Identification of kinetic order parameters for non-equilibrium dynamics

Cite as: J. Chem. Phys. **150**, 164120 (2019); https://doi.org/10.1063/1.5083627
Submitted: 29 November 2018 . Accepted: 04 April 2019 . Published Online: 26 April 2019

Fabian Paul, Hao Wu, Maximilian Vossel, Bert L. de Groot, and Frank Noé

The Journal of Chemical Physics ARTICLE scitation.org/journal/jcp

Identification of kinetic order parameters for non-equilibrium dynamics

Cite as: J. Chem. Phys. 150, 164120 (2019); doi: 10.1063/1.5083627
Submitted: 29 November 2018 • Accepted: 4 April 2019 •
Published Online: 26 April 2019

Fabian Paul,[1,2] Hao Wu,[3] Maximilian Vossel,[4] Bert L. de Groot,[4] and Frank Noé[1,5,a)]

AFFILIATIONS

[1] Department of Mathematics and Computer Science, FU Berlin, Arnimallee 6, 14195 Berlin, Germany
[2] University of Chicago, 929 East 57th Street, Chicago, Illinois 60637, USA
[3] Tongji University, School of Mathematical Sciences, Shanghai 200092, People's Republic of China
[4] Max Planck Institute for Biophysical Chemistry, Am Fassberg 11 D-37077 Göttingen, Germany
[5] Department of Chemistry, Rice University, Houston, Texas 77005, USA

Note: This article is part of the Special Topic "Markov Models of Molecular Kinetics" in J. Chem. Phys.
[a)] **Electronic mail:** frank.noe@fu-berlin.de

ABSTRACT

A popular approach to analyze the dynamics of high-dimensional many-body systems, such as macromolecules, is to project the trajectories onto a space of slowly varying collective variables, where subsequent analyses are made, such as clustering or estimation of free energy profiles or Markov state models. However, existing "dynamical" dimension reduction methods, such as the time-lagged independent component analysis (TICA), are only valid if the dynamics obeys detailed balance (microscopic reversibility) and typically require long, equilibrated simulation trajectories. Here, we develop a dimension reduction method for non-equilibrium dynamics based on the recently developed Variational Approach for Markov Processes (VAMP) by Wu and Noé. VAMP is illustrated by obtaining a low-dimensional description of a single file ion diffusion model and by identifying long-lived states from molecular dynamics simulations of the KcsA channel protein in an external electrochemical potential. This analysis provides detailed insights into the coupling of conformational dynamics, the configuration of the selectivity filter, and the conductance of the channel. We recommend VAMP as a replacement for the less general TICA method.

I. INTRODUCTION

Much understanding about molecular kinetics has been gained by modeling kinetics with Markov state models (MSMs),[1–4] rate equation models,[5] or diffusion map-based models.[6–8] A key element in all of these methods is that the dynamics are modeled in a low-dimensional space of collective variables.[9–11] In MSMs and rate equation models, there is a direct link between the kinetic model, consisting of a set of states and transition probabilities, and the underlying microscopic dynamical equations via the spectral decomposition of Markov operators.[12] As a result of this theory, the natural collective variables to describe the long-time dynamics are the eigenfunctions of the Markov operator.[11] In practice, the eigenfunctions and eigenvalues of the (Markov) transfer operator can be approximately computed directly from molecular dynamics (MD) simulation data by means of the Noé-Nüske variational approach.[13,14] This has led to a wide application of spectral methods in the molecular dynamics community, in particular, the time-lagged independent component analysis (TICA).[15–17]

However, the application of TICA is only truly justified if the dynamics fulfill the principle of detailed balance (microscopic reversibility) and if the dynamical equations are stationary, i.e., do not change as a function of time. Moreover, since TICA is a data-driven method, the reversibility and stationarity must be approximately met in the (finite) simulation data. Ignoring this limitation can result in systematic errors. For instance, if TICA is applied to non-equilibrium data (such as data that consist of short trajectories that were not initialized from the equilibrium distribution), the computed eigenvalues and eigenfunctions incur large biases.[18] A valid alternative method that can be applied to non-equilibrium

data is Koopman reweighting.[18] This method removes estimation bias but empirically induces a large variance, as seen in the results of Ref. 18.

The restriction to equilibrium data impedes the analysis of interesting and biologically relevant molecular systems whose function relies on nonreversible dynamics. Ion conduction in channel proteins is an example for such a process since the ion current is driven by external and perhaps time-dependent electric fields and chemical potentials. Therefore, a dynamical dimension reduction method that is similar to TICA but is directly applicable to non-equilibrium dynamics or even to non-equilibrium data would be desirable.

Recently, a new approach for dimensionality reduction in dynamic systems was proposed by Wu and Noé.[19] The *variational approach to Markov processes* (VAMP) dispenses with the assumptions of stationarity and reversibility. This was made possible by reformulating the problem of dimensionality reduction as a regression problem. Similarly to the reversible methods like TICA, VAMP can be directly applied to MD simulation data; it is, hence, a data-driven method. Wu and Noé[19] showed that there exists an optimal low-rank approximation to the solution of the above regression problem. This gives rise to a low-dimensional space of order parameters that are chosen such that the regression error is minimized. Mathematically this space can be found by performing a restricted singular value decomposition (SVD)[20,21] of a regression matrix learned from the simulation data.

Mardt *et al.*[22] showed that VAMP can be used to train a deep neural network to find informative order parameters and derive a coarse-grained MSM for the conformational dynamics of the alanine dipeptide and the folding of the N-terminal domain of ribosomal protein L9 (NTL9). In their work, Mardt *et al.* focused on demonstrating that VAMP can be successfully used to select highly nonlinear transformations to approximate the singular functions. However, the MD simulations that they used were reversible and stationary.

In this work, we show that VAMP works as a dimension reduction method for non-equilibrium data that may or may not originate from an equilibrium system. The goal is to establish an alternative to TICA which can be applied to reduce the dimension of the data and keep the slow processes, no matter whether the data are too short to be equilibrated or if the underlying process is fundamentally out of equilibrium. We also extend the Chapman-Kolmogorov test, which is frequently used to validated MSMs,[4] to validate the Markov property of the dimensionality-reduced model obtained with VAMP. We demonstrate VAMP by identifying the slow collective variables for two non-equilibrium systems: (1) the asymmetric simple exclusion process (ASEP) which is a simple model of single file diffusion and (2) non-equilibrium MD simulation data[23] of the KcsA potassium ion channel in which an ion current is driven through the channel pore. Furthermore, using a simple model of diffusion in a low-dimensional energy landscape, we compare the biases of VAMP and TICA when applied to an ensemble of short trajectories that were initiated from a non-equilibrium distribution.

II. THEORY

We first lay a theoretical framework with the most important mathematical results. The more practically inclined reader is advised to skip to the Sec. III. The theoretical framework is formulated in the language of dynamical operators. The advantage of this formulation is that theoretical properties can be obtained by using linear methods—albeit in infinite many dimensions. The main theoretical result is a variational principle, which can then be used for the formulation of linear or nonlinear solvers (such as VAMPnets[22]). In order to go into even more theoretical detail, please refer to Ref. 19.

A. Exact dynamics in full configuration space

Let $\mathbf{x}$ be the coordinates in which the MD algorithm is Markovian (atom positions, velocities, box coordinates, etc.) Let $p(\mathbf{x}, t)$ be the probability density of finding the system in state $\mathbf{x}$ at time t. We are interested in $p(\mathbf{x}, \tau n)$, the density at times τn that are integer multiples of some lag time τ. At these times, the time evolution of p can be described with the following integral equation:

$$p(\mathbf{x}', t+\tau) = \int p(\mathbf{x}' \mid \mathbf{x}) p(\mathbf{x}, t)\, \mathrm{d}x = \mathscr{P}_\tau[p]. \tag{1}$$

Here, $\mathscr{P}_\tau$ stands for the propagation operator (or propagator) which can be thought of as the discrete-time analog of the Fokker-Planck operator. $p(\mathbf{x}'|\mathbf{x})$ denotes the conditional probability density of visiting an infinitesimal phase space volume around point $\mathbf{x}'$ at time $t + \tau$ given that the phase space point $\mathbf{x}$ was visited at the earlier time t.

An equivalent description of the time evolution is given by the following integral equation which defines the Koopman operator $\mathscr{K}_\tau$:

$$g(\mathbf{x}, t+\tau) = \int p(\mathbf{x}' \mid \mathbf{x}) f(\mathbf{x}')\, \mathrm{d}\mathbf{x}' = \mathscr{K}_\tau[f], \tag{2}$$

where f is an observable, i.e., in general, a function of positions and momenta. The result $g(\mathbf{x}, t + \tau)$ can be interpreted as the expectation value of f at time $t + \tau$ computed from an ensemble that was propagated for a time τ after having been started at time t from the single point $\mathbf{x}$

$$\mathscr{K}_\tau[f](\mathbf{x}) = \mathbb{E}_{t+\tau}(f \mid p(\mathbf{x}, t) = \delta(\mathbf{x})). \tag{3}$$

Here, δ is the (vectorial) Dirac delta function.

Both the propagator and the Koopman operator fulfill the Chapman–Kolmogorov equation

$$\mathscr{P}_{\tau_1+\tau_2} = \mathscr{P}_{\tau_1}\mathscr{P}_{\tau_2}, \tag{4}$$

$$\mathscr{K}_{\tau_1+\tau_2} = \mathscr{K}_{\tau_1}\mathscr{K}_{\tau_2}. \tag{5}$$

For stationary dynamics, this implies that expectations of any observable f can be computed for all times from the Koopman operator

$$g(\mathbf{x}, n\tau) = \mathbb{E}_{n\tau}(f \mid p(\mathbf{x}, 0) = \delta(\mathbf{x})) = \mathscr{K}_\tau^n f. \tag{6}$$

Expectations for an ensemble that was started from an arbitrary probability density p_0 can be computed from the following scalar product:

$$\mathbb{E}_{n\tau}(f \mid p(\mathbf{x}, 0) = p_0(\mathbf{x})) = \int p_0(\mathbf{x}) g(\mathbf{x}, n\tau)\, \mathrm{d}x. \tag{7}$$

For the computation of instantaneous and time-lagged variances and covariances,[24] similar equations that use the Koopman operator can be derived.

B. Formulation of ensemble propagation as a regression problem

Can a dynamic model be built using only expectation values that were computed from simulation data? This question has been addressed in a series of papers preceding our VAMP theory that have developed the so-called Koopman analysis.[25,26] We seek a small matrix $\mathbf{K}_\tau \in \mathbb{R}^{k\times k}$, called the Koopman matrix, that fulfills the equation

$$\mathscr{K}_\tau \mathbf{g} \approx \mathbf{K}_\tau^\top \mathbf{f} \tag{8}$$

in a sense that we explain below. Here, $\mathbf{f}$ and $\mathbf{g}$ are vectors of observables that can be arbitrary functions of the conformation. $\mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), f_2(\mathbf{x}), \dots)^\top$ and similarly for $\mathbf{g}$. We use the shorthand notation $(\mathscr{K}_t \mathbf{g})_i := \mathscr{K}_t g_i$ which means that the Koopman operator is applied element-wise to $\mathbf{g}$.

More formally, for fixed $\mathbf{f}$ and $\mathbf{g}$, the optimal data-dependent matrix $\mathbf{K}_\tau$ can be computed by minimizing the following error:

$$\varepsilon = \mathbb{E}_{\rho_0}\left[\left\|\mathscr{K}_\tau \mathbf{g} - \mathbf{K}_\tau^\top \mathbf{f}\right\|^2\right], \tag{9}$$

where ρ_0 is the empirical distribution of the simulation data, excluding time steps $t_i > T - \tau$, where T is the length of the (single) time series. By inserting the definitions of ρ_0 and $\mathscr{K}_\tau$ into (9), one finds that

$$\varepsilon = \sum_{\substack{t \text{ s.t.} \\ 0\le t\le T-\tau}} \left\|\mathbf{g}(\mathbf{x}(t+\tau)) - \mathbf{K}_\tau^\top \mathbf{f}(\mathbf{x}(t))\right\|^2, \tag{10}$$

which shows that the error is purely data-dependent.

Equations (9) and (10) have the form of a regression problem: a future window from the time series is regressed against the current window of the time series. This formulation avoids any assumption of microscopic reversibility.

C. Optimal low-dimensional observables

Unlike the Koopman matrix, the observable functions $\mathbf{f}$ and $\mathbf{g}$ cannot be chosen by only minimizing the regression error defined by (9) because the minimal $\varepsilon = 0$ can be trivially obtained by an uninformative model with $\mathbf{f}(\mathbf{x}) \equiv \mathbf{g}(\mathbf{x}) \equiv 1$.

In Ref. 19, the error of the approximate Koopman operator provided by model (8) was analyzed. It was shown that the model with the smallest approximation error in Hilbert-Schmidt norm is given by $\mathbf{f} = \boldsymbol{\psi} = (\psi_1, \dots, \psi_k)^\top$, $\mathbf{g} = \boldsymbol{\phi} = (\phi_1, \dots, \phi_k)^\top$, and $\mathbf{K}_\tau = \mathrm{diag}(\boldsymbol{\sigma}) = \mathrm{diag}(\sigma_1, \dots, \sigma_k)$ for a given k, and the corresponding approximation of the Koopman operator is

$$\mathscr{K}_\tau g \approx \sum_{i=1}^{k} \sigma_i \langle g, \phi_i\rangle_{\rho_1} \psi_i, \tag{11}$$

where σ_i is the ith largest singular value of $\mathscr{K}_\tau$ and ψ_i and ϕ_i are the corresponding left and right singular functions, respectively. (The singular value decomposition is to be understood of being applied after a whitening transformation of $\mathbf{f}$ and $\mathbf{g}$. See Subsection III A for details.) ρ_1 is the empirical distribution of simulation data excluding time steps $t_i < \tau$, and $\langle f, g\rangle_{\rho_1} = \int f(\mathbf{x}) g(\mathbf{x}) \rho_1(\mathbf{x})\, \mathrm{d}\mathbf{x}$.

It can be shown that the largest singular value σ_1 is always 1 and that the corresponding left and right singular functions are constant and identical to 1 for all $\mathbf{x}$.[18] Only the singular components σ_i, ψ_i, and ϕ_i with $i > 1$ contain kinetic information.

If $\boldsymbol{\psi}$ and $\boldsymbol{\phi}$ are approximated with a finite linear combination of ansatz functions, a corresponding finite-dimensional singular value decomposition of the whitened Koopman matrix can be used to compute the optimal superposition coefficients (see Subsection III A for details).

D. The kinetic map induced by the singular functions

For a Markov process, we can measure the difference between two points $\mathbf{x}$ and $\mathbf{y}$ by the kinetic distance[27] $D_\tau(\mathbf{x}, \mathbf{y})$, where

$$D_\tau^2(\mathbf{x}, \mathbf{y}) = \int \frac{(p(\mathbf{z}|\mathbf{x}) - p(\mathbf{z}|\mathbf{y}))^2}{\rho_1(\mathbf{z})}\, \mathrm{d}\mathbf{z}. \tag{12}$$

$D_\tau(\mathbf{x}, \mathbf{y}) = 0$ means $\mathbf{x}$ and $\mathbf{y}$ are equivalent for predicting the future evolution of the process. By using the singular components of $\mathscr{K}_\tau$, the square of the kinetic distance can be written as

$$D_\tau^2(\mathbf{x}, \mathbf{y}) = \sum_i \sigma_i^2 (\psi_i(\mathbf{x}) - \psi_i(\mathbf{y}))^2 \tag{13}$$

(see Appendix A for proof). If all but the k leading singular values are close 0, we have

$$D_\tau^2(\mathbf{x}, \mathbf{y}) \approx \|\mathrm{diag}(\boldsymbol{\sigma})\boldsymbol{\psi}(\mathbf{x}) - \mathrm{diag}(\boldsymbol{\sigma})\boldsymbol{\psi}(\mathbf{y})\|^2, \tag{14}$$

where $\mathrm{diag}(\boldsymbol{\sigma})$ denotes the $k \times k$ diagonal matrix with the singular values on its diagonal. This means that all the points $\mathbf{x}$ can be embedded into a k-dimensional Euclidean space by the kinetic map $\mathbf{x} \to \mathrm{diag}(\boldsymbol{\sigma})\boldsymbol{\psi}(\mathbf{x})$ with the structure of the kinetic distance preserved. Note that the extension of Ref. 27 to the commute distance[28] is not directly applicable to VAMP because the commute distance relies on the computation of relaxation time scales, which relies on the eigenvalue decomposition of the Markov operator and cannot be directly done with the singular value decomposition.

Also note that the kinetic distance defined in Eq. (12) depends on the empirical distribution of the data ρ_1. Therefore, $D_\tau^2(\mathbf{x}, \mathbf{y})$, in general, depends on how the system dynamics were sampled. For systems that possess a unique stationary distribution (see, for example, the ASEP model in Subsection IV A), ρ_1 can be set to the stationary distribution to define a kinetic distance that is independent from the sampling.

Furthermore, it is worth noting that the coherent sets of nonreversible Markov processes can also be identified from the k dominant singular components, and more details can be seen in Ref. 29. Also, note that the right singular functions $\boldsymbol{\phi}$ induce a kinetic map with respect to time-reversed propagation of the dynamics (unlike the kinetic map induced by $\boldsymbol{\psi}$ that uses conventional forward-time propagation).

III. METHODS

In this work, we use VAMP as a method for computing optimal kinetic order parameters for non-equilibrium dynamics using a linear combination of input features. However, in general, the scope of VAMP is larger: order parameters are not restricted to be linear combinations but can also be formed from a nonlinear combination of features as was demonstrated by Mardt *et al.*[22] by training a deep neural network (VAMPnet) to capture the conformational

dynamics of the alanine dipeptide and the N-terminal domain of ribosomal protein L9 (NTL9). Another application of VAMP is the scoring of input features (see publication "Variational Selection of Features for Molecular Kinetics" by Scherer *et al.* in this issue).

Using VAMP to find kinetic order parameters from a linear combination of molecular features is also called time-lagged canonical covariance analysis (TCCA)[22] and works as follows.

A. Dimension reduction using the variational approach for Markov processes (VAMPs)

Let $\boldsymbol{\chi}(t)$ be a multivariate time series where every element $\chi_i(t)$ is the time series of one molecular feature. Features can be Cartesian or internal coordinates (such as distances or dihedral angles) of the molecular system or functions thereof (such as the sine and cosine of dihedral angles or a step function that converts a distance into a contact). From the input features $\boldsymbol{\chi}(t)$, first the means $\boldsymbol{\mu}_0$ and $\boldsymbol{\mu}_1$ are computed from all data excluding the last and first τ steps of every trajectory, respectively,

$$\boldsymbol{\mu}_0 := \frac{1}{T-\tau}\sum_{t=0}^{T-\tau}\boldsymbol{\chi}(t), \tag{15}$$

$$\boldsymbol{\mu}_1 := \frac{1}{T-\tau}\sum_{t=\tau}^{T}\boldsymbol{\chi}(t). \tag{16}$$

Next, the instantaneous covariance matrices $\mathbf{C}_{00}$ and $\mathbf{C}_{11}$ and the time-lagged covariance matrix $\mathbf{C}_{01}$ are computed as follows:

$$\mathbf{C}_{00} := \frac{1}{T-\tau}\sum_{t=0}^{T-\tau}\left[\boldsymbol{\chi}(t)-\boldsymbol{\mu}_0\right]\left[\boldsymbol{\chi}(t)-\boldsymbol{\mu}_0\right]^{\top}, \tag{17}$$

$$\mathbf{C}_{11} := \frac{1}{T-\tau}\sum_{t=\tau}^{T}\left[\boldsymbol{\chi}(t)-\boldsymbol{\mu}_1\right]\left[\boldsymbol{\chi}(t)-\boldsymbol{\mu}_1\right]^{\top}, \tag{18}$$

$$\mathbf{C}_{01} := \frac{1}{T-\tau}\sum_{t=0}^{T-\tau}\left[\boldsymbol{\chi}(t)-\boldsymbol{\mu}_0\right]\left[\boldsymbol{\chi}(t+\tau)-\boldsymbol{\mu}_1\right]^{\top}. \tag{19}$$

After that, a Koopman matrix $\bar{\mathbf{K}}$ is computed in the basis of whitened[19,30] input features

$$\bar{\mathbf{K}} := \mathbf{C}_{00}^{-\frac{1}{2}}\mathbf{C}_{01}\mathbf{C}_{11}^{-\frac{1}{2}}. \tag{20}$$

Then, the singular value decomposition (SVD) of $\bar{\mathbf{K}}$ is performed, giving orthonormal matrices $\mathbf{U}'$ and $\mathbf{V}'$ as well as $\mathbf{S} = \mathrm{diag}(\boldsymbol{\sigma})$ such that

$$\bar{\mathbf{K}} = \mathbf{U}'\mathbf{S}\mathbf{V}'^{\top}. \tag{21}$$

Finally, the input conformations are mapped to the left singular functions $\boldsymbol{\psi}$ and right singular functions $\boldsymbol{\phi}$ as follows:

$$\boldsymbol{\psi}(t) := \mathbf{U}'^{\top}\mathbf{C}_{00}^{-\frac{1}{2}}\left[\boldsymbol{\chi}(t)-\boldsymbol{\mu}_0\right], \tag{22}$$

$$\boldsymbol{\phi}(t) := \mathbf{V}'^{\top}\mathbf{C}_{11}^{-\frac{1}{2}}\left[\boldsymbol{\chi}(t)-\boldsymbol{\mu}_1\right], \tag{23}$$

where $\boldsymbol{\psi}(t)$ and $\boldsymbol{\phi}(t)$ are the sought-after kinetic order parameters. Since the left singular functions $\boldsymbol{\psi}(t)$ induce a kinetic map for the (conventional) forward-time propagator, they are the natural choice of order parameters if one wants to perform a clustering of space to obtain state definitions. For simplicity, we will call them VAMP components.

Note that the algorithm above performs a Canonical Correlation Analysis (CCA)[31] in time and is, hence, also called Time-lagged CCA (TCCA).[19] The singular value decomposition in the whitened basis (20) and (21) is also called the generalized[21] or restricted[20] SVD of $\mathbf{C}_{01}$ under constraints imposed by $\mathbf{C}_{00}$ and $\mathbf{C}_{11}$.

B. The variational score

In Subsection III A, VAMP was used to linearly combine molecular features to compute kinetic order parameters. A question that remained unanswered is how to select the best molecular features to use as input. This question can be answered by computing the variational score of the dimensionality-reduced kinetic model. The VAMP-r score is defined as the sum of the leading m largest singular values that have been taken to the power of r (see Ref. 19 and the publication "Variational Selection of Features for Molecular Kinetics" by Scherer *et al.* in this issue)

$$\mathrm{VAMP}_{r,\mathrm{train}} = \sum_{i=1}^{m}\sigma_i^r. \tag{24}$$

In a situation with infinite sampling, where the singular values are known without statistical error, the best selection of molecular features is the one that maximizes the VAMP-r score. In a practical setting, however, where the time series data is finite, direct maximization of the VAMP-r score is not possible due to model overfitting.[32] That is why the VAMP-r score needs to be computed in a cross-validated manner.

Cross-validation works by splitting the trajectory data into two sets: the training set, from which a dimensionality-reduced model is estimated, and the test set, against which the model is tested. From the training set, the matrices $\mathbf{U}^{\mathrm{train}} = \mathbf{C}_{00}^{-\frac{1}{2}}\mathbf{U}'$ and $\mathbf{V}^{\mathrm{train}} = \mathbf{C}_{11}^{-\frac{1}{2}}\mathbf{V}'$ are computed, where $\mathbf{C}_{00}$, $\mathbf{C}_{11}$, $\mathbf{U}'$, and $\mathbf{V}'$ are computed from the training data according to Eqs. (17), (18), and (21). Next, the test score is computed from the equation

$$\mathrm{VAMP}_{r,\mathrm{test}} = \sum_{i=1}^{m}\varkappa_i^r, \tag{25}$$

where $\varkappa_i$ is the ith singular value of the matrix product $\mathbf{ABC}$[33] with

$$\mathbf{A} = \left(\mathbf{U}^{\mathrm{train}\top}\mathbf{C}_{00}^{\mathrm{test}}\mathbf{U}^{\mathrm{train}}\right)^{-\frac{1}{2}}, \tag{26}$$

$$\mathbf{B} = \mathbf{U}^{\mathrm{train}\top}\mathbf{C}_{01}^{\mathrm{test}}\mathbf{V}^{\mathrm{train}}, \tag{27}$$

$$\mathbf{C} = \left(\mathbf{V}^{\mathrm{train}\top}\mathbf{C}_{11}^{\mathrm{test}}\mathbf{V}^{\mathrm{train}}\right)^{-\frac{1}{2}}, \tag{28}$$

and where $\mathbf{C}_{00}^{\mathrm{test}}$, $\mathbf{C}_{01}^{\mathrm{test}}$, and $\mathbf{C}_{11}^{\mathrm{test}}$ have been computed from the test data via Eqs. (17), (19), and (18) (with the caveat that the means $\boldsymbol{\mu}_0$ and $\boldsymbol{\mu}_1$ of the *training* data have to be subtracted). Finally, the k-fold cross-validated test score is computed by repeating the splitting of the data into test and training data k times, computing one test score for each partition of the data and then taking the average of the individual test scores.

C. The non-equilibrium Chapman-Kolmogorov test

For stationary (but possibly nonreversible) dynamics, the full-state-space Koopman operator fulfills the Markov property (5). It shares this property with the propagator and with the transition matrix of MSMs.

In the context of molecular dynamics simulation, the Markov property is often exploited to calculate long-time-scale properties ($\mathscr{K}_{n\tau}$) from short-lag-time estimates ($\mathscr{K}_\tau$). One of the most important long-time-scale properties is the stationary distribution that can be computed from a MSM by applying the transition matrix an infinite number of times to an initial probability distribution.

To extrapolate to higher multiples of the lag-time, the Markov property needs to hold. While this property is guaranteed for the full-state-space dynamical operators, it is not necessarily fulfilled for dimensionality-reduced dynamical models like the transition matrix of a MSM or an approximated Koopman operator. Therefore, the Markov property is typically tested by comparing $\mathscr{K}_{n\tau}$ to $\mathscr{K}_\tau^n$ for the multiples of the lag time $n\tau$.

The standard way to perform this test is to compare the direct estimate of a time-lagged covariance

$$\mathrm{cov}_{\mathrm{est}}(f,g;\, n\tau) = \langle f, \mathscr{K}_{n\tau} g\rangle_\rho \tag{29}$$

from the simulation data to the model-prediction of the same covariance

$$\mathrm{cov}_{\mathrm{pred}}(f,g;\, n\tau) = \langle f, \mathscr{K}_\tau^n g\rangle_\rho, \tag{30}$$

where f and g are some functions of the configuration-space coordinates. When f and g are indicator functions this test is known under the name Chapman-Kolmogorov test.[34] Here, we propose to perform the same comparison for the data-driven estimate of the dimensionality-reduced Koopman operator. See Appendix C 1 for details.

To make the test independent on the subjective choice of the functions f and g, the left and right singular functions of the Koopman operator estimated at the lowest multiple of the lag time $1 \times \tau$ can be used as f and g, respectively. This choice is in the spirit of the Chapman-Kolmogorov test as it is typically applied to Markov models of metastable molecular kinetics. There, the test is typically applied to the probability of staying in one of the metastable states, which constitutes a particular hard test that requires data that thoroughly samples exit and entry events into the metastable states.[4,34] Characteristic (indicator) functions of the metastable states are related by a linear transform to the eigenfunctions of the transfer operator.[12] By analogy, we assume here that using the singular functions in the Chapman-Kolmogorov test also constitutes a particular hard test.

D. Interpretation of the VAMP components and spectral clustering

For reversible dynamics, the theory of conformational dynamics describes how the leading eigenfunctions can be used to understand which structural changes are associated with the slowest processes and to find the metastable states via spectral clustering.[12,35]

For non-equilibrium dynamics, we can replace the eigenfunctions by the left singular functions found by VAMP. As for TICA,[16] we can interpret the ith kinetic order parameter in terms of structural changes by computing its correlation with all features χ_j

$$\mathrm{corr}(\psi_i, \chi_j) = \frac{\frac{1}{T-\tau}\sum_{0\le t<T-\tau}\psi_i(t)(\chi_j(t)-\bar{\chi}_j)}{\sqrt{\frac{1}{T-\tau}\sum_{0\le t<T-\tau}(\chi_j(t)-\bar{\chi}_j)^2}}, \tag{31}$$

$$= \frac{(\mathbf{C}_{00}^{\frac{1}{2}}\mathbf{U}')_{ji}}{\sqrt{(\mathbf{C}_{00})_{jj}}}, \tag{32}$$

and by visualizing the most-correlated features. In the last equation, $\bar{\chi}_j$ denotes the empirical mean of feature χ_j computed from the data in time steps $0 \le t_i < T - \tau$.

Furthermore, we can compute the long-lived states of non-equilibrium dynamics by performing spectral clustering in the VAMP components in a similar way as it is done with the dominant eigenspace for equilibrium dynamics in Ref. 36. Let $\boldsymbol{\psi}$ be the vector that contains n_{spec} leading singular functions (with singular values close to one, including the constant singular function). Let $\mathbf{A} \in \mathbb{R}^{n_{\mathrm{spec}} \times n_{\mathrm{spec}}}$. Then, the vector of macrostate memberships $\mathbf{m} \in \mathbb{R}^{n_{\mathrm{spec}}}$ is given by

$$\mathbf{m}(t) = \mathbf{A}\boldsymbol{\psi}(t). \tag{33}$$

See Appendix C 2 or Ref. 36 for the algorithm to compute $\mathbf{A}$. The element $m_i(t)$ encodes the degree of membership of the conformation sampled at time t in the macrostate i. The memberships at every time step always sum to one (which expresses the necessity of belonging to some macrostate with certainty) and, depending on the specific algorithm that was used to compute $\mathbf{A}$, are confined between 0 and 1[35] or not.[36] The memberships define the macrostates in a fuzzy manner; that is, every conformation belongs to macrostate i with a degree of membership given by $m_i(t)$. Fuzzy states can be converted into crisp states by imposing a cutoff on the memberships and treating conformations with memberships larger than the cutoff as being part of the crisp state. Structural differences between states can be found using significant distance analysis.[37]

IV. RESULTS

A. Model of single file diffusion: The asymmetric simple exclusion process

The asymmetric simple exclusion process (ASEP) is a generic model for single file diffusion. It was originally formulated by MacDonald *et al.*[38] as a model for the kinetics of protein synthesis and was independently introduced by Spitzer[39] in the mathematical literature. Since then, it has been extensively analyzed and applied to model phenomena such as macromolecular transport, conductivity, traffic flow, sequence alignment, and molecular motors (see Refs. 40 and 41 and references therein).

The ASEP consists of a linear chain of N_{sites} sites each of which can either be empty or occupied by exactly one particle, resulting in a large state space with $2^{N_{\mathrm{sites}}}$ elements [Fig. 1(a)]. If the first site is empty, a particle is inserted with a rate α. Particles can move to adjacent unoccupied sites with rate p in the forward and rate q in the backward direction. In the last site, particles are annihilated with rate β. Hence, the ASEP is a driven (nonreversible) Markovian multiparticle system. Here, we show that VAMP can be used

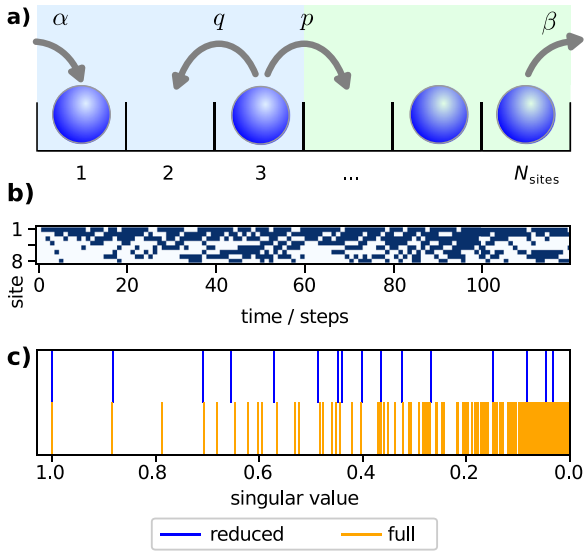

FIG. 1. (a) The asymmetric simple exclusion process is a model for single file diffusion. It consists of a linear chain of N_{sites} sites along which particles can move. Particles are inserted at position 1 with rate α and annihilated at site N_{sites} with rate β. (b) the first 120 time steps of an exemplary realization of a occupancy time trace for the ASEP with parameters $N_{\text{sites}} = 8$, $\alpha = \beta = p = 1$, and $q = 1/3$. (c) Singular values of the 17×17 Koopman model trained on the time series (blue, upper spectrum) and singular values of the full ASEP model (orange, lower spectrum).

to train a low-dimensional model that allows us to reproduce the time-lagged covariances and autocovariances for a large range of lag-times.

We use the master equation formulation of the ASEP as our true reference (Appendix B 1). The model parameters are chosen as $N_{\text{sites}} = 8$, $\alpha = \beta = p = 1$, and $q = 1/3$. Since VAMP works with a finite lag-time, we convert the master equation model to a transition matrix by taking the matrix-exponential of the master equation coefficient matrix. From the transition matrix, we generate a long trajectory with $N_{\text{steps}} = 10^6$ steps. The trajectory is encoded as a matrix of shape $N_{\text{steps}} \times N_{\text{sites}}$ where every row represents the occupancy pattern at a given time point [see Fig. 1(b) for an example of a transposed trajectory matrix].

We estimate an empirical Koopman matrix using VAMP at a lag time of $\tau = 1$ steps and using a basis consisting of two groups of features. The first group consists of the site occupancy vectors [the columns of the matrix shown in Fig. 1(b)]. The second set of features is a 9 dimensional vector that contains the "one-hot" encoded number of occupied sites. That is, element i in the second feature set is 1 if and only if there are i occupied sites. Our selection of features already constitutes a dimensionality-reduction since we estimate the Koopman model in the 8 + 9-dimensional space of feature vectors and not in the 2^8-dimensional state space. As a consequence, the spectrum of the empirical Koopman model consists of only 17 singular values. Not all singular values of the true model can be reproduced [see Fig. 1(c)], still large singular values approximately agree. The singular values of the empirical model decay quickly with increasing rank [see top part of Fig. 1(c)]. Therefore, we discard the very small singular components and further reduce the rank of the model to 11 dimensions (not counting the first singular function which is the constant function). Despite these two dimension reductions, we will show later that physically interesting observables are correctly captured by the 11-dimensional model.

To gain some physical understanding of the true singular functions of the ASEP model, we cluster the space of the leading 9 VAMP components of the true transition matrix with the PCCA+ algorithm without using any further approximation (see KcsA application below for more details on PCCA+ clustering). This allows us to group all the possible site occupancy patterns (microstates) into 9 macrostates. We select 9 states because in the true spectrum, a relatively large gap follows a denser cluster of singular values at position 9 [see lower part of Fig. 1(c)]. Macrostates are shown in Fig. 2, with the microstates ordered from low macrostate membership to high macrostate membership (from left to right). The top-membership microstates are characterized by long uninterrupted segments with the same occupancy (long occupied/long empty segments) and show only one alternation from occupied to unoccupied (shock) along the queue. Macrostates differ in the position of the shock. Microstates with lower memberships resemble the top membership states but show a noisier shock profile with more alternations between occupied and unoccupied. Macrostates also differ in the average number of occupied sites.

Next, we test whether our choice of the 17-dimensional basis that consists of the occupancy vector and the one-hot encoded occupancy affects the capability of PCCA+ to find the correct macrostates. Therefore, we repeat the PCCA+ clustering using the singular functions that were approximated with the Koopman model. Since the simple basis does not allow us to capture all leading 9 singular components correctly, we perform the comparison in the space of the leading 3 components (counting the constant component). Results are shown in the supplementary material, Fig. 5, and show good agreement between macrostates computed from the true and the approximate model. With an increased number of macrostates, the results deviate.

To test the predictive power of the reduced model, we compare observed time-lagged covariances to the model prediction

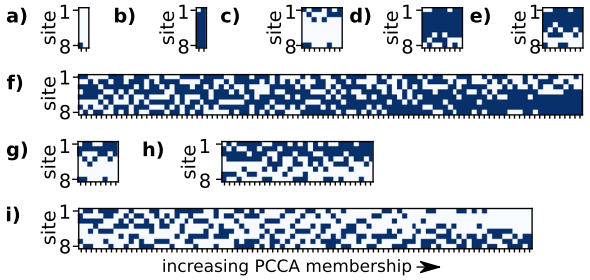

FIG. 2. The nine dominant long-lived macrostates of the ASEP with parameters $N_{\text{sites}} = 8$, $\alpha = \beta = p = 1$, and $q = 1/3$. In each inset (a)–(i), one macrostate is shown. Occupancy vectors of all microstates in a macrostate are ordered along the microstate axis (x-axis) with increasing memberships. Dark blue squares mark occupied sites, and white squares mark empty sites. The microstates with the highest macrostate memberships (right-most patterns) are characterized by long uninterrupted segments with the same occupancy and a single jump from occupied to unoccupied (shock) for states (c)–(i). Macrostates differ in the location of the shock.

of the same covariance using the Chapman-Kolmogorov test. We pick one of the observables $f = N_{\text{front}}$ to be the number of particles in the first half of the queue and the second observable $g = N_{\text{back}}$ the number particles in the second half. Estimates for the observed and the predicted time-lagged covariance of N_{front} and N_{back} computed from Eqs. (29) and (30) for multiple lag times are shown in Figs. 3(b)–3(e). For comparison, we also show the true covariances computed from the full ASEP model without using the VAMP approximation (shown in gray in Fig. 3). The Chapman-Kolmogorov test shows that predictions from the dimensionality-reduced VAMP model agree with the observed covariances computed from the time series data as well as with the results from the full model.

To make the Chapman-Kolmogorov test less dependent on the subjective choice of observables f and g, we repeat the test but this time selecting the observables to be identical to the singular functions $\psi_i^{(1)}$ and $\phi_i^{(1)}$, respectively, that were estimated from the dimensionality-reduced model estimated at lag time $\tau = 1$ steps. That is, we compare $\text{cov}_{\text{est}}(\psi_i^{(1)}, \phi_i^{(1)}; n\tau)$ to $\text{cov}_{\text{pred}}(\psi_i^{(1)}, \phi_i^{(1)}; n\tau)$. Results are shown in Fig. 4. The figure shows that the Chapman-Kolmogorov test succeeds for the all pairs of singular functions, that is, model predictions of covariances are consistent with the re-estimated covariances for all lag times. Predictions from the VAMP model are in good agreement with the true covariances that were computed from the full ASEP model. The dimensionality-reduced model does not correctly reproduce the second and third true singular function but reproduces the fourth true singular function (see the supplementary material, Fig. 1). To obtain this approximate agreement of the leading singular functions, it was necessary to include the one-hot-encoded count of occupied sites into the set of input features to VAMP. The mismatch between the remaining singular functions and singular values of the true and reduced model [see Fig. 1(c) and the supplementary material, Fig. 1] is a consequence of the very simple set of input features that was used to estimate $\mathbf{K}_\tau$. Had the reduced model not been trained on the 17-dimensional occupancy vectors but on the 2^8-dimensional full state space, the agreement would have been exact. Also using a more expressive set of basis functions[22,42,43] could have produced a richer reduced model that captures more singular components of the full model. Despite the simple approach, some observables can be modeled correctly.

Besides the estimation of a Koopman model, the typical use of VAMP will be to compute kinetic order parameters for

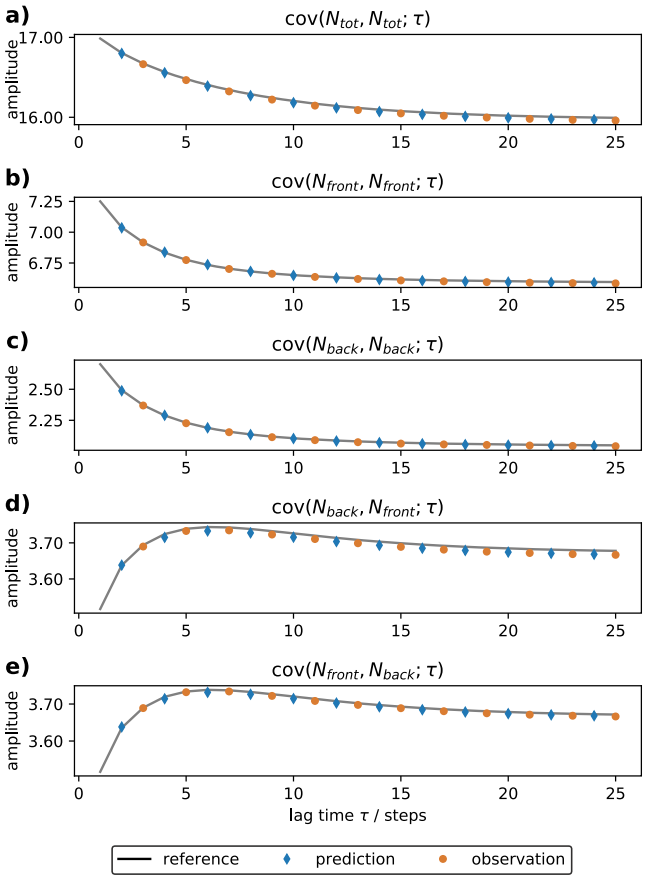

FIG. 3. Chapman-Kolmogorov test results for the low-dimensional Koopman matrix estimated from the ASEP model with parameters $N = 8$, $\alpha = \beta = p = 1$, and $q = 1/3$. (a) cov(N_{tot}, N_{tot}; τ), (b) cov(N_{front}, N_{front}; τ), (c) cov(N_{back}, N_{back}; τ), (d) cov(N_{back}, N_{front}; τ), and (e) cov(N_{front}, N_{back}; τ). N_{tot} is the total particle count in the queue. N_{front} is the total particle count in the first half of the queue [blue shaded area in Fig. 1(a)]. N_{back} is the particle count in the second half [green shaded area in Fig. 1(a)]. The true reference is computed from the full ASEP model by using Eq. (30) with the true Koopman operator.

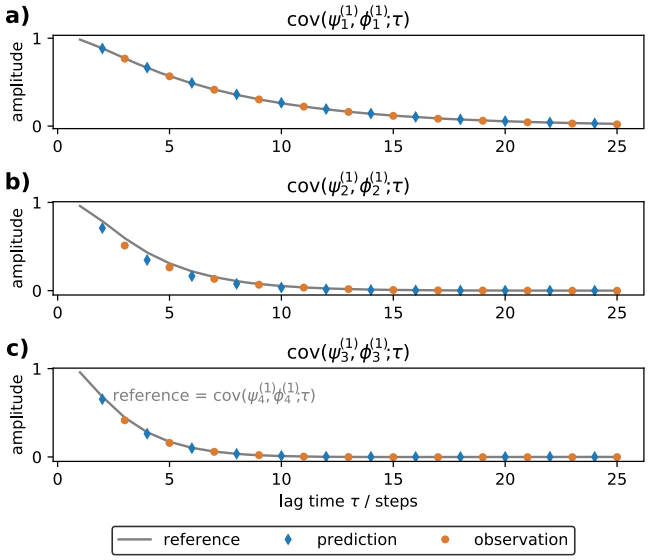

FIG. 4. Same as Fig. 3 but with singular functions chosen as observables. (a) $\text{cov}(\psi_1^{(1)}, \phi_1^{(1)}; \tau)$, (b) $\text{cov}(\psi_2^{(1)}, \phi_2^{(1)}; \tau)$, and (c) $\text{cov}(\psi_3^{(1)}, \phi_3^{(1)}; \tau)$. $\psi_i^{(1)}$ and $\phi_i^{(1)}$ are left and right singular functions, respectively, of the Koopman matrix estimated at the smallest lag time $\tau = 1$. In (c), $\text{cov}(\psi_4^{(1)}, \phi_4^{(1)}; \tau)$ of the true ASEP transition matrix serves as the reference since the third singular function of the dimensionality-reduced VAMP model predominantly matches the fourth true singular function (see the supplementary material, Fig 1).

nonreversible kinetics. To assess the improvement of these order parameters over the independent components obtained from TICA, we compare the kinetic distance obtained from TICA and VAMP to the true reference. We compute the true reference of the kinetic distance by applying Eq. (12) to the true ASEP transition matrix in a complete basis. We set ρ_1 in Eq. (12) to the true stationary distribution. We compare this reference to the VAMP estimate computed from Eq. (14) using the same full basis as well as to the TICA estimate. The TICA estimate of the kinetic distance is computed from a modified Eq. (14) with the singular values replaced by the TICA eigenvalues and the right singular functions replaced by the TICA eigenfunctions. This version is the default in the PyEMMA software.[44] Results are shown in Fig. 5. As implied by VAMP theory, the VAMP estimate converges to the true reference as the number of singular components is increased. In contrast to this, the TICA estimate does not converge to the true reference. This is expected since for non-equilibrium dynamics, the kinetic distance cannot be expressed using only the right eigenfunctions alone that TICA provides.[27] Full kinetic distances between all states are given in the supplementary material, Fig. 4.

In summary, the application of VAMP to the ASEP model shows that VAMP can accurately capture the dominant singular functions and can be used to accurately compute time-lagged auto-covariances and cross-covariances of physical quantities like the occupancy of the first and second half of the queue. The ASEP is a genuinely nonreversible model. Therefore, its dimension-reduction can only be accomplished with methods like VAMP that are capable of modeling nonreversible processes and do not rely on detailed balance. A decomposition of the state space into 9 macrostates shows that the location of the shock (jump from occupied to unoccupied segments) allows us to approximately distinguish the macrostates for the ASEP parameter settings that we chose.

B. Application of VAMP to a reversible system in the limit of non-equilibrium sampling

While the ASEP system is intrinsically non-equilibrium as its dynamical equations violate detailed balance, we now investigate the performance of VAMP when the underlying dynamics obey detailed balance, but the data does not reflect the equilibrium distribution. In cases where the metastable states are reversibly connected, reweighting methods[18,45] and reversible maximum-likelihood MSMs[4,46] have been shown to provide unbiased estimates and to recover the equilibrium kinetics from non-equilibrium data. When transitions between states have only been sampled in one direction, the current MSM practice is simply to discard the not reversibly connected states.[4,46] In VAMP, this is not necessary because VAMP does not require a stationary distribution to be computed.

Here, we study the performance of VAMP on non-equilibrium data generated from a 1-D double-well energy landscape [Fig. 6(a)]. Trajectories were generated from the transition matrix which is provided in the PyEMMA example datasets/models package[44] using a lag time of $\tau = 6$ steps. To produce non-equilibrium sampling, we start all trajectories from the left well. The trajectory lengths are 500 τ to 4000 τ, which is on the order of the mean-first-passage time to the right well. For each trajectory length, the aggregate data over all trajectories is 90 000 τ. Each run is repeated 100 times to compute means and uncertainties.

We compare VAMP with TICA in terms of the kinetic distance between the two energy wells. The kinetic distance is one of the few quantities that can be computed from both VAMP and TICA, whereas eigenvalues and singular values cannot directly be compared.

All estimates converge to the true reference value as the trajectory length is increased, when the sampling becomes increasingly representative of the equilibrium kinetics (Fig. 6). The true reference is computed with Eq. (12) and with ρ_1 set to the true stationary distribution of the model transition matrix. For non-equilibrium data (short trajectories), neither TICA nor VAMP reproduce the equilibrium kinetic distance. In VAMP, this is due to the weighting of points with respect to the empirical distributions ρ_1 which is, in general, different from the stationary distribution. Strikingly, the results from VAMP and TICA are almost identical, both in terms of the medians and of their statistical errors. This example indicates no particular advantage of using VAMP over using TICA but also no disadvantage.

Both VAMP and TICA can handle completely disconnected datasets (if transitions in both directions between a pair of states

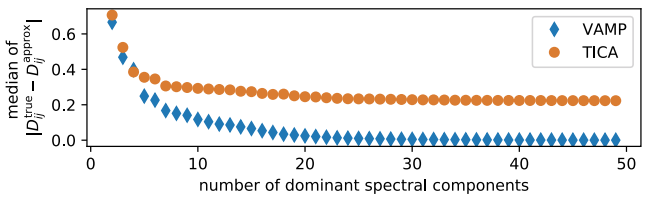

FIG. 5. Comparison of the median difference between the exact kinetic distance computed with Eq. (12) to its low-rank VAMP approximation and to its low-rank TICA approximation as a function on the retained number of spectral components. ρ_1 in Eq. (12) was set to the stationary distribution of the ASEP model.

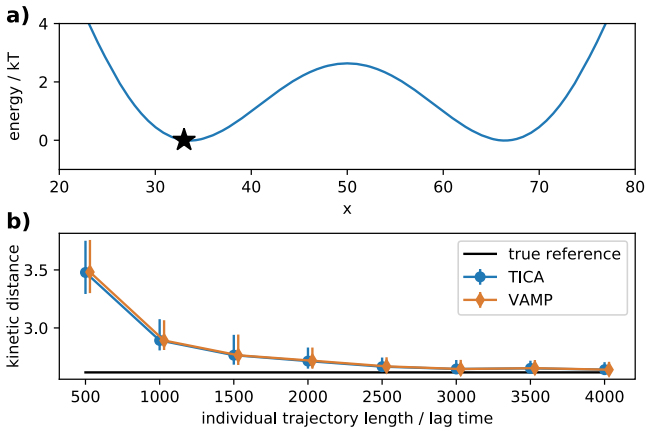

FIG. 6. (a) Double-well energy landscape (parameters from PyEMMA[44]) used to test VAMP in the limit of non-equilibrium sampling. Trajectories were all started from the minimum of the left well (star). (b) Kinetic distance between the two local minima of the energy landscape depending on the trajectory length used for its estimation with VAMP or TICA. Plot markers mark the median. Tips of the error bars mark the 10th and the 90th percentile.

are missing). Every disconnected set leads to an additional singular value/eigenvalue of value 1 (or close to one due to projection errors). However, the strength of VAMP lies elsewhere—in the analysis of inherently non-equilibrium systems such as driven ion motion as exemplified by the ASEP model.

C. Conformational changes of the KcsA potassium ion channel

Ion channels are pore-forming transmembrane proteins that enable ions to cross biomembranes. Ion channels are found both in the outer cell membrane and in the membranes of the cell organelles. They are important for functions such as cellular signaling, the regulation of osmotic activity, and the propagation of action potentials in nerves and muscle cells.[47]

The first potassium channel protein to be crystallized is the bacterial channel KcsA.[48] The structure can be subdivided into three consecutive parts: following the pore from the extracellular to the intracellular side, one finds (1) the selectivity filter, (2) a hydrophobic cavity, and (3) the intracellular gate. The selectivity filter [see Figs. 11(a) and 11(b)] is formed by a conserved Thr-Val-Gly-Tyr-Gly motif. The backbone carbonyls of this motif and the O_γ-atoms of the Thr side chains form five cubic cages each of which is able to coordinate one potassium ion. The structure of the selectivity filter found in KcsA is conserved even in eukaryotic channels. That is why KcsA acts as a general model system that is used to study potassium channel function.

Many channels can open and close their pore via a conformational change. This so-called gating takes place in a controlled way and can be provoked by the interaction of the pore-forming protein domain with other domains, other molecules, or in response to electric forces.[47,49] In the KcsA channel and its homologs, gating can take place via the intracellular gate or via conformational changes in the selectivity filter.[50] Here, we investigate the motions of the filter and their influence on conductance. The intracellular gate remains in the open state.

We reanalyze the non-equilibrium molecular dynamics simulation data of the KcsA channel protein that were previously published by Köpfer *et al.*[23] and consists of a total amount of 15.1 μs of MD simulation in 20 short trajectories with individual lengths ranging between 541.4 ns and 793.5 ns. In their simulations, a steady potassium ion current is maintained by the computational electrophysiology approach of Kutzner *et al.*[51] The simulations are therefore intrinsically nonreversible and the applications of methods that were developed for reversible dynamics, like TICA, are not justified.

In the following, we compute the VAMP components, define long-lived states in this space using PCCA+, and characterize the thus-obtained states.

1. Dynamic modes of the selectivity filter and surrounding residues

We compute the leading singular functions of the Koopman operator that describe the KcsA dynamics at a lag time of 40 ns using VAMP. The input features (ansatz functions) $\chi(t)$ for VAMP consist of two groups: (a) all inverse pairwise distances between heavy atoms of the selectivity filter (residues 75–79 in the first subunit using the numbering scheme of PDB file 1K4C[52] and their corresponding residues in the other three subunits) and (b) the inverse distance to the closest potassium cation for every heavy atom in the selectivity filter. This results in a total number of 7750 features. In the computation of distances, atoms that are symmetric under a rotation of the side chain dihedral by π are treated as one atom. In this analysis, this applies to the atom pairs $(C_{\delta1}, C_{\delta2})$ and $(C_{\varepsilon1}, C_{\varepsilon2})$ in tyrosine residues, $(C_{\gamma1}, C_{\gamma2})$ in valine residues, the pair $(O_{\delta1}, O_{\delta2})$ in aspartic acid residues, and the pair $(O_{\varepsilon1}, O_{\varepsilon2})$ in glutamic acid residues.

We discarded the first 18 ns of every trajectory. That is because in the first 18 ns, we observed conformational changes at the N-terminal end of the intracellular gate. We see these conformational changes at the beginning of every trajectory. We suspect that this might be due the pulling procedure that was used to prepare the open-gate conformation in Ref. 23.

The spectrum of singular values [Fig. 7(a)] shows jumps at positions 1, 2, 6, 7, 8, and 14 (not counting the constant singular value $\sigma_0 = 1$; see the supplementary material, Fig. 12) and become quasicontinuous afterward. We therefore restrict the analysis to the dynamics within the space of the leading 14 singular functions. To validate this decomposition, we perform the non-equilibrium Chapman-Kolmogorov test. Results show (see the supplementary material, Fig. 13) good agreement between estimates and

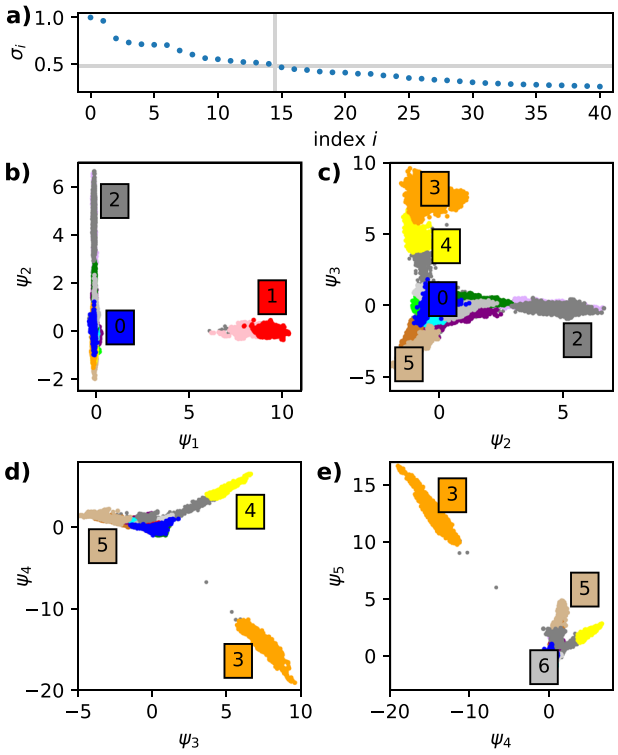

FIG. 7. (a) Leading 40 singular values obtained with VAMP. (b)–(e) Projection of the simulation data on pairs of singular functions (VAMP components). Data points were colored according to the macrostate to which they have the highest membership. Data points that do not clearly belong to any of the macrostates (maximum membership to any state <0.6) are shown as small gray points.

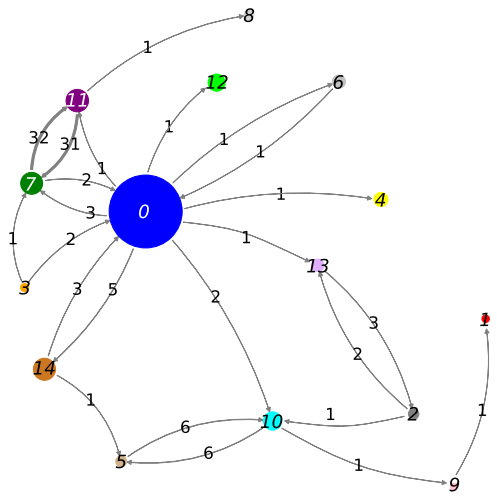

FIG. 8. Connectivity network for the 15 long-lived states that were identified with VAMP and PCCA+. Long-lived states are shown as disks with areas proportional to the frequency of the state in the MD simulation data. Macrostates that are kinetically connected by transitions in the data are connected by an arrow in this figure. Numbers on the arrows denote the number of transition events observed in the MD data. Numbers inside the disks are state labels.

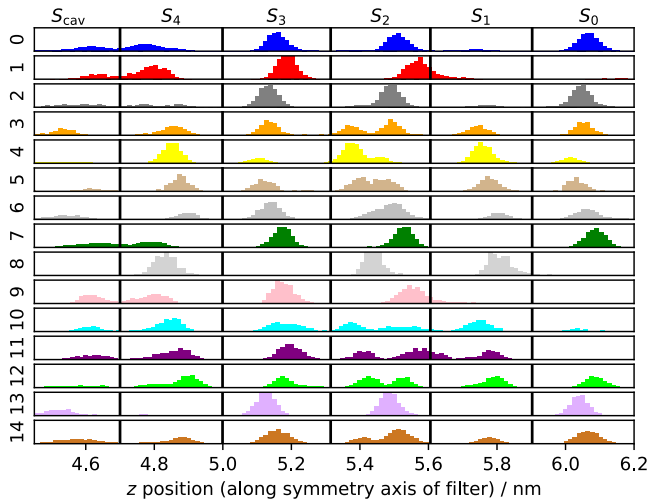

FIG. 9. Histogram of ion occupancy along the channel pore for each macrostate. Black vertical lines mark the positions of the carbonyl oxygen atoms that demarcate the ion binding sites S_{cav}, S_4, S_3, S_2, S_1, and S_0. In state 8 (light gray), the ion binding site S_3 is occupied by a water molecule. In states 1 (red), 9 (pink), and 10 (cyan), the ion binding site S_1 is occupied by water.

predictions for the fast processes (with smaller singular values) and deviations between predictions and estimates for the slower processes (with large singular values). As elaborated in the next paragraphs, the KcsA trajectories contain many unique transition events, which explain the failure of the dynamic model to provide accurate predictions of the long time scale kinetics.

2. Detection of long-lived states with PCCA+

Projections of the MD data points $\mathbf{x}$ (conformations) onto pairs of left singular functions ($\psi_i(\mathbf{x})$, $\psi_j(\mathbf{x})$) show that the data points form clearly separated clusters [see Figs. 7(b)–7(e)]. Such clustering has been observed for many other molecular systems and indicates the presence of long-lived states.[53] This motivates us to group the conformations into a small number of macrostates.

We assign the data points to 15 macrostates using the PCCA+ algorithm. We apply the PCCA+ variant of Ref. 53 to the data points in the space of the leading 14 singular functions (see methods Subsection III D and Appendixes C 2 and C 3). We observe that the macrostates defined with PCCA+ match well with the "density blobs" that one would assign intuitively by looking at the projections [see Figs. 7(b)–7(e)]. This indicates that the space of the singular functions is a suitable space for clustering with PCCA+.

3. Transitions between long-lived states and their populations

We compute the number of transitions between the macrostates using the mile-stoning method [also called transition-based

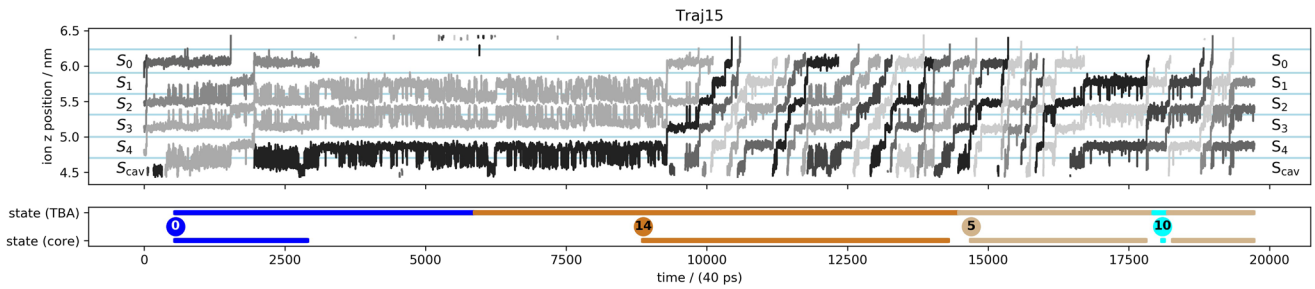

FIG. 10. Exemplary time series of potassium ion positions and assignment to macrostates. A pair of plots is shown, which share the time axis. The top plot shows the z positions of all ions in the selectivity filter. The bottom plot shows the macrostate visited at time t. Macrostates are color-coded, and the index of every state is shown in a circle on the first core entry. Two variants (core-based assignment and transition-based assignment, TBA) of macrostate assignments are shown. The trajectory of cores only shows frames where the conformation can be assigned with a high probability (membership) to a macrostate. Frames that were left unassigned in the core trajectory are assigned to the most recent or most proximate core in the TBA trajectories by splitting transitions at the midpoint.[5]

assignment (TBA) or core set approach; see Appendix C 4 and Refs. 5, 54, and 55]. The network of transitions between the macrostates (Fig. 8) shows that most transitions occur only once. States 0, 6, 7, 11, and 14 are in the reversibly connected set (ergodically visited macrostates). Most of the simulation data are assigned to macrostate number 0 (see Fig. 8).

The present MD simulation data do not allow us to make any statements about asymptotic state occupancies in the steady state equilibrium that might possibly be reached at 100s of microseconds and above. Most conformational changes observed in the MD data occur only in one direction. This might indicate a lack of sampling of state transitions in the short MD data, and longer MD simulation might reveal that the transitions are in fact reversible.

Inspection of the trajectories shows that transitions between the cores can take relatively long (e.g., see the transition from macrostate 0 to macrostate 14 in Fig. 10). For some of the states, the transition in/out of the state can take roughly the same amount of time that the system spends in the state. This may indicate either that the description of the dynamics requires more macrostates or that the approximation of the singular functions with VAMP is not accurate enough. (That is, there exists a better approximation that would lead to more metastable kinetics of the reduced model.)

Ion permeation (except for the blocked states) is faster than the life-times of the macrostates. The time between ion transition events is typically on the order of 10 ns, while dwell times of the macrostates are typically on the order of 100 ns (see the supplementary material, Table 2 and, for example, permeation in macrostate 5 and 14 in Fig. 10). Therefore, transitions between macrostates do not seem to describe the individual ion movement steps in the permeation mechanism. Rather the macrostates appear more related to the protein conformation (see Secs. IV C 4 and IV C 6).

4. Long-lived macrostates differ in the occupancy of the selectivity filter

We compute histograms of the ion occupancy of the selectivity filter (Fig. 9). One histogram is computed for each macrostate separately. The most frequent state (0, blue) and states 2, 7, and 13 have an evacuated ion binding site S_1 that is occupied neither by a potassium ion nor by a water molecule (see the supplementary material, Fig. 6). This means that ions do a long jump from S_2 to S_0 during conduction in these macrostates and S_1 is only visited transiently with a dwell time that is much shorter than the dwell time in S_2 and S_0 (see, for instance, the part of the trajectory that is assigned to core 0 in Fig. 10 and the supplementary material, Figs. 7–10). In other macrostates, e.g., state 11 (violet), S_1 is more frequently occupied. States 1 (red) and 9 (pink) show an ion binding site S_1 that is occupied with water. Furthermore, these states a characterized by a flipped Tyr78 conformation and a drastically distorted selectivity filter [see Fig. 11(1) and Sec. IV C 5]. No ion permeation events are observed in these states (see below). In state 10 (cyan), S_1 is partially occupied by water. Still, this state is conductive. In state 8 (light gray), the ion binding site S_3 is occupied by a water molecule. No ion conduction takes place in this state (see below).

5. Long-lived macrostates differ in ion conduction

We compute the potassium ion current through the pore by counting the net number of forward ion transitions from the S_1 to the S_0 binding site (see Appendix C 5 for details). The ion current for each macrostate is shown in Fig. 12. Computing the ion current from the transitions from S_4 to S_3 gave identical results. Ion current differs significantly between macrostates. However, it should be noted that there is only little simulation data for the different macrostates. For many states, less than 20 permeation events are observed (see also the supplementary material, Table 2). In states 1 (red), 8 (light gray), and 9 (pink), no permeation is observed. In states 1 and 9, the filter is in the Tyr78-flipped conformation. In state 8, a water molecule was threaded into the ion file. The most frequently visited state in the MD simulation (state number 0) conducts little compared to most other states. Since the free energies of

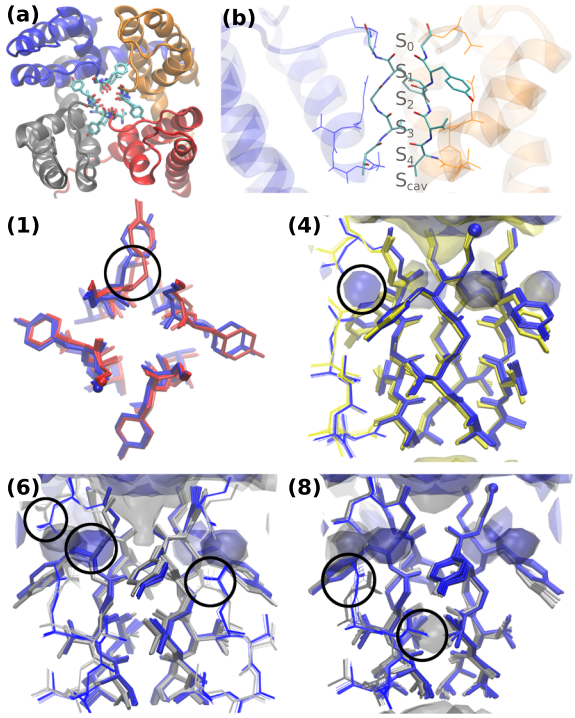

FIG. 11. (a) KcsA seen from the extracellular side, atoms of the selectivity filter are shown as sticks. (b) Cross section through the filter (side view). Filter atoms are shown as sticks, surrounding atoms are shown with lines. Ion binding sites are labeled S_{cav} through S_0. (1) View of the selectivity filter from the extracellular side. Structures (conformations) in blue are drawn from macrostate 0. Superimposed on that are structures from macrostate 1. A remarkable deviation from state 0 in the Tyr78 conformation is marked with a black circle. (4), (6), (8): Side view of the selectivity filter and surrounding amino acids. Structures (conformations) in blue are drawn from macrostate 0. Superimposed on that are structures from different macrostates (4, 6, 8). 20 structures are shown per macrostate where every conformation is an average over 50 conformations drawn randomly from the macrostate. Residues forming the pore of the selectivity filter are shown as sticks, residues surrounding the filter (including Glu71 and Asp80) with thin lines. Water density is shown as semitransparent isosurfaces. The tiny blue sphere marks the oxygen atom in Gly79 of the first subunit of the channel. Significant deviations (compared to state 0) in the Glu71-Asp80 contact and buried water presence are indicated with black circles. All states except 0 show disruption of the Glu71-Asp80 interaction in one or two subunits. This disruption can be accompanied by the absence of a buried water molecule like in state 4 (yellow), state 6 (light gray), and state 5 (tan, not shown).

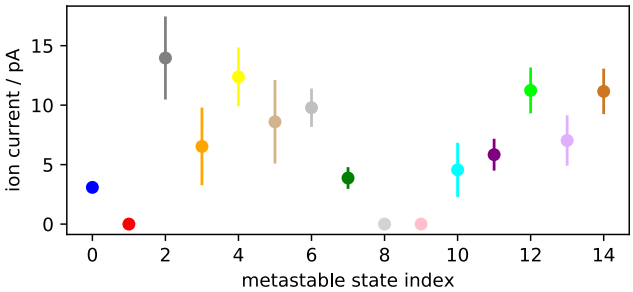

FIG. 12. Potassium ion current for each macrostate. Error bars show standard deviations and were computed by assuming that the number of permeation evens is Poisson-distributed (see Appendix C 5). In states 1, 8, and 9, no permeation events took place. (States 1 and 9 are in the Tyr78-flipped conformation. In state 8, a water molecule is located in S_3.) In the most frequent state 0, the ion current is relatively low.

the states cannot be computed from the present simulation data, no statement can be made about the contribution of the different states to the overall conduction.

6. Differences in Tyr78 conformation, Glu71-Asp80 interaction, and presence of buried water

To characterize the structural features of the macrostates, we sampled randomly with replacement 1000 conformations from every state. Average conformations for a subset of macrostates are shown in Fig. 11.

In states 1 (red) and 9 (pink), Tyr78 is in a flipped conformation (compared to state 0). The Tyr78 flip coincides with a disruption of the selectivity filter structure and zero conduction.

The Glu71-Asp80 contact[56] can either be formed or be broken in one or in two subunits of the channel. It is formed in all subunits in state 0. It is open in one subunit in states 2, 3, 4, 7, 8, 10, 12, and 14. The contact is open in two subunits in states 1, 5, 6, 9, 11, and 13 (see the supplementary material, Table 1). Opening and closing of the Glu71-Asp80 contact is a reversible process (since there are reversible transitions between states 0 and 6, 7, 11, 14). We observe opening of two Glu71-Asp80 contacts only in adjacent subunits of the channel which may hint to cooperativity.

In contrast to what is known about the KcsA channel,[56] opening of the Glu71-Asp80 contact in these MD simulations does not inactivate the channel. Instead ion current in the open-contact states is slightly increased (exceptions: 1, 8, 9; see explanations above) over the ion current in the closed-contact state 0.

The water molecule that is involved in the Glu71-Asp80 interaction can either be present or absent. We observe at most one absent water molecule. Breaking of the Glu71-Asp80 interaction does not strictly coincide with the absence of water.

In state 2, we observe a highly tilted Glu71 side chain conformation where the side chain points toward the filter pore. The tilt is larger than in the crystal structures and NMR structures of the KcsA protein that are available in the Protein Data Bank. This conformation might be stabilized by electrostatic attraction between the carboxyl groups of Glu71 and the potassium ions.

In summary, our analysis of the MD data of Köpfer *et al.*[23] reveals 15 long-lived states. While the most frequent state shows a crystal-like conformation of the selectivity filter, other states show flipping of the Tyr86 side chain or opening of the Glu71-Asp80 contact in one or more subunits. The different identified states display distinct ion conductances, establishing a direct link between channel function and the conformations identified with VAMP and PCCA+.

V. CONCLUSION

We have used the Variational Approach to Markov Processes (VAMP) to formulate a dynamical dimension reduction method for identifying the collective variables of the "slow" or "rare" processes in many-body systems. In this formulation, VAMP can replace the TICA method that is only defined for statistically reversible and stationary dynamics and, in practice, often only usable when the probability distribution sampled by the simulation trajectories is close to equilibrium.

We have applied VAMP-based dimension reduction to the asymmetric simple exclusion process toy model for single file ion diffusion and to non-equilibrium molecular dynamics data of the KcsA potassium channel protein. Both systems have high-dimensional state spaces and follow non-equilibrium dynamics that do not comply with the principle of detailed balance (microscopic reversibility). For both systems, we could construct a low-dimensional model that captures physically interesting processes.

We have demonstrated that VAMP is superior to TICA in correctly estimating kinetic distances for the intrinsically nonreversible ASEP model. Based on theoretical insights, we expect this to be true for any nonreversible system. For the analysis of non-equilibrium data that originates from simulating a reversible system with a non-equilibrium initial condition, we empirically showed that TICA and VAMP give similar results.

We have shown that the space of the leading singular functions is a suitable space for identification of long-lived macrostates even for the case of nonreversible dynamics. This was confirmed twice for the KcsA protein data: (1) the PCCA+ macrostates appear as well-separated density-clusters when projected to the space of the singular functions and (2) counting exit end entry events with the core-set (or transition-based) approach confirms that transitions between macrostates are rare events.

We proposed to extend the scope of the Chapman-Kolmogorov test from an application to probabilities[4] to general observables. We further proposed to use the singular functions as observables for the Chapman-Kolmogorov test. In fact it has been shown that the singular functions span the space of indicator functions for *coherent sets*.[29] Coherent sets are particular stable sets in time-space.[57] Examples for coherent sets are oceanic[57,58] or atmospheric[59] eddies. Reliably simulating their formation and dissolution should be equally challenging as sampling the exit from metastable states in systems with reversible dynamics. Testing whether a reduced dynamical model captures these rare events seems worthwhile.

SUPPLEMENTARY MATERIAL

See supplementary material for comparisons of VAMP to TICA and for comparisons of VAMP results obtained using a complete basis to VAMP results obtained using a reduced basis for the ASEP

model. Also see the supplementary material for more detailed properties of the long-lived KcsA states, of the MD trajectories, and of the Koopman model of KcsA filter motion.

ACKNOWLEDGMENTS

F.P. acknowledges funding from the Yen Post-Doctoral Fellowship in Interdisciplinary Research and from the National Cancer Institute of the National Institutes of Health (NIH) through Grant No. CAO93577. F.P. and F.N. acknowledge funding by European Commission (ERC CoG 772230 "ScaleCell"), Deutsche Forschungsgemeinschaft (No. SFB1114/C03). F.N. acknowledges funding by MATH+ (AA1-6). H.W. acknowledges funding from 1000-Talent Program of Young Scientists in China. B.L.d.G. and F.N. gratefully acknowledge financial support from the Deutsche Forschungsgemeinschaft (No. FOR 2518 "DynIon"). We are grateful to Brooke E. Husic and Erik H. Thiede for their corrections to the manuscript, to Simon Olsson for the intensive exchange about the modeling of many-particle systems, and to Cecilia Clementi for insightful discussions about the ASEP model.

The linear variational approach to Markov processes has been implemented in the publicly available PyEMMA software package http://emma-project.org.

APPENDIX A: PROOF OF (13)

The SVD of $\mathscr{K}_\tau$ is

$$\mathscr{K}_\tau g = \sum_i \sigma_i \langle g, \phi_i \rangle_{\rho_1} \psi_i, \tag{A1}$$

and then the transition density can be expressed as

$$p(\mathbf{z}|\mathbf{x}) = \mathscr{K}_\tau \delta_{\mathbf{z}}(\mathbf{x}), \tag{A2}$$

$$= \sum_i \sigma_i \psi_i(\mathbf{x}) \phi_i(\mathbf{z}) \rho_1(\mathbf{z}). \tag{A3}$$

Considering the orthonormality of singular functions, we have

$$\begin{aligned} D_\tau^2(\mathbf{x},\mathbf{y}) &= \int \frac{\left(\sum_i \sigma_i (\psi_i(\mathbf{x}) - \psi_i(\mathbf{y}))\phi_i(\mathbf{z})\rho_1(\mathbf{z})\right)^2}{\rho_1(\mathbf{z})} \mathrm{d}\mathbf{z} \\ &= \sum_{i,j} \sigma_i \sigma_j (\psi_i(\mathbf{x}) - \psi_i(\mathbf{y}))(\psi_j(\mathbf{x}) - \psi_j(\mathbf{y}))\langle \phi_i, \phi_j \rangle_{\rho_1} \\ &= \sum_i \sigma_i^2 (\psi_i(\mathbf{x}) - \psi_i(\mathbf{y}))^2. \end{aligned} \tag{A4}$$

APPENDIX B: MODELS

1. Koopman matrix for the ASEP model

Let $\wedge$ denote the bitwise AND operator. Let $\mathbf{L} \in \mathbb{R}^{2^N \times 2^N}$. For all $0 \le i < 2^N$, $0 \le j < 2^N$, $i \ne j$, let

$$L_{ij} = \alpha \quad \text{if } i \wedge 1 = 0 \text{ and } j \wedge 1 = 1, \tag{B1}$$

$$L_{ij} = \beta \quad \text{if } i \wedge 2^{N-1} = 1 \text{ and } j \wedge 2^{N-1} = 0, \tag{B2}$$

$$L_{ij} = p \quad \text{if } \exists\, 0 \le k < N-1 : i \wedge 2^k = 1 \text{ and } i \wedge 2^{k+1} = 0 \text{ and } j \wedge 2^k = 0 \text{ and } j \wedge 2^{k+1} = 1, \tag{B3}$$

$$L_{ij} = q \quad \text{if } \exists\, 0 \le k < N-1 : i \wedge 2^k = 0 \text{ and } i \wedge 2^{k+1} = 1 \text{ and } j \wedge 2^k = 1 \text{ and } j \wedge 2^{k+1} = 0, \tag{B4}$$

$$L_{ij} = 0 \quad \text{otherwise}, \tag{B5}$$

and $L_{ii} = -\sum_{j \ne i} L_{ij}$.

The model transition matrix $\mathbf{T}_\tau \in \mathbb{R}^{2^N \times 2^N}$ is computed by taking the matrix exponential of $\tau\mathbf{L}$, where τ is the lag time. The full Koopman operator $\mathscr{K}_\tau$ is finite-dimensional for this model and is identical to $\mathbf{T}_\tau$

APPENDIX C: METHODS

1. Implementation of the non-equilibrium Chapman-Kolmogorov test

In the Chapman-Kolmogorov test, the estimate of the time-lagged cross correlation $\mathrm{cov}_{\mathrm{est}}(f, g; n\tau)$ and its model prediction $\mathrm{cov}_{\mathrm{pred}}(f, g; n\tau)$ are compared.

Using Eq. (11), it is possible to express the covariance at the unit lag time $1 \times \tau$ as

$$\mathrm{cov}_{\mathrm{pred}}(f,g;\tau) = \mathrm{cov}_{\mathrm{est}}(f,g;\tau), \tag{C1}$$

$$= \langle f, \mathscr{K}_\tau g \rangle_{\rho_0}, \tag{C2}$$

$$= \sum_i \langle g, \phi_i \rangle_{\rho_1} \sigma_i \langle \psi_i, f \rangle_{\rho_0}, \tag{C3}$$

$$= \mathbf{q}^\top \mathbf{r}, \tag{C4}$$

where we have defined $q_i := \langle g, \phi_i \rangle_{\rho_1}$ and $r_i := \sigma_i \langle \psi_i, f \rangle_{\rho_0}$.

By combining Eqs. (11) and (5), the prediction for higher multiples $n > 1$ of the lag time of can be computed as

$$\mathrm{cov}_{\mathrm{pred}}(f, g; n\tau) := \langle f, \mathscr{K}_\tau^n g \rangle_{\rho_0} = \mathbf{q}^\top \mathbf{P}^{n-1} \mathbf{r}, \tag{C5}$$

where $P_{ij} := \sigma_i \langle \psi_i, \phi_j \rangle_{\rho_1}$. The quantities $\mathbf{P}$, $\mathbf{q}$, and $\mathbf{r}$ can all be computed from the data and from the spectral quantities that VAMP provides an approximation for.

2. Computing the metastable memberships

We use the "inner simplex" algorithm of the PCCA+ method[36] to compute the linear map $\mathbf{A}$ from the space of singular functions to the space of macrostate memberships. The "inner simplex" algorithm was motivated by the observation that reversible metastable systems show a clustering of data points close to the N most distant points in the space of the dominant N eigenfunctions (counting the constant eigenfunction). This so-called "simplex structure" forms the basis for many spectral clustering algorithms.[36]

The algorithm in the version of Ref. 36 consist of two stages: (1) localizing the N most distant points (the *vertices*) $\{\boldsymbol{\psi}_1^{(\mathrm{ex})}, \ldots, \boldsymbol{\psi}_{N-1}^{(\mathrm{ex})}\}$ in the $N-1$-dimensional space of the dominant eigenfunctions (excluding the constant eigenfunction) and (2) computing barycentric coordinates for every point $\boldsymbol{\psi}(t)$ with respect to the vertices by solving the following equations for $m_i(t)$

$$\boldsymbol{\psi}(t) = \sum_{i=1}^{N} m_i(t)\boldsymbol{\psi}_i^{(\mathrm{ex})}, \tag{C6}$$

$$1 = \sum_{i=1}^{N} m_i(t). \tag{C7}$$

The solution of this linear problem implicitly defines the linear map **A** from $\boldsymbol{\psi}$ to **m**.

If the data points $\{\boldsymbol{\psi}(t)\}_t$ indeed form clusters close to the vertices, the coefficient $m_i(t)$ can be understood as the membership of point $\boldsymbol{\psi}(t)$ in the macrostate number i. Here, we apply the "inner simplex" algorithm not in the space of the eigenfunctions of the MSM transition matrix as initially suggested in Ref. 36 but in the space of the singular functions.

Note that in its conventional use, PCCA+ is applied to cluster the space of MSM eigenvectors. For MSMs, only one representative value of the eigenfunctions is needed for every microstate because the approximations to the eigenfunctions are constant on every microstate by definition. This is different in our application of PCCA+ to the continuous order parameters computed with VAMP. Hence, all data points have to be used here.

3. Defining macrostates and core sets

In order to interpret the macrostates that originate from VAMP and PCCA+, we investigate representative molecular conformations from every macrostate. To this, we first define the core set C_i of the macrostate number i as

$$C_i := \{\mathbf{x}(t) \mid i = \mathrm{argmax}_j m_j(t) \text{ and } m_i(t) \geq f\}, \tag{C8}$$

where $0.5 \leq f \leq 1$ is some arbitrary cutoff on the memberships. We chose $f = 0.6$. From each core set, we draw 1000 random samples of molecular conformations with replacement. A subset of these conformations is shown Fig. 11.

4. Counting transitions and finding the largest connected set of macrostates

We count transitions at a lag time of $\tau = 40$ ps according to the transition-based assignment (TBA) algorithm or the mile-stoning algorithm.[5,54] In the TBA algorithm, every conformation $\mathbf{x}(t)$ is first assigned to the either the last core set that was hit by the trajectory or the next core set that will hit by the trajectory, whichever is closer in time. We thus obtain a sequence of core labels $\{s(t)\}_{t=0,\ldots,T}$, $s(t) \in \mathbb{N}$ for every MD trajectory. For every trajectory, we compute a count matrix **c** from $\{s(t)\}_t$ using the standard approach[4] as follows:

$$c_{ij} = \sum_{t=0}^{T-\tau} \delta_{is(t)}\delta_{js(t+\tau)}, \tag{C9}$$

where δ_{ij} is the Kronecker delta. The count matrix for all trajectories **C** is computed by summing the individual count matrices of each trajectory.

The largest connected set of macrostates[4] is computed from **C** and consists of the five states 0, 6, 7, 11, and 14.

5. Computing the ion current

We estimate the potassium ion current by computing the number of times some ion transitions from binding site S_1 to binding site S_0 of the selectivity filter. We estimate the number of transitions using a core-set approach.[5] The core region of the S_1 site is defined as $5.7\ \mathrm{nm} \leq z_{\mathrm{ion}} \leq 5.85\ \mathrm{nm}$, and the core region of the S_0 size is defined as $6.0\ \mathrm{nm} \leq z_{\mathrm{ion}} \leq 6.14\ \mathrm{nm}$ (in the coordinate system of the MD data from Ref. 23). For all trajectory segments that are assigned to macrostate i, we compute n_i the number of transitions from the core of S_1 to the core of S_0 minus the number of reverse transitions (summing up the number of transitions of all the ions). The error Δn_i of n_i is computed as $\sqrt{n_i}$ by assuming that n_i is Poisson distributed. The ion current is computed as $I_i = \frac{en_i}{\Delta t_i}$, where e is the elementary charge and Δt_i is the length of the trajectory pieces assigned to macrostate i. The error is estimated as $\Delta I_i = \frac{e\Delta n_i}{\Delta t_i}$.

REFERENCES

[1] C. Schütte, A. Fischer, W. Huisinga, and P. Deuflhard, J. Comput. Phys. **151**, 146 (1999).
[2] N. Singhal, C. D. Snow, and V. S. Pande, J. Chem. Phys. **121**, 415 (2004).
[3] F. Noé, I. Horenko, C. Schütte, and J. C. Smith, J. Chem. Phys. **126**, 155102 (2007).
[4] 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).
[5] N.-V. Buchete and G. Hummer, J. Phys. Chem. B **112**, 6057 (2008).
[6] M. A. Rohrdanz, W. Zheng, M. Maggioni, and C. Clementi, J. Chem. Phys. **134**, 124116 (2011).
[7] J. Preto and C. Clementi, Phys. Chem. Chem. Phys. **16**, 19181 (2014).
[8] W. Zheng, B. Qi, M. A. Rohrdanz, A. Caflisch, A. R. Dinner, and C. Clementi, J. Phys. Chem. B **115**, 13065 (2011).
[9] M. A. Rohrdanz, W. Zheng, and C. Clementi, Annu. Rev. Phys. Chem. **64**, 295 (2013).
[10] B. Peters and B. L. Trout, J. Chem. Phys. **125**, 054108 (2006).
[11] F. Noé and C. Clementi, Curr. Opin. Struct. Biol. **43**, 141 (2017).
[12] C. Schütte, W. Huisinga, and P. Deuflhard, "Transfer operator approach to conformational dynamics in biomolecular systems," Technical Report No. SC 99-36, Konrad-Zuse-Zentrum für Informationstechnik, Berlin-Dahlem, Germany, 1999.
[13] F. Noé and F. Nüske, Multiscale Model. Simul. **11**, 635 (2013).
[14] F. Nüske, B. G. Keller, G. Pérez-Hernández, A. S. J. S. Mey, and F. Noé, J. Chem. Theory Comput. **10**, 1739 (2014).
[15] L. Molgedey and H. G. Schuster, Phys. Rev. Lett. **72**, 3634 (1994).
[16] G. Pérez-Hernández, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, J. Chem. Phys. **139**, 015102 (2013).
[17] C. R. Schwantes and V. S. Pande, J. Chem. Theory Comput. **9**, 2000 (2013).
[18] H. Wu, F. Nüske, F. Paul, S. Klus, P. Koltai, and F. Noé, J. Chem. Phys. **146**, 154104 (2017).
[19] H. Wu and F. Noé, "Variational approach for learning Markov processes from time series data," e-print arXiv:1707.04659 (2017).
[20] B. L. R. De Moor and G. H. Golub, "The restricted singular value decomposition: Properties and applications," Technical Report No. MA-89-03, Department of Computer Science, Stanford University, Stanford, CA, USA, 1989.
[21] H. Abdi, 'Singular value decomposition (SVD) and generalized singular value decomposition (GSVD)," in *Encyclopedia of Measurement and Statistics* (SAGE Publications, Thousand Oaks, CA, USA, 2007), pp. 907–912.
[22] A. Mardt, L. Pasquali, H. Wu, and F. Noé, Nat. Commun. **9**, 4443 (2018).
[23] D. A. Köpfer, C. Song, T. Gruene, G. M. Sheldrick, U. Zachariae, and B. L. de Groot, Science **346**, 352 (2014).
[24] F. Noé, S. Doose, I. Daidone, M. Löllmann, M. Sauer, J. D. Chodera, and J. C. Smith, Proc. Natl. Acad. Sci. U. S. A. **108**, 4822 (2011).
[25] I. Mezić, Nonlinear Dyn. **41**, 309 (2005).
[26] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz, J. Comput. Dyn. **1**, 391 (2014).

[27]F. Noé and C. Clementi, J. Chem. Theory Comput. **11**, 5002 (2015).
[28]F. Noé, R. Banisch, and C. Clementi, J. Chem. Theory Comput. **12**, 5620 (2016).
[29]P. Koltai, H. Wu, F. Noé, and C. Schütte, Computation **6**, 22 (2018).
[30]A. Hyvärinen, J. Karhunen, and E. Oja, in *Independent Component Analysis* (John Wiley & Sons, New York, USA, 2001), Chap. 6.4, pp. 140–141.
[31]T. R. Knapp, Psychol. Bull. **85**, 410 (1978).
[32]R. T. McGibbon and V. S. Pande, J. Chem. Phys. **142**, 124105 (2015).
[33]This sum is by definition the r'th power of the so-called r-Schatten norm $\|.\|_r^r$.
[34]F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl, Proc. Natl. Acad. Sci. U. S. A. **106**, 19011 (2009).
[35]P. Deuflhard and M. Weber, Linear Algebra Appl. **398**, 161 (2005), special Issue on Matrices and Mathematical Biology.
[36]M. Weber and T. Galliat, "Characterization of transition states in conformational dynamics using fuzzy sets," Technical Report No. 02-12, Konrad-Zuse-Zentrum für Informationstechnik Berlin, Berlin-Dahlem, Germany, 2002.
[37]S. Stolzenberg, Bioinformatics **2018**, bty818.
[38]C. T. MacDonald, J. H. Gibbs, and A. C. Pipkin, Biopolymers **6**, 1 (1968).
[39]F. Spitzer, Adv. Math. **5**, 246 (1970).
[40]O. Golinelli and K. Mallick, J. Phys. A: Math. Gen. **39**, 12679 (2006).
[41]A. B. Kolomeisky, G. M. Schütz, E. B. Kolomeisky, and J. P. Straley, J. Phys. A **31**, 6911 (1998).
[42]F. Nüske, R. Schneider, F. Vitalini, and F. Noé, J. Chem. Phys. **144**, 054105 (2016).
[43]C. R. Schwantes and V. S. Pande, J. Chem. Theory Comput. **11**, 600 (2015).
[44]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).
[45]F. Nüske, H. Wu, J.-H. Prinz, C. Wehmeyer, C. Clementi, and F. Noé, J. Chem. Phys. **146**, 094104 (2017).
[46]G. R. Bowman, K. A. Beauchamp, G. Boxer, and V. S. Pande, J. Chem. Phys. **131**, 124101 (2009).
[47]*Molecular Biology of the Cell*, edited by B. Alberts, A. Johnson, J. Lewis, D. Morgen, M. Raff, K. Roberts, and P. Walter (Garland Science, Taylor & Francis Group, New York, USA, 2008), Chap. 11, pp. 611–640.
[48]D. A. Doyle, J. M. Cabral, R. A. Pfuetzner, A. Kuo, J. M. Gulbis, S. L. Cohen, B. T. Chait, and R. MacKinnon, Science **280**, 69 (1998).
[49]G. Yellen, Nature **419**, 35 (2002).
[50]T. Hoshi and C. M. Armstrong, J. Gen. Physiol. **141**, 151 (2013).
[51]C. Kutzner, H. Grubmüller, B. L. de Groot, and U. Zachariae, Biophys. J. **101**, 809 (2011).
[52]Y. Zhou, J. H. Morais-Cabral, A. Kaufman, and R. MacKinnon, Nature **414**, 43 (2001).
[53]M. Weber, "Clustering by using a simplex structure," Technical Report No. 04-03, Konrad-Zuse-Zentrum für Informationstechnik Berlin, Berlin-Dahlem, Germany, 2004.
[54]C. Schütte, F. Noe, J. Lu, M. Sarich, and E. Vanden-Eijnden, J. Chem. Phys. **134**, 204105 (2011).
[55]A. Jain and G. Stock, J. Chem. Theory Comput. **8**, 3810 (2012).
[56]J. F. Cordero-Morales, L. G. Cuello, Y. Zhao, V. Jogini, D. M. Cortes, B. Roux, and E. Perozo, Nat. Struct. Mol. Biol. **13**, 311 (2006).
[57]R. Banisch and P. Koltai, Chaos **27**, 035804 (2017).
[58]G. Froyland, K. Padberg, M. H. England, and A. M. Treguier, Phys. Rev. Lett. **98**, 224503 (2007).
[59]N. Santitissadeekorn, G. Froyland, and A. Monahan, Phys. Rev. E **82**, 056311 (2010).