until L_c changes of the slow variable occur. The rightmost column shows the recovery errors for our proposed Markov-based approach.

	$L_c = 100$	$L_c = 1,000$	$L_c = 10,000$	$L_c = 100,000$	ADM-CLE
CS-I	0.40385	0.1225	0.055016	0.053468	0.003389
CS-II	0.68339	0.19916	0.12937	0.080234	0.028239

CMA method proposed in [8], which relies on the CSSA algorithm, and estimate the coefficients of the effective SDE via a proposed constrained multiscale algorithm. We compare the results of the two methods with the ground truth and record the error defined as

$$(5.4) \quad \text{Error}(\pi, \mathbb{P}(S = s)) = \sum_{s \in S} |\pi(s) - \mathbb{P}(S = s)|,$$

where $\mathbb{P}(S = s)$ denotes the ground truth probability distribution of the slow variable, and π denotes the estimated solution, either by the CMA and or the ADM-CLE. As Table 5.2 shows, the ADM-CLE approach yields lower errors compared to the CMA algorithm, even when we run the latter with the parameter L_c as large as 100,000. Note that for the chemical system CS-I, the ground truth probability distribution of the slow variable $\mathbb{P}(S = s)$ can be easily computed using the multivariate Poisson distribution, as discussed in section 2.1. For the second chemical system CS-II, we consider as ground truth the solution obtained by solving the associated CME of the full model in a large (truncated) domain.

In Table 5.2 we show numerical results that highlight the accuracy improvement of the ADM-CLE approach compared to the CMA approach of [8]. For the latter method, we run the CSSA algorithm [8] for each value of the slow variable, until L_c changes of the slow variable occur. As expected, the accuracy of the CMA algorithm improves as L_c increases, at the cost of additional computational running time of the method. In comparison, our stochastic simulation-free approach yields significantly more accurate results, with errors that are one order of magnitude lower than the CMA method with $L = 10,000$, and even $L = 100,000$. We plot in Figures 5.5 and 5.6 (top rows) the estimated stationary distribution of the CMA method for both chemical systems considered throughout this paper for several values of the L parameter. The bottom rows of Figures 5.5 and 5.6 show the estimated distribution, after smoothing out by the kernel density estimation (KDE).




FIG. 5.5. (a)–(d) The stationary distribution of the slow variable computed by the CMA for CS-I, using knowledge of the slow variable (blue histograms). The red solid line is the exact solution, $\mathbb{P}(S = s)$, obtained by solving the CME of the full system CS-I. The CMA approach runs the CSSA algorithm for each value of the slow variable $S = s$ until $L_c = \{10^2, 10^3, 10^4, 10^5\}$ changes of the slow variable occur. Panels (e)–(f) show the CMA-computed distribution after smoothing out by the KDE procedure.




FIG. 5.6. (a)–(d) The stationary distribution of the slow variable computed by the CMA for CS-II, using knowledge of the slow variable (blue histograms). The red solid line is the exact solution, $\mathbb{P}(S = s)$, obtained by solving the CME of the full system CS-II. The CMA approach runs the CSSA algorithm for each value of the slow variable $S = s$ until $L_c = \{10^2, 10^3, 10^4, 10^5\}$ changes of the slow variable occur. Panels (e)–(f) show the CMA-computed distribution after smoothing out by the KDE procedure.

5.5. Computational complexity. Computing the local covariance matrices can be done in linear time in number of states N in the system $\mathcal{O}(Nl^2)$, where l denotes the number of species. The main drawback of our approach is that the number of states is exponential in the number of species in the system. If r_i denotes the size of the range of the molecule X_i , the total number of states in the system

renders $N = \prod_{i=1}^l r_i$. If each molecule has roughly r states in its range, then $N = r^l$, which renders the current version of the ADM-CLE approach prohibitive for systems with large l .

From a computational point of view, all steps of the algorithm can be implemented in a distributed fashion and scaled linearly in the number of states N , except for the eigenvector computation, which is nearly linear. This is due to the fact that every iteration of the power method is linear in the number of edges of the graph G (which is very sparse), but the number of iterations is greater than $O(1)$, as it depends on the spectral gap.

6. Summary and discussion. In this paper we have introduced an ADM-CLE approach for detecting intrinsic slow variables in high-dimensional dynamic data, generated by stochastic dynamical systems. In the original ADM framework, the local bursts of simulations initiated at each data point to estimate the local covariances are computationally expensive, a shortcoming we avoid by using an approximation of the CLE. A second innovation that further improved the computational performance relates to the underlying similarity graph, a starting point for the diffusion map approach. By exploiting the spectrum of each local covariance matrix, we built a sparse ellipsoid-like neighborhood graph at each point in the data set, with the end result of being able to build a sparse similarity graph that requires the computation of a much smaller number of distances, which makes the ADM-CLE approach scalable to networks with thousands or even tens of thousands of nodes. For the two illustrative examples considered in this paper, the sizes of the resulting graphs are $N = 10,201$ for CS-I and $N = 12,100$ for CS-II, respectively. Had these graphs been complete graphs, the number of resulting weighted edges would be over 50 million, while in our computations, the number of edges is approximately 2.9 million for CS-I and 3.9 million for CS-II.

We have proposed a spectral-based method for inferring the slow variable present within the chemical system without any prior knowledge of its structure and a Markov-based approach for estimating its stationary distribution. We augment the proposed algorithmic approach with numerical simulations that confirm that the ADM-CLE approach can compare favorably for some systems to the CMA for estimating the stationary distribution of such slow variables. The ADM-CLE approach can also be applied to systems with a low number of states of slow variables. The CMA, as introduced in [8], is more suitable for systems where the slow variable(s) can take many different values, because the CMA uses an underlying SDE approximation for the behavior of the slow variables. One option to overcome this problem is to estimate effective propensity functions of the slow subsystem [9]. An open question is to extend the ADM-CLE to systems where the range of the X_i variables is very large. In the ADM-CLE approach applied to CS-I or CS-II, we associate a state (i.e., node in the initial graph) to each possible combination of pairs of states (x_1, x_2) , an approach no longer feasible whenever the range of the variables is large. To bypass this problem, one could change the discretization of the state space and modify accordingly the Markov chain-based approach (Figure 5.2) used in the ADM-CLE.

The main issue with the extension of the ADM-CLE to chemical systems with a higher number of species stems from the large computational running time due to the number of states growing exponentially in the number of species. In particular, the presented ADM-CLE approach which is based on standard matrix-vector computations cannot be applied to such systems. Exponential scaling can be avoided for some systems by storing the computed data in a tensor format, with smaller memory

requirements than in traditional approaches. The resulting tensor data can be analyzed using algebraic operations with computational complexity which scales linearly with dimension [31]. In principle, the tensor-based computations could be combined with the ADM-CLE approach to improve the scaling of its computational complexity with the dimension.

The ADM-CLE couples a method for finding slow variables (ADM) with an approach to compute the stationary distribution of a multiscale chemical reaction network. Chemical systems depend on a number of parameters (e.g., kinetic rate constants), and an open question is to extend the ADM approach to situations where one (or more) parameters are varied, i.e. to perform bifurcation analysis of multiscale stochastic chemical systems [31].

Finally, we point out recent work of Dsilva et al. [11], who also rely on the ADM framework to discover nonlinear intrinsic variables in high-dimensional data in the context of multiscale simulations, with the task of merging different simulation ensembles or partial observations of such ensembles in a globally consistent manner. Their work is motivated by the fact that often one is not merely interested in extracting the hidden (slow) variables from the underlying low-dimensional manifold, given partial observations $\mathbf{x}^{(i)}$, $i = 1, 2, \dots, N$, as in the ADM setting, but also in extending high-dimensional functions on a set of points lying in a low-dimensional manifold. Their proposed approach relies on the so-called Laplacian Pyramids [36], a multiscale algorithm for extending a high-dimensional function defined on a set of points in the space of intrinsic variables to a second set of points not in the data set, by using Laplacian kernels of decreasing bandwidths (the ε parameter in (3.5)).

REFERENCES

- [1] A. ARKIN, J. ROSS, AND H. MCADAMS, *Stochastic kinetic analysis of developmental pathway bifurcation in phage *l*-infected *Escherichia Coli* cells*, *Genetics*, 149 (1998), pp. 1633–1648.
- [2] M. BELKIN AND P. NIYOGI, *Laplacian eigenmaps for dimensionality reduction and data representation*, *Neural Comput.*, 15 (2003), pp. 1373–1396, <https://doi.org/10.1162/089976603321780317>.
- [3] Y. CAO, D. GILLESPIE, AND L. PETZOLD, *Multiscale stochastic simulation algorithm with stochastic partial equilibrium assumption for chemically reacting systems*, *J. Comput. Phys.*, 206 (2005), pp. 395–411.
- [4] Y. CAO, D. T. GILLESPIE, AND L. R. PETZOLD, *The slow-scale stochastic simulation algorithm*, *J. Chem. Phys.*, 122 (2005), 014116.
- [5] Y. CAO, H. LI, AND L. PETZOLD, *Efficient formulation of the stochastic simulation algorithm for chemically reacting systems*, *J. Chem. Phys.*, 121 (2004), pp. 4059–4067.
- [6] R. R. COIFMAN AND S. LAFON, *Diffusion maps*, *Appl. Comput. Harmon. Anal.*, 21 (2006), pp. 5–30.
- [7] R. R. COIFMAN, S. LAFON, A. B. LEE, M. MAGGIONI, B. NADLER, F. WARNER, AND S. W. ZUCKER, *Geometric diffusions as a tool for harmonic analysis and structure definition of data: Diffusion maps*, *Proc. Natl. Acad. Sci. USA*, 102 (2005), pp. 7426–7431.
- [8] S. COTTER, K. ZYGALAKIS, I. KEVREKIDIS, AND R. ERBAN, *A constrained approach to multiscale stochastic simulation of chemically reacting systems*, *J. Chem. Phys.*, 135 (2011), 094102.
- [9] S. L. COTTER AND R. ERBAN, *Error analysis of diffusion approximation methods for multiscale systems in reaction kinetics*, *SIAM J. Sci. Comput.*, 38 (2016), pp. B144–B163, <https://doi.org/10.1137/14100052X>.
- [10] D. L. DONOHO AND C. GRIMES, *Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data*, *Proc. Natl. Acad. Sci. USA*, 100 (2003), pp. 5591–5596.
- [11] C. J. DSILVA, R. TALMON, N. RABIN, R. COIFMAN, AND I. KEVREKIDIS, *Nonlinear intrinsic variables and state reconstruction in multiscale simulations*, *J. Chem. Phys.*, 139 (2013), 184109.
- [12] A. DUNCAN, R. ERBAN, AND K. ZYGALAKIS, *Hybrid framework for the simulation of stochastic chemical kinetics*, *J. Comput. Phys.*, 326 (2016), pp. 398–419, <https://doi.org/10.1016/j.jcp.2016.08.034>

- [13] A. DUNCAN, S. LIAO, T. VEJCHODSKY, R. ERBAN, AND R. GRIMA, *Noise-induced multistability in chemical systems: Discrete vs. continuum modeling*, Phys. Rev. E (3), 91 (2015), 042111.
- [14] W. E, D. LIU, AND E. VANDEN-EIJNDEN, *Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates*, J. Chem. Phys., 123 (2005), 194107.
- [15] R. ERBAN, I. KEVREKIDIS, AND H. OTHMER, *An equation-free computational approach for extracting population-level behavior from individual-based models of biological dispersal*, Phys. D, 215 (2006), pp. 1–24.
- [16] R. ERBAN, I. KEVREKIDIS, D. ADALSTEINSSON, AND T. ELSTON, *Gene regulatory networks: A coarse-grained, equation-free approach to multiscale computation*, J. Chem. Phys., 124 (2006), 084106.
- [17] R. ERBAN, T. FREWEN, X. WANG, T. ELSTON, R. COIFMAN, B. NADLER, AND I. KEVREKIDIS, *Variable-free exploration of stochastic models: A gene regulatory network example*, J. Chem. Phys., 126 (2007), 155103.
- [18] F. FOGEL, R. JENATTON, F. BACH, AND A. D’ASPROMONT, *Convex relaxations for permutation problems*, in NIPS: Proceedings of the 26th Annual Conference on Advances in Neural Information Processing Systems, C. J. C. Burges, L. Bottou, Z. Ghahramani, and K. Q. Weinberger, eds., 2013, pp. 1016–1024.
- [19] W. C. GEAR, J. M. HYMAN, P. G. KEVREKIDIS, I. G. KEVREKIDIS, O. RUNBORG, AND C. THEODOROPOULOS, *Equation-free, coarse-grained multiscale computation: Enabling macroscopic simulators to perform system-level analysis*, Commun. Math. Sci., 1 (2003), pp. 715–762.
- [20] M. GIBSON AND J. BRUCK, *Efficient exact stochastic simulation of chemical systems with many species and many channels*, J. Phys. Chem. A, 104 (2000), pp. 1876–1889.
- [21] D. T. GILLEPSIE, *The chemical Langevin equation*, J. Chem. Phys., 113 (2000), pp. 297–306.
- [22] D. T. GILLEPSIE, *A general method for numerically simulating the stochastic time evolution of coupled chemical reactions*, J. Comput. Phys., 22 (1976), pp. 403–434.
- [23] D. T. GILLEPSIE, *Exact stochastic simulation of coupled chemical reactions*, J. Phys. Chem., 81 (1977), pp. 2340–2361.
- [24] R. GRIMA, P. THOMAS, AND A. STRAUBE, *How accurate are the chemical Langevin and Fokker-Planck equations?*, J. Chem. Phys., 135 (2011), 084103.
- [25] M. HAATAJA, D. J. SROLOVITZ, AND I. G. KEVREKIDIS, *Apparent hysteresis in a driven system with self-organized drag*, Phys. Rev. Lett., 92 (2004), 160603.
- [26] E. L. HASELTINE AND J. B. RAWLINGS, *Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics*, J. Chem. Phys., 117 (2002), pp. 6959–6969.
- [27] T. JAHNKE AND W. HUISINGA, *Solving the chemical master equation for monomolecular reaction systems analytically*, J. Math. Biol., 54 (2007), pp. 1–26.
- [28] S. KAR, W. BAUMANN, M. PAUL, AND J. TYSON, *Exploring the roles of noise in the eukaryotic cell cycle*, Proc. Natl. Acad. Sci. USA, 16 (2009), pp. 6471–6476.
- [29] G. KLINGBEIL, R. ERBAN, M. GILES, AND P. K. MAINI, *STOCHSIMGPU: Parallel stochastic simulation for the Systems Biology Toolbox 2 for MATLAB*, Bioinformatics, 27 (2011), pp. 1170–1171.
- [30] H. W. KUHN, *The Hungarian method for the assignment problem*, Naval Res. Logist. Quart., 2 (1955), pp. 83–97.
- [31] S. LIAO, T. VEJCHODSKY, AND R. ERBAN, *Tensor methods for parameter estimation and bifurcation analysis of stochastic reaction networks*, J. R. Soc. Interface, 12 (2015), 20150233.
- [32] L. LOVÁSZ, *Random walks on graphs: A survey*, in Combinatorics, Paul Erdős Is Eighty, Vol. 2, János Bolyai Mathematical Society, Budapest, 1996, pp. 353–397.
- [33] M. MEILA AND J. SHI, *A random walks view of spectral segmentation*, in AI and STATISTICS (AISTATS), 2001.
- [34] M. NEWMAN, *Finding community structure in networks using the eigenvectors of matrices*, Phys. Rev. E (3), 74 (2006), 036104.
- [35] A. Y. NG, M. I. JORDAN, AND Y. WEISS, *On spectral clustering: Analysis and an algorithm*, in NIPS: Proceedings of the 14th Annual Conference on Advances in Neural Information Processing Systems, 2001, pp. 849–856.
- [36] N. RABIN AND R. R. COIFMAN, *Heterogeneous datasets representation and learning using diffusion maps and Laplacian pyramids*, in Proceedings of the 2012 SIAM International Conference on Data Mining, Anaheim, CA, 2012, pp. 189–199, <https://doi.org/10.1137/1.9781611972825.17>.
- [37] S. ROWEIS AND L. SAUL, *Nonlinear dimensionality reduction by local linear embedding*, Science, 290 (2000), pp. 2323–2326.
- [38] H. SALIS AND Y. KAZNESSIS, *Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions*, J. Chem. Phys., 122 (2005), 054103.

- [39] J. SHI AND J. MALIK, *Normalized cuts and image segmentation*, IEEE Trans. Pattern Anal. Mach. Intell., 22 (2000), pp. 888–905.
- [40] A. SINGER AND R. R. COIFMAN, *Non linear independent component analysis with diffusion maps*, Appl. Comput. Harmon. Anal., 25 (2008), pp. 226–239.
- [41] A. SINGER, R. ERBAN, I. G. KEVREKIDIS, AND R. R. COIFMAN, *Detecting intrinsic slow variables in stochastic dynamical systems by anisotropic diffusion maps*, Proc. Natl. Acad. Sci. USA, 106 (2009), pp. 16090–16095.
- [42] D. SPIELMAN AND S.-H. TENG, *Spectral partitioning works: Planar graphs and finite element meshes*, in Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science (FOCS '96), 1996, pp. 96–107.
- [43] D. SPIELMAN AND S.-H. TENG, *Spectral partitioning works: Planar graphs and finite element meshes*, Linear Algebra Appl., 421 (2007), pp. 284–305.
- [44] Y. SUN, W. LIN, AND R. ERBAN, *Time delay can facilitate coherence in self-driven interacting particle systems*, Phys. Rev. E (3), 90 (2014), 062708.
- [45] E. VANDEN-EIJNDEN, *Numerical techniques for multi-scale dynamical systems with stochastic effects*, Commun. Math. Sci., 1 (2003), pp. 385–391.
- [46] J. VILLAR, H. KUEH, N. BARKAI, AND S. LEIBLER, *Mechanisms of noise-resistance in genetic oscillators*, Proc. Natl. Acad. Sci. USA, 99 (2002), pp. 5988–5992.
- [47] U. VON LUXBURG, *A tutorial on spectral clustering*, Statist. and Comput., 17 (2007), pp. 395–416, <https://10.1007/s11222-007-9033-z>.
- [48] J. WILKIE AND Y. M. WONG, *Positivity preserving chemical Langevin equations*, Chem. Phys., 353 (2008), pp. 132–138.
- [49] X. ZHOU, M. BELKIN, AND N. SREBRO, *An iterated graph Laplacian approach for ranking on manifolds*, in Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '11), New York, 2011, pp. 877–885.