Coarse-graining molecular systems by spectral matching Ⓕ

Feliks Nüske ⓘD , Lorenzo Boninsegna, and Cecilia Clementi ⓘD

COLLECTIONS

Ⓕ This paper was selected as Featured

View Online   Export Citation   CrossMark

ARTICLES YOU MAY BE INTERESTED IN

Two methods developed for efficiently reproducing kinetic properties in molecular systems
Scilight **2019**, 310003 (2019); https://doi.org/10.1063/1.5121291

A hybrid, bottom-up, structurally accurate, Gō-like coarse-grained protein model
The Journal of Chemical Physics **151**, 044111 (2019); https://doi.org/10.1063/1.5108761

Unsupervised machine learning in atomistic simulations, between predictions and understanding
The Journal of Chemical Physics **150**, 150901 (2019); https://doi.org/10.1063/1.5091842

The Journal of Chemical Physics

Submit Today

The Emerging Investigators Special Collection and Awards
Recognizing the excellent work of early career researchers!

Coarse-graining molecular systems by spectral matching (F) (SCI)

Feliks Nüske,[a] (iD) Lorenzo Boninsegna, and Cecilia Clementi[b] (iD)

AFFILIATIONS

Center for Theoretical Biological Physics and Department of Chemistry, Rice University, Houston, Texas 77005-1892, USA

[a] feliks.nueske@rice.edu
[b] cecilia@rice.edu

ABSTRACT

Coarse-graining has become an area of tremendous importance within many different research fields. For molecular simulation, coarse-graining bears the promise of finding simplified models such that long-time simulations of large-scale systems become computationally tractable. While significant progress has been made in tuning thermodynamic properties of reduced models, it remains a key challenge to ensure that relevant kinetic properties are retained by coarse-grained dynamical systems. In this study, we focus on data-driven methods to preserve the rare-event kinetics of the original system and make use of their close connection to the low-lying spectrum of the system's generator. Building on work by Crommelin and Vanden-Eijnden [Multiscale Model. Simul. **9**, 1588 (2011)], we present a general framework, called spectral matching, which directly targets the generator's leading eigenvalue equations when learning parameters for coarse-grained models. We discuss different parametric models for effective dynamics and derive the resulting data-based regression problems. We show that spectral matching can be used to learn effective potentials which retain the slow dynamics but also to correct the dynamics induced by existing techniques, such as force matching.

Published under license by AIP Publishing. https://doi.org/10.1063/1.5100131

I. INTRODUCTION

Coarse-graining (or *model reduction*) is the process of describing a high-dimensional and complex dynamical system by a smaller set of variables and of providing a new set of governing equations for this reduced description. Coarse-graining has become a fundamental challenge in many different areas of science, such as finance, atmospheric science, or molecular biophysics. Two of the central reasons for the importance of coarse-graining are that first, analysis or numerical simulation of high-dimensional systems is often challenging or simply infeasible, and second, not all detailed features of the full system are needed in order to answer questions of scientific interest. For molecular systems, important contributions to the field include coarse-graining in structural space, where the physical representation of a system is simplified. Several approaches have been proposed to design coarse-grained models for large molecular systems that either reproduce structural features of fine-grained (atomistic) models (bottom-up)[1–6] or reproduce experimentally measured properties for one or a range of systems (top-down).[7–13]

An alternative approach is coarse-graining in configurational space, where a transformation of variables is applied to arrive at a smaller set of descriptors. Notable examples along these lines are the Mori–Zwanzig formalism,[14–18] conditional expectations,[19–21] averaging and homogenization,[22] Markov state models and related techniques,[23,24] and diffusion maps.[25]

The starting point in the design of a coarse molecular model is the definition of the variables. The choice of the coarse coordinates is usually made by replacing a group of atoms by one effective particle and is usually based on physical and chemical intuition. Because of the modularity of a protein backbone or a DNA molecule, popular models coarse-grain a macromolecule to a few interaction sites per residue or nucleotide, e.g., the C_α and C_β atoms for a protein.[26–28] Alternative methods have been proposed to design coarse variables more systematically.[29–31]

In this study, we are concerned with the inference of governing equations for the reduced set of variables, given that these descriptors have already been selected and that simulation data of the full system are available. Several methods to learn the parameters of

an effective dynamics from data of the full system have been proposed, most notably the force matching scheme[2,3] and the relative entropy method[4] (the two approaches are connected[32]). Most of the previous work aimed at recovering correct thermodynamics of the reduced system, that is, the distribution sampled by the effective dynamics should equal the distribution of the projected original process. However, these methods do not determine the equations for a system's dynamical evolution (of course there are exceptions, such as the adaptive resolution approach; see Ref. 33). This means that most methods may not be able to design coarse-grained models that can reproduce molecular dynamical mechanisms (e.g., large conformational changes or assembly mechanisms in protein systems). Here, we shift the focus to designing coarse-grained models that reproduce slow dynamical mechanisms of a fine-grained system, that is, time scales, metastable states, and transitions in between them.

In principle, if the dynamical equations of the fine-grained model are known, the dynamics of the corresponding coarse-grained variables is given by the Mori-Zwanzig projection formalism,[14–17] which introduces a memory term. Even if the memory term can be simplified with the assumption of a separation of time scales, the estimation of the quantities involved in the Mori-Zwanzig approach is nontrivial.[18] However, here we are not interested in reproducing all time scales of the system but just the slowest processes (e.g., conformational changes or assembly processes in protein systems). This fact allows us to bypass the Mori-Zwanzig approach to directly define an effective coarse-grained potential to satisfy this requirement. Instead of focusing on the learning of dynamical equations (see, e.g., Refs. 34 and 35), here we build on a framework for parameter estimation of stochastic dynamics introduced in Refs. 36 and 37. It exploits the fact that slow dynamical processes are directly related to low-lying eigenvalues and associated eigenfunctions of the generator of the dynamics.[38,39] A broad range of methods is available to approximate these spectral components from simulation data of the full system.[23,40–44] Hence, we argue that the same loss function can be used to learn coarse-grained dynamics if the focus is on reproducing the slow kinetics and call the resulting framework the *spectral matching* estimator. We derive the optimization problems for two specific use cases of spectral matching: the first is to recover an effective potential within an overdamped dynamics, and the second is to correct dynamics obtained from force matching by learning a position-dependent diffusion. Applications to several toy systems and molecular dynamics simulations of alanine dipeptide illustrate the capabilities of the method and highlight practical details, especially the importance of regularization.

II. THEORY

A. Full space dynamics

We consider a stochastic process X_t attaining values in d-dimensional space $\mathbb{R}^d$, where d is typically large. In the case of a molecular system, X_t represents the coordinates of the molecule at time t. Even though the methods below are also applicable in more general contexts, we assume for the purpose of illustration that the process solves a reversible stochastic differential equation (SDE),

$$dX_t = [-A(X_t)\nabla_x F(X_t) + \nabla_x \cdot A(X_t)]dt + \sqrt{2}\Sigma(X_t)dW_t. \quad (1)$$

In this formulation, F is a scalar potential, while $A = \Sigma\Sigma^T$ is the diffusion matrix, and W_t is Brownian motion. Equation (1) is the general form of a reversible SDE, where reversibility holds with respect to the invariant density of X_t, given by $\mu(x) \propto \exp(-F(x))$. A widely used example of Eq. (1) is the overdamped Langevin dynamics defined by a potential energy function V and a constant diffusion $a > 0$, which is related to temperature T and friction γ by $a = \frac{k_B T}{\gamma}$,

$$dX_t = -a\nabla_x V(X_t)dt + \sqrt{2a}\,dW_t. \quad (2)$$

The generator associated with Eq. (1) is the second order differential operator

$$\mathcal{L}f = [-A\nabla_x F + \nabla_x \cdot A] \cdot \nabla_x f + A : \nabla_x^2 f, \quad (3)$$

with the colon indicating the Frobenius inner product between matrices. In most practically relevant cases, the operator $-\mathcal{L}$ possesses a discrete set of increasing eigenvalues $0 = \kappa_0 < \kappa_1 < \cdots$, with associated eigenfunctions ψ_i, $i = 0, 1, \ldots$. Each eigenvalue κ_i corresponds to the relaxation rate of a dynamical process in the system,[45] with *characteristic (implied) time scale*

$$t_i = \frac{1}{\kappa_i}. \quad (4)$$

We are particularly interested in metastable processes; in terms of the eigenspectrum of (3), it means that there is a cluster of M eigenvalues $\kappa_1, \ldots, \kappa_M$ close to $\kappa_0 = 0$, separated from all higher eigenvalues by a spectral gap.[38–40,46] We can see from Eq. (4) that these low-lying eigenvalues correspond to slow relaxation processes. We thus assume that there is a separation of time scales between fast and slow processes in the molecular system of interest.

B. Reduced dynamics

In this study, we consider coarse-graining of system (1) by projecting the state space $\mathbb{R}^d$ into a lower-dimensional space $\mathbb{R}^m$, $m \leq d$. Following Refs. 19 and 47, this projection is realized by a *coarse-graining map* $\xi : \mathbb{R}^d \mapsto \mathbb{R}^m$, $x \mapsto z = \xi(x)$. Our objective is to replace (1) by another reversible SDE defined only on the lower-dimensional space $\mathbb{R}^m$, with corresponding scalar potential F^ξ and diffusion A^ξ. These parameters should meet the following two requirements:

1. **Thermodynamic consistency**: The coarse-grained invariant density $\nu = \exp(-F^\xi)$ should be related to the full state invariant density by averaging, which we write as[19,47,48]

$$\nu(z) = \int \mu(x)\delta(\xi(x) - z)\,dx. \quad (5)$$

The corresponding scalar potential is then called the *potential of mean force (pmf)*,

$$F_{mf}^\xi(z) = -\log \nu(z). \quad (6)$$

2. **Kinetic consistency**: The coarse-grained dynamics should retain the metastable part of the original dynamics (that is, the slow processes). If $\mathcal{L}^\xi$ is the reduced generator analogous to (3), its leading eigenvalues $\kappa_1^\xi, \ldots, \kappa_M^\xi$ and eigenfunctions $\psi_1^\xi, \ldots, \psi_M^\xi$ should match the corresponding eigenvalues and eigenfunctions of the original system, that is, they should satisfy $\kappa_i^\xi \approx \kappa_i$ and $\psi_i^\xi \approx \psi_i$ (in an appropriate sense, please see the discussion in Sec. II D).

C. Force matching

The thermodynamic consistency can be enforced by *force matching*,[2,3] a powerful technique to extract the potential of mean force from simulation data of the full system. It is based on the fact that the gradient of the pmf solves the following minimization problem:[3,48]

$$\nabla_z F_{mf}^\xi = \mathrm{argmin}_f \int_{\mathbb{R}^d} \|\nabla_z f(\xi(x)) - F_{lmf}^\xi(x)\|^2 \, d\mu(x), \qquad (7)$$

where the minimization is over all suitable potential functions f and F_{lmf}^ξ is the *local mean force*.[48] The integral in (7) can be replaced by a data-based regression if sufficient simulation data are available. One of the difficulties in the practical application of this method has been that, in general, a coarse-grained potential satisfying the thermodynamic consistency includes many-body terms that are not easily modeled in the energy functions. Recently, machine learning methods have been used to alleviate this problem.[49–52]

D. Spectral matching

In order to achieve kinetic consistency, we build on an idea presented in Refs. 36 and 37 for parameter estimation in stochastic dynamics like Eq. (1). The goal is to identify the parameters $\theta \in \mathbb{R}^N$ of a parametric model $\mathcal{L}_\theta^\xi$ of the reduced generator. Given the leading eigenpairs κ_i, ψ_i, or an approximation of those, the idea is to minimize the discrepancy in the leading eigenvalue equations with respect to the model parameters. We probe these equations against a set of test functions f_j on the coarse-grained space. The loss function to be minimized is thus

$$\theta^* = \mathrm{argmin}_\theta \frac{1}{2} \sum_{i=1}^M \sum_{j=1}^P \langle \mathcal{L}_\theta^\xi \psi_i + \kappa_i \psi_i, f_j \rangle_\nu^2. \qquad (8)$$

Alternatively, we can also have the model generator act on the test functions in (8). These two setups are identical if the model dynamics are reversible with respect to the measure ν. It is important to note that we assume the eigenfunctions ψ_i to be functions of the reduced variables z. It follows then that the inner product in (8) can be approximated using simulation data of the full process.[47] For a justification of the assumption that all ψ_i are functions of z, please see Refs. 20 and 53.

III. METHODS

In this section, we discuss two specific use-cases for the spectral matching estimator as well as details of the practical implementation of the method.

A. Estimation of a scalar potential

The first use-case arises from the assumption that the effective dynamics can be modeled using overdamped Langevin dynamics in a scalar potential F_θ^ξ at a fixed constant diffusion $a > 0$. The resulting model generator is

$$\mathcal{L}_\theta^\xi f = -a\nabla_z F_\theta^\xi \cdot \nabla_z f + a\Delta_z f. \qquad (9)$$

In particular, if the model is a linear expansion into given basis functions, i.e., $F_\theta^\xi = \sum_{n=1}^N w_n g_n$, $\theta = (w_1, \ldots, w_N) \in \mathbb{R}^N$, spectral matching becomes a linear regression. As the eigenfunctions are generally constant along metastable sets but can display strong variations in poorly sampled transition regions, the numerical calculation of eigenfunction derivatives tends to be unstable. Hence, we apply the model generator to the test functions in (8). We also include the zeroth spectral pair, as the corresponding matrix entries are nonzero, to obtain a regression matrix $X \in \mathbb{R}^{(M+1)P \times N}$ and a data vector $y \in \mathbb{R}^{(M+1)P}$,

$$E(w) = \frac{1}{2}\|Xw - y\|^2, \qquad (10)$$

$$X_{i,j;n} = -\langle \psi_i, a\nabla_z g_n \cdot \nabla_z f_j \rangle_\nu, \qquad (11)$$

$$y_{i,j} = -\langle \psi_i, a\Delta_z f_j + \kappa_i f_j \rangle_\nu. \qquad (12)$$

We have found that regression (10) is often ill-conditioned and requires regularization. Below, we will use elastic net regularization[54] with parameters α, ρ, that is,

$$E^\alpha(w) = \frac{1}{2}\|Xw - y\|^2 + \alpha\rho\|w\|_1 + \frac{1}{2}\alpha(1-\rho)\|w\|_2^2. \qquad (13)$$

B. Estimation of a diffusion

Next, we consider the situation where an estimate of a thermodynamically consistent scalar potential F_{mf}^ξ is already available, for example, by application of the force matching scheme. As a result, any parametric model A_θ^ξ for the diffusion results in a symmetric model generator, and the parameter dependent term in the loss function can be evaluated as

$$\langle \mathcal{L}_\theta^\xi \psi_i, f_j \rangle_\nu = -\int A_\theta^\xi \cdot \nabla_z \psi_i \cdot \nabla_z f_j \, d\nu, \qquad (14)$$

provided that A_θ^ξ is always symmetric. This formulation only requires first order derivatives, but we need to estimate those derivatives for the eigenfunctions. As a special case, we focus on a linear expansion into scalar multiples of the identity matrix, i.e., $A_\theta^\xi = \left[\sum_{n=1}^N w_n g_n\right]\mathrm{Id}$, $\theta = (w_1, \ldots, w_N) \in \mathbb{R}^N$. Again, spectral matching leads to a linear regression,

$$w^* = \mathrm{argmin}_w \frac{1}{2}\|Xw - y\|^2, \qquad (15)$$

$$X_{i,j;n} = -\langle g_n \nabla_z \psi_i \cdot \nabla_z f_j \rangle_\nu, \qquad (16)$$

$$y_{i,j} = -\kappa_i \langle \psi_i, f_j \rangle_\nu. \qquad (17)$$

Positive definiteness of the diffusion is easily enforced in this setting by requiring positivity of the scalar prefactor at all data points.

In fact, we will restrict the diffusion to satisfy preselected upper and lower bounds $0 \leq a_{min}(z) \leq a_{max}(z)$, which can in principle be position dependent. If K data points $Z_{t_1} = \xi(X_{t_1}), \ldots, Z_{t_K} = \xi(X_{t_K})$ are given, we add linear inequality constraints

$$a_{max}(Z_{t_k}) \geq \sum_{n=1}^{N} w_n g_n(Z_{t_k}) \geq a_{min}(Z_{t_k}) \geq 0 \qquad (18)$$

for all $k = 1, \ldots, K$. The full optimization problem (15) then turns into a quadratic programming problem.

C. Variational approach

The spectral matching approach discussed above requires an estimate of the eigenfunctions and eigenvalues of the original system. In order to compute approximations of a system's dominant eigenfunctions ψ_i and corresponding eigenvalues κ_i and also to analyze simulation data of the effective dynamics, we make use of the *Variational Approach to Conformational Dynamics* (VAC),[41,43] which is a data-driven method to represent the dominant eigenfunctions from a given library of basis functions. In most of the examples below, the basis will consist of piecewise constant functions, which is also referred to as a *Markov state model* (MSM).[23,40,42,55] In a few cases, we will also use a basis of Gaussian functions, for example, if we also require derivatives of the slow eigenfunctions. VAC models are validated using the implied time scale test as described in Ref. 23. In addition, we employ Bayesian Markov state models[56] in order to obtain errorbars for all estimates of implied time scales. In order to extract a decomposition of state space into metastable sets, based on the dominant eigenfunctions, we use the PCCA algorithm.[57,58] The software implementation we use is the pyEmma package.[59]

IV. RESULTS

A. Three well potential

We first illustrate the idea of spectral matching to model the overdamped Langevin dynamics [Eq. (9)] in a two-dimensional toy potential. The energy function is given as a sum of three Gaussians and a harmonic confinement, that is,

$$V(x,y) = \sum_{q=1}^{3} a_q \exp\left(-\frac{1}{2s_q^2}\left[(x - m_q^x)^2 + (y - m_q^y)^2\right]\right) + a_4 x^2 + a_5 y^2. \qquad (19)$$

The actual values of the parameters are $(a_1, a_2, a_3, a_4, a_5) = (-3, -5, -4, 0.1, 0.2)$, $(m_1^x, m_1^y) = (0,0)$, $(m_2^x, m_2^y) = (1,2)$, $(m_3^x, m_3^y) = (-4,-1)$, and $(s_1, s_2, s_3) = (0.5, \sqrt{5/6}, 0.5)$; a contour of the energy is shown in Fig. 1(a). The diffusion constant is $a = 1$. We generate a long equilibrium simulation of these dynamics at integration time step $\Delta_t = 10^{-4}$, spanning $K = 10^7$ time steps.

We use a Markov state model based on a regular grid discretization of the two-dimensional state space to compute the slow eigenfunctions and eigenvalues of the system. The slowest process corresponds to the transition out of the shallow minimum on the left into the center, and it occurs at a time scale of $t_1 \approx 13$. The rest of the spectrum is separated from this process, and there is no other distinct slow motion to be identified; see the black dots in Fig. 1(c). We thus set $M = 1$ (that is, we use only the first nontrivial eigenfunction) and extract the approximate eigenfunction ψ_1 from the MSM.

We attempt to reproduce the energy V in two-dimensional space from data while not applying any model reduction. We employ spectral matching and a linear model in combination with the elastic
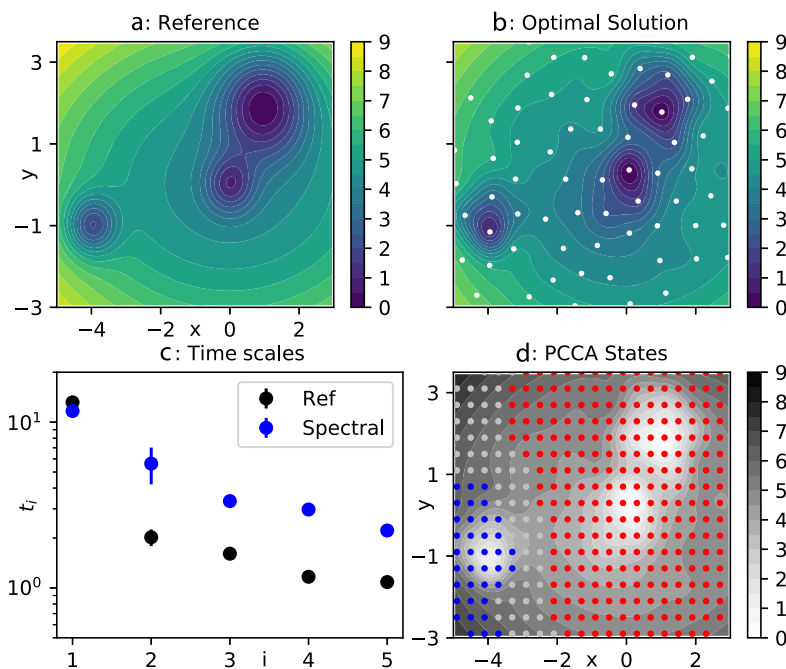

a: Reference   b: Optimal Solution
c: Time scales   d: PCCA States

FIG. 1. Application of spectral matching with a linear model, cf. Eq. (13), to a two-dimensional test system. (a) The reference potential Eq. (19). (b) Optimized potential found by spectral matching. White dots indicate the locations of 82 centers used to define the Gaussian basis and test functions. (c) Comparison of the first five implied time scales of the reference dynamics (black) and the corresponding time scales of the learned potential (blue). The spectral matching was designed to reproduce the slowest time scale, as indeed appears in the figure. (d) Metastable decomposition obtained by applying 2-state PCCA to an MSM of the learned dynamics. States are assigned to a macrostate if their membership is larger than 0.6. Silver dots represent transition states that cannot be assigned in this way.

net as in Eq. (13). A regular space clustering of the data at minimum distance 0.8 helps us define a set of 82 spherical Gaussian functions centered at the cluster centers, as shown in Fig. 1(b). These centers serve to define the test set, while the basis set consists of the same Gaussians plus the quadratic functions x^2, y^2 ($P = 82$, $N = 84$). We use uniform widths σ_{base} and σ_{test} for all Gaussians in the basis set and the test set, respectively. We fix $\sigma_{test} = 0.8$, while σ_{base} and both elastic net parameters α, ρ are treated as hyperparameters. They are determined by 3-fold cross-validation (CV) for all triples of the form $(\sigma_{base}, \alpha, \rho) \in \{0.2, 0.4, 0.6, 0.8, 1.0\} \times \{10^{-6}, 10^{-4}, 10^{-2}, 10^0\} \times \{0.5, 0.8, 0.9, 0.95, 1.0\}$. The optimal parameter set is $(\sigma_{base}, \alpha, \rho) = (0.4, 10^{-4}, 0.5)$; the corresponding solution is displayed in Fig. 1(b) and agrees well with the original energy landscape in panel (a).

Using the same simulation parameters as for the original system, we generate a realization of the dynamics defined by the optimized potential. An MSM analysis confirms that the slowest time scale is indeed well reproduced by those dynamics; see Fig. 1(c). As only the first nontrivial eigenfunction was used in the spectral matching, we do not expect to reproduce additional time scales besides the slowest one. Moreover, application of 2-state PCCA to this Markov model also shows that the decomposition of state space into metastable sets is the same as for the original potential, as displayed in panel (d) of Fig. 1.

B. Three well potential with roughness

Here, the performance of the spectral matching is tested on a perturbed version V_{per} of the three well potential investigated in Sec. IV A to which 100 small amplitude Gaussians were added, i.e.,

$$V_{per}(x,y) = V(x,y) + \sum_{q=1}^{100} w_q \exp\left(-\frac{1}{2\sigma^2}\left[(x - x_q^0)^2 + (y - y_q^0)^2\right]\right),$$

(20)

with $\sigma = 0.2$ and (x_q^0, y_q^0) indicating the center of the qth perturbing Gaussian; a contour plot of such a potential is shown in Fig. 2(a). The diffusion constant is $a = 1$. We generate a long equilibrium simulation of these dynamics at integration time step $\Delta_t = 10^{-4}$, spanning $K = 6 \times 10^7$ time steps.

The dominant time scales t_i and eigenfunctions ψ_i were numerically approximated using the VAC and a basis of 196 Gaussian features.

Spectral matching was applied to the data set, as exactly the same functions entering the linear combination Eq. (20) were employed both as basis and test functions in the procedure. Eventually, the full set of 105 coefficients $\{a_i\}_1^5 \bigcup \{w_q\}_1^{100}$ is approximated and compared with the exact result. Regression was solved using standard implementations, with and without elastic net regularization. The first two nontrivial eigenpairs were used in the spectral matching [that is, $M = 2$ in Eq. (8)].

The nonregularized solution is shown in Fig. 2(b) and almost perfectly reproduces the exact potential, panel (a). The regularized solution is shown in panel (c), and the optimal regularization hyperparameters ρ, $\alpha = (0.5, 2 \times 10^{-4})$ were identified by running 3-fold CV on all tuples of the form $(\rho, \alpha) \in \{0.4, 0.5, 0.6\} \times \{10^{-4}, 2 \times 10^{-4}, 3 \times 10^{-4}\}$. The solution is very sparse; only 17 coefficients (out of the full set of 105) have nonzero values. The position of the three main energy wells is correctly identified, which correlate with the system's slowest dynamics, while most of the roughness of the potential is washed out. Regularization appears to be a physically meaningful
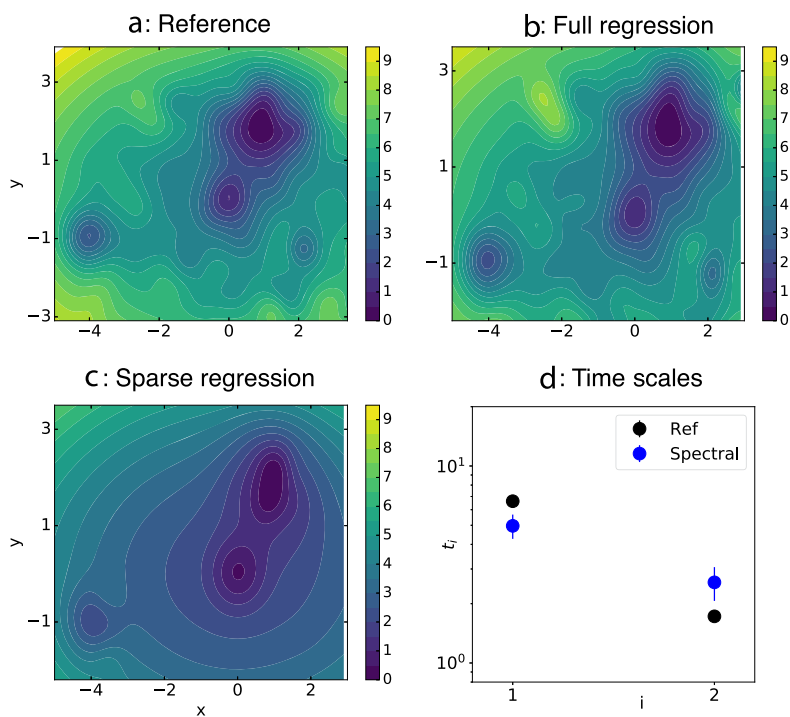

FIG. 2. Application of spectral matching with a linear model, cf. Eq. (13), to a rugged potential. (a) The reference potential is shown and was built by adding 100 small amplitude Gaussians to Eq. (19), which are responsible for the ripples; compare with Fig. 1(a). (b) Resulting potential obtained from running standard linear regression. (c) Optimal solution obtained after optimizing regularizing hyperparameters ρ, α by cross-validation. (d) First two implied time scales t_1 and t_2 estimated from simulations of the reference (black markers) and the learned potential (blue markers).

procedure since it systematically filters perturbations, i.e., rugged local minima, out.

However, the wells are shallower than expected; compare with Fig. 2(a). In order to investigate if this observation has any effect, a long equilibrium trajectory was generated by integrating overdamped Langevin dynamics, using the optimal solution as potential energy ($\Delta_t = 10^{-4}$, $a = 1$), and the associated time scales were estimated using VAC and a Bayesian Markov model in the same setup as before. Results are shown in Fig. 2(d), where the resulting two slowest time scales (blue markers) are compared to the reference ones (black markers): both t_1 and t_2 are approximately retained.

C. Recovery of slow kinetics by nonconstant diffusion

Next, we demonstrate that spectral matching can be used to recover the slow kinetics of a projected system with the aid of a nonconstant diffusion. We consider a modified three-well system [Fig. 3(a)] where the locations of the minima have been changed and additional peaks have been added to the landscape. The slowest process is now the transition out of the shallow minimum on the right. An equilibrium simulation of 10×10^6 frames at integration time step $\Delta_t = 10^{-3}$ serves as the reference data set.

We consider the reduced dynamics along the first coordinate $\xi(x, y) = x$. Due to the perturbations of the landscape, the projection onto x averages both high- and low-energy regions, which significantly affects the potential of mean force along x. We apply force matching using a basis of 71 Gaussian functions centered at grid spacing 0.2 between $x = -7.0$ and $x = 7.0$. A uniform width is selected by means of CV from the set $\sigma \in \{0.1, 0.2, 0.4, 0.6, 0.8\}$. The estimated potential of mean force for the optimal value $\sigma = 0.2$ is represented by the black line in Fig. 3(b).

We generate simulation data of overdamped Langevin dynamics in the potential of mean force at diffusion $a = 1$ (using the same integration time step and number of frames as for the full system). An MSM analysis of these data shows that the slowest time scale is decreased by a factor ten, while the next time scale is almost unaffected by the projection; see Fig. 3(d). Thus, the order of the two time scales is reversed, and the slow kinetics cannot be recovered by simply rescaling the effective diffusion constant a.

Spectral matching in form (15) is applied to estimate a position-dependent diffusion. Approximate rates κ_i, $i = 1, 2$ are extracted from a Markov state model of the full process. As the approximate eigenfunctions ψ_i need to be functions of z in Eq. (15), we build another MSM along x alone and extract those approximate eigenfunctions. A spline interpolation is used to ensure differentiability of the eigenfunctions as required by (15). Even though this MSM does not yield accurate estimates of both time scales, the corresponding eigenfunctions still capture the structure of both slow transitions correctly, as we can see by looking at the corresponding PCCA memberships [solid lines in Fig. 3(c)]. The test set is chosen as the same basis set used for force matching. As a basis set, we choose a set of N piecewise constant functions, where the corresponding discrete sets are obtained by partitioning the interval $[-6, 6]$ into N equal-sized parts. The optimal choice $N = 10$ is again obtained by 3-fold cross-validation on the hyperparameter set $N \in \{5, 10, 20, 30, 40\}$.

It is important to solve Eq. (15) subject to position dependent lower bounds a_{min}. Based on our analysis of the full data and the force matching simulation, we require $a_{min} = 0.9$, that is, close to one, in the left metastable set, while we set $a_{min} = 0.1$ everywhere else. We also use $a_{max} = 1.0$ as a global upper bound. The resulting optimal diffusion is indicated by the blue line in Fig. 3(b). This solution tends to slow down the transition out of the rightmost metastable
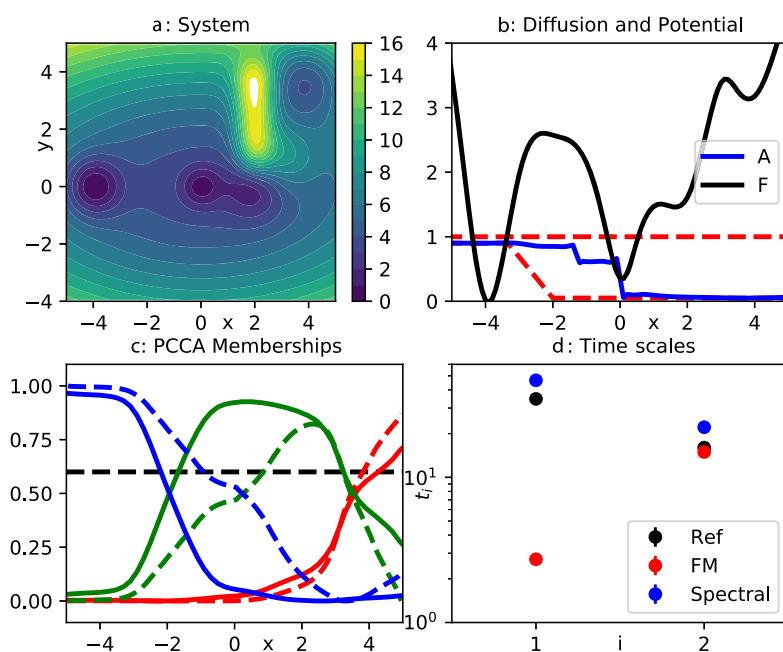

FIG. 3. Results for suboptimal projection of a two-dimensional toy potential. (a) Full potential. (b) Estimates of the potential of mean force (black curve) and position-dependent diffusion (blue curve). Red dashed lines correspond to the lower and upper bounds a_{min}, a_{max} for the diffusion. (c) Metastable memberships extracted by running three-state PCCA on Markov state models of the full dynamics after projection onto x (solid lines) and of the effective simulations (dashed lines). The black line at membership equal to 0.6 is given as a visual aid to determine metastable sets. (d) Comparison of the slowest two time scales obtained from the full simulation (black dots), simulations of the force matching potential (red dots), and simulations of the combined effective dynamics using the force matching potential and the optimized diffusion (blue dots). Note that we have exchanged the order of the force matching time scales to ensure all time scales for index i correspond to the same slow process.

state, while leaving the time scale of transition out of the center state unchanged. After running the effective dynamics for 10×10^6 steps at $\Delta_t = 10^{-3}$ and analyzing these data by an MSM, we find that both time scales are indeed approximately restored; see Fig. 3(d). These effective dynamics are more diffusive than the original system, as is reflected in the corresponding PCCA memberships being less crisp [dashed lines in Fig. 3(c)]. The change in the dynamics to restore the correct time scales perturbs the position of the metastable sets. However, the three metastable sets of the original dynamics are still approximately retained.

D. Alanine dipeptide

Finally, we study model reduction of a small molecular system, alanine dipeptide, which has served as a test case for numerous studies in recent years. The data set at our disposal is the same that was used in Ref. 51; please see Ref. 60 for the detailed simulation setup. It comprises one million frames of Langevin dynamics in explicit water saved every 1 ps.

It is well known that the system's metastable processes are effective functions of its two backbone dihedral angles ϕ, ψ. Therefore, we study reduction of this system into the two-dimensional space of those dihedral angles. Figure 4(a) shows the empirical free energy in this space. It presents a more challenging example than the toy systems, as the coarse-graining map is nonlinear, and the metastability is much more pronounced in the sense that there are large unsampled areas and we only have very few samples in the transition regions. To simplify matters, we shift both dihedrals in order to eliminate almost all periodicity from this representation and use nonperiodic basis functions in what follows.

To obtain the approximate eigenfunctions ψ_i and eigenvalues κ_i, we apply the VAC with 225 spherical Gaussian functions centered

on a regular grid between -2.8 and $+2.8$ at a distance of 0.4 in each direction. Their widths are uniformly set to 0.3. We retain the first three nontrivial slow eigenfunctions, and the corresponding implied time scales are $t_1 \approx 1.2$ ns, $t_2 \approx 63$ ps, and $t_3 \approx 36$ ps. The corresponding transitions in dihedral space occur, respectively, between the left and right half of the plane, between the two minima on the left, and between the two shallow minima on the right.

We use spherical Gaussian functions to represent the force matching potential and to serve as basis and test functions in the spectral matching. The centers are placed on a regular grid with grid spacing 0.4 in all directions, but we remove centers in unpopulated areas of the state space. As a result, a set of 152 centers for Gaussian basis functions is obtained. For force matching, we employ 3-fold cross-validation to determine a uniform width out of the parameter set $\sigma \in \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8\}$ to yield an optimal value $\sigma = 0.6$. To ensure that the resulting potential is confining, we evaluate the solution on a fine grid and replace its values at unsampled grid points by a function that grows quadratically with the distance to the sampled area. A two-dimensional spline is fitted to these data to yield the final approximation.

To evaluate the kinetic properties of the force matching model, we generate a set of 100 overdamped Langevin simulations of 15 ns length at constant diffusion $a = 1$ and estimate an MSM from these data. By comparing the three slowest time scales of these dynamics to the reference values [red and black dots in Fig. 4(d)], we find that the force matching dynamics are uniformly accelerating the kinetics by about a factor three.

For spectral matching, we use the optimal Gaussian functions from force matching to serve as test functions and also use the same centers to define the basis set. These centers are indicated by red markers in Fig. 4(b). We globally set $a_{min} = 0.3$, $a_{max} = 1.0$ and apply 3-fold CV to determine the widths of the basis set, where $\sigma_{base} \in \{0.4,$
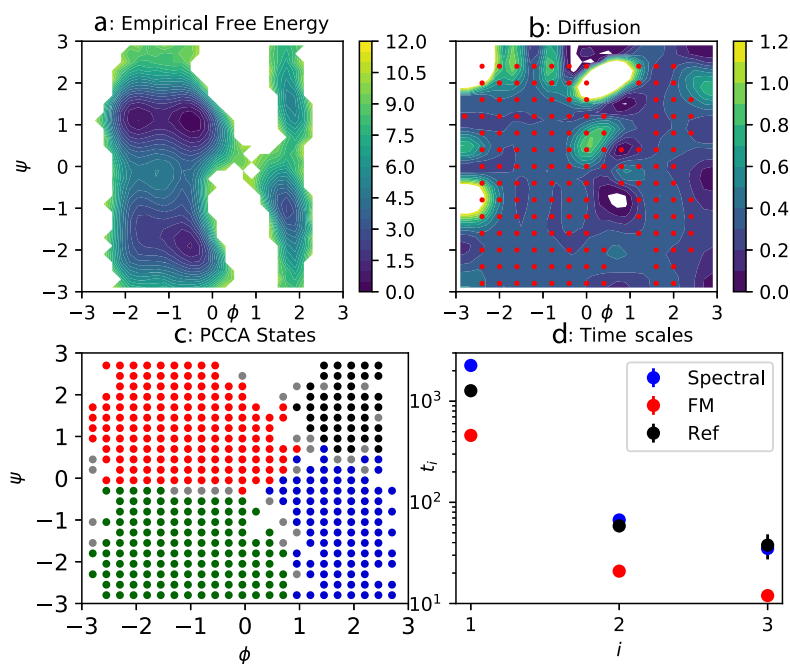

FIG. 4. Results for coarse-graining of alanine dipeptide in the space of its backbone dihedrals ϕ, ψ. (a) Empirical free energy in dihedral space from original MD simulation. (b) Optimized diffusion. Red markers indicate centers of the Gaussian functions defining the basis and the test set. (c) Metastable states as identified by four-state PCCA on an MSM of the effective dynamics. A state is assigned to a PCCA state if its degree of membership is larger than 0.6. Gray dots represent transition states that cannot be assigned to a macrostate in this way. (d) Comparison of the three slowest implied time scales obtained from the reference dynamics (black dots), simulations of the force matching potential (red dots), and simulations combining the optimal diffusion from (b) and the force matching potential (blue dots).

0.45, 0.5, 0.55, 0.6}. The solution of Eq. (15) for the optimal parameter value $\sigma_{base} = 0.55$ is shown in Fig. 4(b). The resulting diffusion is mostly constant at a value close to 0.3 throughout most of the domain, except for some peaks in boundary or unsampled regions. By generating another set of 100 simulations of 15 ns simulation time each, we confirm by an MSM that all three time scales are corrected by applying the optimized diffusion on top of the scalar potential; see blue dots in Fig. 4(d). By applying four-state PCCA to the effective simulations, we also verify the corresponding metastable sets to coincide with those of the original dynamics; see Fig. 4(c). Thus, the spectral matching enables us to find a (mostly constant) diffusion term in order to correct the coarse-grained system's kinetic properties.

As a further consistency check, we also apply spectral matching with just a single basis function, which is the constant, while leaving all other settings unchanged. The resulting expansion coefficient, which is nothing more than an effective diffusion constant, also equals 0.3.

V. CONCLUSIONS

We have introduced spectral matching as a method to define effective coarse-grained models reproducing the dynamics of a fine-grained system. Spectral matching can be used as stand alone or in combination with the existing force matching method. While force matching enforces thermodynamic consistency, spectral matching enforces kinetic consistency. The goal of spectral matching is to retain slow processes of the original dynamics, while faster motions are considered less relevant and may be lost in the process. For two specific settings, we have presented the resulting data-based regression problems that follow from spectral matching. The first setting is the estimation of an effective potential for an overdamped dynamics, and the second is concerned with learning an effective diffusion to correct the kinetics induced by force matching. We have demonstrated by several examples that spectral matching can be used to learn governing equations which retain the slow part of the dynamics. We found that suitable regularization is vital to the success of the method.

At this point, the question arises which of the two setups for spectral matching presented in this study is preferable. The answer certainly depends on the context: oftentimes, being able to define the system in terms of a potential energy alone is desirable, especially with regard to the problem of defining transferable parameters for coarse-grained systems. A position-dependent diffusion is harder to interpret in physical terms, and solving the spectral matching problem for the diffusion seems more involved in practice. On the other hand, the second setup bears the promise of ensuring both thermodynamic and kinetic consistency and may therefore lead to more accurate effective models. Also, the second use-case avoids setting the diffusion to a fixed value, which can be a hard problem in practice. While it may be possible to use the ratio of original and force matching time scales, as in Sec. IV D, the example in Sec. IV C suggests this may not always be possible.

Several questions are left open for future work. While we were able to use relatively simple basis sets here, the application of the method to large molecular systems will require the use of more powerful model classes. Using deep learning, or other state of the art machine learning techniques, in conjunction with spectral matching

is the topic of ongoing research. Along the same lines, the choice of test functions in high-dimensional spaces is another open problem to be considered in future work. A better theoretical understanding of the method is also required. For instance, the limiting behavior of spectral matching if the test and basis sets become exhaustive needs to be addressed, especially in the case where a potential different from the pmf is estimated by spectral matching. Moreover, we found in many cases that multiple different solutions with almost equivalent kinetic properties can be found by spectral matching. A systematic treatment of this phenomenon will also follow in future work.

REFERENCES

[1] D. Reith, M. Pütz, and F. Müller-Plathe, J. Comput. Chem. **24**, 1624 (2003).
[2] S. Izvekov and G. A. Voth, J. Phys. Chem. B **109**, 2469 (2005).
[3] W. Noid, J.-W. Chu, G. S. Ayton, V. Krishna, S. Izvekov, G. A. Voth, A. Das, and H. C. Andersen, J. Chem. Phys. **128**, 244114 (2008).
[4] M. S. Shell, J. Phys. Chem. **129**, 144108 (2008).
[5] Y. Wang, W. G. Noid, P. Liu, and G. A. Voth, Phys. Chem. Chem. Phys. **11**, 2002 (2009).
[6] W. G. Noid, J. Phys. Chem. **139**, 090901 (2013).
[7] A. P. Lyubartsev and A. Laaksonen, Phys. Rev. E **52**, 3730 (1995).
[8] S. O. Nielsen, C. F. Lopez, G. Srinivas, and M. L. Klein, J. Chem. Phys. **119**, 7043 (2003).
[9] S. J. Marrink, A. H. de Vries, and A. E. Mark, J. Phys. Chem. B **108**, 750 (2004).
[10] S. J. Marrink, H. J. Risselada, S. Yefimov, D. P. Tieleman, and A. H. de Vries, J. Phys. Chem. B **111**, 7812 (2007).
[11] W. Shinoda, R. DeVane, and M. L. Klein, Mol. Simul. **33**, 27 (2007).
[12] L. Monticelli, S. K. Kandasamy, X. Periole, R. G. Larson, D. P. Tieleman, and S.-J. Marrink, J. Chem. Theory Comput. **4**, 819 (2008).
[13] C. Clementi, Curr. Opin. Struct. Biol. **18**, 10 (2008).
[14] H. Mori, Prog. Theor. Phys. **33**, 423 (1965).
[15] R. Zwanzig, J. Stat. Phys. **9**, 215 (1973).
[16] A. J. Chorin, O. H. Hald, and R. Kupferman, Proc. Natl. Acad. Sci. U. S. A. **97**, 2968 (2000).
[17] A. J. Chorin, O. H. Hald, and R. Kupferman, Physica D **166**, 239 (2002).
[18] C. Hijón, P. Español, E. Vanden-Eijnden, and R. Delgado-Buscalioni, Faraday Discuss. **144**, 301 (2010).
[19] F. Legoll and T. Lelièvre, Nonlinearity **23**, 2131 (2010).
[20] W. Zhang and C. Schütte, Entropy **19**, 367 (2017).
[21] F. Legoll, T. Lelièvre, and S. Olla, Stochastic Stochastic Appl. **127**, 2841 (2017).
[22] G. Pavliotis and A. Stuart, *Multiscale Methods: Averaging and Homogenization* (Springer Science & Business Media, 2008).
[23] J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. D. Chodera, C. Schütte, and F. Noé, J. Chem. Phys. **134**, 174105 (2011).
[24] F. Noé and C. Clementi, Curr. Opin. Struct. Biol. **43**, 141 (2017).

[25] M. A. Rohrdanz, W. Zheng, M. Maggioni, and C. Clementi, J. Chem. Phys. **134**, 124116 (2011).

[26] C. Clementi, H. Nymeyer, and J. N. Onuchic, J. Mol. Biol. **298**, 937 (2000).

[27] S. Matysiak and C. Clementi, J. Mol. Biol. **343**, 235 (2004).

[28] A. Davtyan, N. P. Schafer, W. Zheng, C. Clementi, P. G. Wolynes, and G. A. Papoian, J. Phys. Chem. B **116**, 8494 (2012).

[29] L. Ponzoni, G. Polles, V. Carnevale, and C. Micheletti, Structure **23**, 1516 (2015).

[30] A. V. Sinitskiy, M. G. Saunders, and G. A. Voth, J. Phys. Chem. B **116**, 8363 (2012).

[31] L. Boninsegna, R. Banisch, and C. Clementi, J. Chem. Theory Comput. **14**, 453 (2018).

[32] J. F. Rudzinski and W. G. Noid, J. Phys. Chem. **135**, 214101 (2011).

[33] L. Delle Site and M. Praprotnik, Phys. Rep. **693**, 1 (2017).

[34] S. L. Brunton, J. L. Proctor, and J. N. Kutz, Proc. Natl. Acad. Sci. U. S. A. **113**, 3932 (2016).

[35] L. Boninsegna, F. Nüske, and C. Clementi, J. Chem. Phys. **148**, 241723 (2018).

[36] D. Crommelin and E. Vanden-Eijnden, Commun. Math. Sci. **4**, 651 (2006).

[37] D. Crommelin and E. Vanden-Eijnden, Multiscale Model. Simul. **9**, 1588 (2011).

[38] E. B. Davies, Proc. London Math. Soc. **s3-45**, 133 (1982).

[39] E. B. Davies, J. London Math. Soc. **s2-26**, 541 (1982).

[40] P. Deuflhard, W. Huisinga, A. Fischer, and C. Schütte, Linear Algebra Appl. **315**, 39 (2000).

[41] F. Noé and F. Nüske, Multiscale Model. Simul. **11**, 635 (2013).

[42] C. Schütte and M. Sarich, *Metastability and Markov State Models in Molecular Dynamics: Modeling, Analysis, Algorithmic Approaches* (American Mathematical Society and Courant Institute of Mathematical Sciences, 2013).

[43] F. Nüske, B. G. Keller, G. Pérez-Hernández, A. S. J. S. Mey, and F. Noé, J. Chem. Theory Comput. **10**, 1739 (2014).

[44] A. Mardt, L. Pasquali, H. Wu, and F. Noé, Nat. Commun. **9**, 5 (2018).

[45] A. Pazy, *Semigroups of Linear Operators and Applications to Partial Differential Equations* (Springer, 1983).

[46] M. Dellnitz and O. Junge, SIAM J. Numer. Anal. **36**, 491 (1999).

[47] W. Zhang, C. Hartmann, and C. Schutte, Faraday Discuss. **195**, 365 (2016).

[48] G. Ciccotti, T. Lelievre, and E. Vanden-Eijnden, Commun. Pure Appl. Math. **61**, 371 (2008).

[49] K. K. Bejagam, S. Singh, Y. An, and S. A. Deshmukh, J. Phys. Chem. Lett. **9**, 4667 (2018).

[50] L. Zhang, J. Han, H. Wang, R. Car, and E. Weinan, J. Chem. Phys. **149**, 034101 (2018); e-print arXiv:1802.08549.

[51] J. Wang, S. Olsson, C. Wehmeyer, A. Pérez, N. E. Charron, G. de Fabritiis, F. Noe, and C. Clementi, ACS Cent. Sci. **5**, 755 (2019).

[52] H. Chan, M. J. Cherukara, B. Narayanan, T. D. Loeffler, C. Benmore, S. K. Gray, and S. K. R. S. Sankaranarayanan, Nat. Commun. **10**, 379 (2019).

[53] F. Nüske, P. Koltai, L. Boninsegna, and C. Clementi, J. Nonlinear Sci. (submitted); e-print arXiv:1901.01557.

[54] H. Zou and T. Hastie, J. R. Stat. Soc.: Ser. B **67**, 301 (2005).

[55] *An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation*, edited by G. R. Bowman, V. S. Pande, and F. Noé (Springer Netherlands, 2014).

[56] B. Trendelkamp-Schroer, H. Wu, F. Paul, and F. Noé, J. Chem. Phys. **143**, 174101 (2015).

[57] P. Deuflhard and M. Weber, Linear Algebra Appl. **398**, 161 (2005).

[58] S. Röblitz and M. Weber, Adv. Data Anal. Classif. **7**, 147 (2013).

[59] M. K. Scherer, B. Trendelkamp-Schroer, F. Paul, G. Pérez-Hernández, M. Hoffmann, N. Plattner, C. Wehmeyer, J.-H. Prinz, and F. Noé, J. Chem. Theory Comput. **11**, 5525 (2015).

[60] F. Nüske, H. Wu, J.-H. Prinz, C. Wehmeyer, C. Clementi, and F. Noé, J. Chem. Phys. **146**, 094104 (2017).