

Estimation and uncertainty of reversible Markov models

Benjamin Trendelkamp-Schroer, Hao Wu, Fabian Paul, and Frank Noé

Citation: *The Journal of Chemical Physics* **143**, 174101 (2015); doi: 10.1063/1.4934536

View online: <http://dx.doi.org/10.1063/1.4934536>

View Table of Contents: <http://scitation.aip.org/content/aip/journal/jcp/143/17?ver=pdfcov>

Published by the [AIP Publishing](#)

Articles you may be interested in

[Efficient Bayesian estimation of Markov model transition matrices with given stationary distribution](#)

J. Chem. Phys. **138**, 164113 (2013); 10.1063/1.4801325

[Bayesian uncertainty quantification and propagation in molecular dynamics simulations: A high performance computing framework](#)

J. Chem. Phys. **137**, 144103 (2012); 10.1063/1.4757266

[Improved coarse-graining of Markov state models via explicit consideration of statistical uncertainty](#)

J. Chem. Phys. **137**, 134111 (2012); 10.1063/1.4755751

[Markov state models based on milestoning](#)

J. Chem. Phys. **134**, 204105 (2011); 10.1063/1.3590108

[Probability distributions of molecular observables computed from Markov models. II. Uncertainties in observables and their time-evolution](#)

J. Chem. Phys. **133**, 105102 (2010); 10.1063/1.3463406



AIP | APL Photonics

APL Photonics is pleased to announce
Benjamin Eggleton as its Editor-in-Chief



Estimation and uncertainty of reversible Markov models

 Benjamin Trendelkamp-Schroer,^{a)} Hao Wu,^{a)} Fabian Paul, and Frank Noé^{b)}
Institut für Mathematik und Informatik, FU Berlin, Arnimallee 6, 14195 Berlin, Germany

(Received 19 July 2015; accepted 12 October 2015; published online 2 November 2015)

Reversibility is a key concept in Markov models and master-equation models of molecular kinetics. The analysis and interpretation of the transition matrix encoding the kinetic properties of the model rely heavily on the reversibility property. The estimation of a reversible transition matrix from simulation data is, therefore, crucial to the successful application of the previously developed theory. In this work, we discuss methods for the maximum likelihood estimation of transition matrices from finite simulation data and present a new algorithm for the estimation if reversibility with respect to a given stationary vector is desired. We also develop new methods for the Bayesian posterior inference of reversible transition matrices with and without given stationary vector taking into account the need for a suitable prior distribution preserving the meta-stable features of the observed process during posterior inference. All algorithms here are implemented in the PyEMMA software — <http://pyemma.org> — as of version 2.0. © 2015 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution 3.0 Unported License. [<http://dx.doi.org/10.1063/1.4934536>]

I. INTRODUCTION

Markov models, Markov state models (MSMs), or master-equation models are a powerful framework to reduce the great complexity of bio-molecular dynamics to a simple kinetic description that represents the underlying transitions between distinct conformations.^{1–7} These models allow us to analyze the longest-living (metastable) sets of structures,⁸ the effective transition rates between them,^{9,10} the kinetic relaxation processes and their relationship to equilibrium kinetics experiments,^{7,11–14} and the thermodynamics and kinetics over multiple thermodynamic states.^{15–18} A key advantage of MSMs is that they are estimated from conditional transition statistics between states, and they thus do not require the data to be in global equilibrium across all states. As a result, they are an excellent tool to integrate the data of multiple simulation trajectories that have been run independently and from different initial states into a single informative model.^{19–21}

A variety of complex molecular processes have been successfully described using MSMs. Examples include the folding of proteins into their native folded structure,^{20,22,23} the dynamics of natively unstructured proteins,^{24,25} and the binding of a ligand to a target protein.^{21,26–30}

There are two key steps in the construction of a MSM. At first, a suitable discretization of the continuous conformation space has to be obtained. In most cases, no good *a priori* discretization is known and the discretization has to be found based on the simulation data. The appropriate choice of discretization is a topic of ongoing research.^{24,31,32} The error incurred by the discretization and by the subsequent approximation of the jump-process as a Markov process can be systematically controlled and evaluated.^{7,33–35}

In the second step, one estimates the transition probabilities between pairs of states based on the transition statistics. The most common approach to estimating Markov models from data is by means of a Bayesian framework. One first harvests the transition counts, $c_{ij}(\tau)$ from the data, i.e., how often trajectories were found in discrete state i at some time t and in discrete state j at some later time $t + \tau$. The parameter τ is called lag time and is crucial for the quality of the Markov model.^{7,33} Next, one computes the transition matrix either by maximizing the likelihood, i.e., the probability over all possible Markov model transition (or rate) matrices that may have generated the observed transition counts^{7,36,37} or by sampling Markov models from the posterior distribution.^{38–42} A maximum likelihood estimate gives a single-point estimate, i.e., a single Markov model that is “most representative” given the data. However, if some transition events are rare compared to the total simulation length — and this is the typical case in molecular dynamics simulation — this maximum likelihood model might be very uncertain and thus far away from the model that the one would converge to by increasing amount of simulation data. The Bayesian posterior ensemble is a natural approach to quantify such statistical uncertainties and thus to make meaningful comparisons between a Markov model obtained from different sets of simulations or to experimental data.

A key property of molecular dynamics at thermal equilibrium, and a necessary consequence of the second law of thermodynamics, is microscopic reversibility of the equations of motion. This property is ensured by many simulation procedures^{43,44} and carries over to detailed balance between discrete states, i.e., formally leads to a (time-) reversible Markov model. A reversible Markov model is not only physically desirable, it offers statistical advantages as it has only about half as many independent parameters compared to a nonreversible model, and it allows the equilibrium kinetics to be analyzed in a straightforward and meaningful manner. Furthermore,

^{a)}B. Trendelkamp-Schroer and H. Wu contributed equally to this work.

^{b)}Electronic mail: frank.noë@fu-berlin.de



imposing detailed balance with respect to a given stationary vector can be used to aid the efficient estimation of rare-event processes from MSMs.⁴⁵

Algorithms imposing detailed balance during likelihood maximization have been discussed in Refs. 7, 16, 36, and 37. First methods for sampling the posterior distribution of reversible transition matrices have been suggested in Ref. 39 and later in Ref. 46. A method working with natural priors for reversible chains was proposed in Ref. 47. The sampling of transition matrices reversible with respect to a fixed stationary distribution was also presented in Ref. 39, while a Gibbs sampling algorithm with a significantly improved convergence rate has been developed in Ref. 42. Reference 48 discusses methods for goodness-of-fit tests for Markov chains.

This article is deliberately broad and presents new concepts, insights, and algorithms for reversible Markov model estimation in general, maximum likelihood estimators, and Bayesian estimators that mutually benefit from each other. For this reason, we first give a survey of principles and consequences of reversible Markov models. We then extend the framework of maximum likelihood estimation of transition matrices by giving a simplified maximum likelihood estimator (MLE) for reversible transition matrices and a new estimator for reversible transition matrices with a fixed equilibrium distribution. The main part of the paper comprises new algorithms for the full Bayesian analysis of the posterior ensemble of reversible Markov models. As yet, three fundamental problems have not been satisfactorily solved: (i) How can one harvest statistically uncorrelated transition counts from trajectories in which subsequent transitions are correlated, so as to give rise to meaningful uncertainty intervals? (ii) Which prior should be used in a Bayesian analysis so as to get error intervals that envelop the true value even for Markov models with many states? (iii) How can we design efficient sampling algorithms for the reversible posterior ensemble, i.e., algorithms that allow to quickly compute reliable error bars for Markov models with many states? In this paper we discuss (i) and give a rather complete treatment of problems (ii) and (iii). Efficient sampling algorithms are derived for reversible Markov models and reversible Markov models with fixed equilibrium distribution.

II. REVERSIBLE MARKOV MODELS

In this section, we will show that microscopic reversibility carries over to the discretized situation and discuss the desirable properties of a reversible Markov state model.

A. From microscopic reversibility to discrete-state detailed balance

Let $\mu(x)$ denote the equilibrium distribution on the microscopic degrees of freedom $x \in \Omega$, e.g., all-atom coordinates of the system of interest, and let $p_\tau(x, y)$ denote the conditional transition density of the MD implementation. $p_\tau(x, y)$ is the probability that the system is found in state y at time $t + \tau$ given that it has been in state x at time t . The MD implementation fulfills microscopic reversibility if the following detailed balance

equation

$$\mu(x) p_\tau(x, y) = \mu(y) p_\tau(y, x) \quad (1)$$

holds for all pairs of states $x, y \in \Omega$. Hence, the terms “detailed balance” and “reversible” are equivalent in our context. Since $\mu(x) p_\tau(x, y)$ is the unconditional probability to find the transition $(x, t) \rightarrow (y, t + \tau)$, Eq. (1) means that the system is on average time-reversible — the absolute number of transitions from x to y is equal to the reverse. Microscopic detailed balance is desirable to have in any MD implementation when the aim is to perform simulations in thermodynamic equilibrium. If (1) would be violated, that would imply the existence of cycles $x \rightarrow y \rightarrow z \rightarrow x$ along which there is a net transport. Since such cycles could be used to generate work, their existence in a system that is driven by purely thermal energy would be inconsistent with the second law of thermodynamics.

Dynamical models that are commonly employed in MD implementations fulfill detailed balance. Brownian (overdamped Langevin) dynamics fulfills detailed balance. Hamiltonian and non-overdamped Langevin dynamics fulfill generalized detailed balance with respect to momentum inversion in phase space, but when integrating over the distribution of momenta they do fulfill ordinary detailed balance in position space.⁴⁹ In practice, some finite time-stepping integrators do not obey exact detailed balance with respect to the Boltzmann distribution, but we here consider that the MD implementation has been chosen such that detailed balance is at least approximately fulfilled.

Now suppose that the state space Ω is partitioned into non-overlapping subsets S_1, \dots, S_n that we shall call *discrete states* here. Each set has an equilibrium probability given by

$$\pi_i = \int_{S_i} dx \mu(x) \quad (2)$$

and the transition density gives rise to a discrete state transition matrix $P(\tau)$ with entries

$$p_{ij}(\tau) = \frac{\int_{S_i} dx \int_{S_j} dy \mu(x) p_\tau(x, y)}{\int_{S_i} dx \mu(x)}. \quad (3)$$

Using (2) and (3) and microscopic detailed balance, (1), it is straightforward to verify that

$$\pi_i p_{ij}(\tau) = \pi_j p_{ji}(\tau). \quad (4)$$

Note that (4) holds independently of the choice of the lag time τ . Moreover, (4) implies that π is the equilibrium probability vector of $P(\tau)$. By defining the diagonal matrix $\Pi = \text{diag}(\pi_1, \dots, \pi_n)$, we can alternatively write (4) as a matrix equation,

$$\Pi P = (\Pi P)^\top, \quad (5)$$

$$P = \Pi^{-1} P^\top \Pi. \quad (6)$$

As a result, if the microscopic dynamics are reversible, the Markov model transition matrix must also be reversible. However, a direct estimate of P from a finite amount of simulation data cannot be expected to fulfill (4) exactly. Thus, the principle validity of detailed balance motivates us to *enforce* (4) in the process of estimating P .

TABLE I. Number of independent variables (degrees of freedom, dof) and their approximated values for transition matrices depending on the constraints.

Constraints	dof	\approx
None	$n(n-1)$	n^2
Reversible	$\frac{1}{2}n(n-1) + n - 1$	$\frac{1}{2}n^2$
Reversible, fixed π	$\frac{1}{2}n(n-1)$	$\frac{1}{2}n^2$

Enforcing detailed balance helps to reduce the statistical error of an estimator for P because it reduces the number of independent variables roughly by approximately one half (see Table I). However, there are other consequences of having (4): Given detailed balance, we can compute the molecular equilibrium kinetics in a physically meaningful way and employ some useful analysis tools that are not defined for nonreversible Markov models. Moreover, we can employ more efficient and robust matrix algebra routines when exploiting that P is a reversible matrix.

B. Eigenvalues and eigenvectors

Many methods to analyze the molecular kinetics based on a Markov model transition matrix rely on the eigenvalue decomposition of P . Using the diagonal eigenvalue matrix $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$, we can formulate a right eigenvalue problem with right column eigenvectors $R = (r_1, \dots, r_n)$, $r_i \in \mathbb{R}^n$, and left row eigenvectors $L = (l_1, \dots, l_n)^\top$,

$$PR = R\Lambda, \quad (7)$$

$$LP = \Lambda L. \quad (8)$$

From (7), we can obtain a generalized eigenvalue problem,

$$\Pi PR = \Pi R\Lambda.$$

Π is symmetric positive definite and as a result of detailed balance, ΠP is symmetric. Hence, all eigenvalues $\lambda_1, \dots, \lambda_n$ are real, and the eigenvectors are orthogonal with respect to the equilibrium distribution⁵⁰

$$\langle r_i, r_j \rangle_\pi \propto \delta_{ij},$$

where we have used the weighted scalar product $\langle u, v \rangle_\pi = \sum_i \pi_i u_i v_i$. We can make R orthonormal by scaling an arbitrarily obtained eigenvector r_i by $\langle r_i, r_i \rangle_\pi^{-1/2}$.

Inserting detailed balance formulation (6) into decomposition (7) immediately gives

$$P^\top \Pi R = \Pi R\Lambda, \\ (\Pi R)^\top P = \Lambda (\Pi R)^\top$$

which is a left eigenvalue problem with the choice

$$L = (\Pi R)^\top, \\ l_i^\top = \Pi r_i. \quad (9)$$

Thus, detailed balance establishes a 1-to-1 relation between the left and the right eigenvectors. We can decompose the transition matrix into its spectral components by just using one

set of eigenvectors and the equilibrium distribution, such as

$$P = R\Lambda R^\top \Pi = \sum_{i=1}^n \lambda_i r_i r_i^\top \Pi. \quad (10)$$

1. Example 1

Consider the following reversible 3×3 transition matrix:

$$P = \begin{pmatrix} 0.5 & 0.34 & 0.16 \\ 0.28 & 0.5 & 0.22 \\ 0.15 & 0.25 & 0.6 \end{pmatrix}. \quad (11)$$

Suppose we generate a Markov chain of length 20 starting from state 1, resulting in the count matrix at lag $\tau = 1$,

$$C = \begin{pmatrix} 4 & 3 & 0 \\ 1 & 4 & 3 \\ 1 & 1 & 2 \end{pmatrix}. \quad (12)$$

Now we conduct a nonreversible and a reversible maximum likelihood estimation of the transition matrix given C . Eigenvalues for the exact transition matrix in (11) and both nonreversible and reversible estimates for the given count matrix in (12) are shown in Fig. 1.

It is seen that the nonreversible estimate contains complex eigenvalues. These generally come in complex conjugate pairs. Fig. 1 shows a much higher accuracy of the reversible estimate compared to the nonreversible estimate. In order to explore the statistical significance of this observation, we run $N = 1000$ chains of length $L = 20$ using transition matrix (11). The reversible and nonreversible estimation results, together with the true eigenvalues, are reported in Table II.

It is seen that the reversible estimates do not only have the correct real-valued structure but can also have smaller uncertainties (here, especially for λ_3). This is expected to be a general result due to the smaller number of degrees of freedom in the reversible estimate.




FIG. 1. Eigenvalues of 3×3 example system. The eigenvalues obtained from the reversible estimate (green) are a closer approximation to the true eigenvalues (red) than the eigenvalues obtained from the non-reversible estimate (blue). The unique eigenvalue $\lambda = 1$ is faithfully reproduced by both estimates.

TABLE II. Results for non-reversible and reversible transition matrix estimation using count matrix (12).

	λ_1	$\text{Re}\{\lambda_2\}$	$\text{Im}\{\lambda_2\}$	$\text{Re}\{\lambda_3\}$	$\text{Im}\{\lambda_3\}$
Exact	1	0.42	0.0	0.18	0.0
Reversible	1	0.36 ± 0.31	0.0	0.18 ± 0.04	0.0
Non-reversible	1	0.32 ± 0.29	-0.04 ± 0.08	0.07 ± 0.21	0.04 ± 0.09

2. Example 2

Fig. 2 shows the distribution of eigenvalues from nonreversible and reversible Markov models from simulation data for the alanine-dipeptide molecule (see Sec. V B). The eigenvalues of the reversible estimate are purely real while the non-reversible estimate has eigenvalues with non-zero imaginary part.

C. Equilibrium kinetics analyses

Since detailed balance is a consequence of a system simulated at dynamical equilibrium, it is not surprising that detailed balance in the transition matrix P is a prerequisite to analyze the equilibrium kinetics given P . Since kinetics are related to slow processes, we will here only consider the m largest eigenvalues $\lambda_1, \dots, \lambda_m$ and assume that they are positive. Here are a few examples for equilibrium kinetics properties computed from reversible transition matrices:

1. The dominant relaxation rates of the molecular system are

$$\kappa_i = -\frac{1}{\tau} \ln \lambda_i, \quad (13)$$

where $i \geq 2$ ($i = 1$ has a relaxation rate of zero and corresponds to the equilibrium distribution). The inverse quantities are the relaxation time scales $t_i = \kappa_i^{-1}$. These rates or time scales are of special interest because they are often detectable in kinetic experiments such as fluorescence time-correlation spectroscopy, two-dimensional IR spectroscopy, or temperature jump experiments — see Ref. 11 for a discussion.

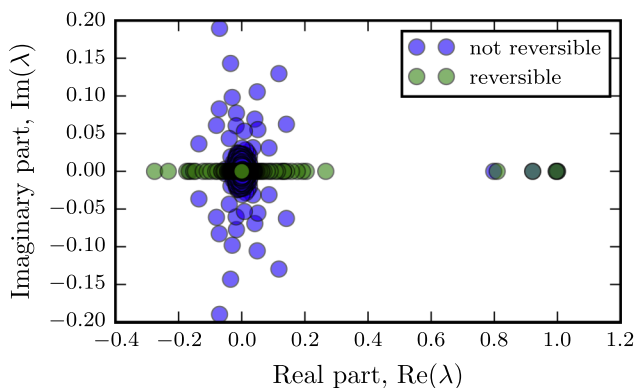


FIG. 2. Eigenvalues for alanine dipeptide. The cluster of dominant eigenvalues indicates that the slowest processes are faithfully reproduced by the non-reversible (blue) as well as by the reversible (green) estimate. Only the eigenvalues of the reversible estimate are purely real.

2. Decomposition (10) can be used to write kinetic experimental observables in an illuminating form.^{11,13,14} For example, the long-time scale part of the autocorrelation of a molecular observable $a \in \mathbb{R}^n$, e.g., containing the fluorescence values of every Markov state of a molecule, can be written as

$$\text{acf}(a; \tau) \approx \langle a, \pi \rangle^2 + \sum_{i=2}^m \langle a, r_i \rangle_{\pi}^2 e^{-\kappa_i \tau}. \quad (14)$$

3. The PCCA+ method for seeking m metastable sets of Markov states and its variants^{8,51} assumes m real-valued eigenvalues and eigenvectors. It is thus only reliably applicable to reversible transition matrices.
4. Discrete transition path theory^{20,52,53} computes the statistics of transition pathways from a set of states A to a set of states B given a transition matrix. Discrete TPT can be used with nonreversible and reversible transition matrices. However, in the reversible case, we get that the forward committor and the backward committor are complementary,

$$q_i^+ = 1 - q_i^-$$

and the net fluxes, when ordering states such that $q_i^+ \leq q_j^+$, are given by

$$f_{ij}^+ = (q_j^+ - q_i^+) \pi_i p_{ij}$$

which is analogous to an electric current $I = UG$, where $I = f_{ij}^+$, $U = (q_j^+ - q_i^+)$ is the potential difference and $\pi_i p_{ij}$ is the conductivity.⁵⁴

III. LIKELIHOOD, COUNTING, AND MAXIMUM LIKELIHOOD ESTIMATION

We restate the transition matrix likelihood and formulate the maximum likelihood estimation problem for Markov model transition matrices. We present new estimation algorithms for reversible Markov models with known or unknown equilibrium distribution π .

A. Likelihood

Suppose we have a discrete sequence $S = \{s_1, \dots, s_N\}$ with $s_i \in \{1, \dots, n\}$. If we assume that this sequence is the realization of a Markov chain with lag time $\tau = 1$, the probability that a transition matrix P has generated X is proportional to the product of individual transition probabilities along the trajectory

$$\mathbb{P}(S | P) \propto \prod_{t=1}^N p_{s_t s_{t+1}}. \quad (15)$$

We have neglected the proportionality constant because we will not need it in order to maximize or sample from (15). This is very handy because one component in this constant is the probability of generating the first state, p_{s_0} , which is often unknown, but is constant for a fixed data set S .

Suppose we have c_{ij} transitions from i to j . Then, we can group all p_{ij} terms together and get a factor $p_{ij}^{c_{ij}}$. Doing this for

all pairs results in the equivalent likelihood formulation

$$\mathbb{P}(C|P) \propto \prod_i \prod_j p_{ij}^{c_{ij}}. \quad (16)$$

We can see that the count matrix $C \in \mathbb{R}^{n \times n}$ is a sufficient statistics for the Markov model likelihood $\mathbb{P}(S | P)$ — it generates the same likelihood although we have discarded the information in which sequence the transitions have occurred.

If multiple trajectories are available, their count matrices are simply added up.

B. Counting

How should we count transitions for longer lag times $\tau > 1$, or if S is not Markovian at lag time τ ? Regarding the first case, if S is Markovian at lag time τ , a safe approach seems to subsample the trajectory at time steps of τ and then treat the subsampled trajectory as above.^{7,9} However, this approach is statistically inefficient: If S is also Markovian for shorter lag times than τ , then we are using less information than we could. Even if S only becomes Markovian at lag times of τ or longer, transitions such as $1 \rightarrow \tau + 1$ and $\tau/2 \rightarrow \tau + \tau/2$ are usually only partially correlated, such that discarding the second transition is also not fully exploiting the data. In practical MD simulations, the lag times required such that a Markov model is a good approximation need to be quite long (often in the range of nanoseconds), such that subsampling the data at τ will create severe problems with data and connectivity loss. Regarding the second case, if S is not Markovian at lag time τ , then treating every c_{ij} as an independent count is incorrect.

Both cases can, in principle, be treated with the following formalism: We always obtain the count matrix c_{ij} in a *sliding window* mode,⁷ i.e., we harvest all $N - \tau$ available transition counts from time pairs $(1 \rightarrow \tau)$, $(2 \rightarrow \tau + 1)$, \dots , $(N - \tau \rightarrow N)$. Unless S is Markovian at lag time 1, we will now harvest more transition counts than are statistically independent. We can formally correct for this by introducing a statistical inefficiency $I_{ij}(\tau)$ for every count at a given lag time, such that $c_{ij}^{\text{eff}}(\tau) = I_{ij}(\tau) c_{ij}(\tau)$ is the effective number of counts, resulting in the likelihood

$$\mathbb{P}(C|P) \propto \prod_{i,j} p_{ij}^{c_{ij}^{\text{eff}}}. \quad (17)$$

The determination of statistical inefficiencies for univariate signals is well established.⁵⁵ Determining $I_{ij}(\tau)$ for transition count matrices is an open problem. A first approach that allows for the first time to estimate consistent, although somewhat too small uncertainty intervals for practical MD data is discussed in Ref. 56. Note that the validity of the estimation algorithms described in the present paper is independent of the choice of the count matrix, such that future methods for estimating the effective count matrix can be adopted without changing the estimation algorithms.

C. Maximum likelihood estimation

We will now assume that the effective counts are given. For better readability, we will subsequently omit the superscript

eff and just use $C = (c_{ij})$ to indicate counts. Now we ask the question what is the most likely transition matrix for the observation C , i.e., we seek the MLE that maximizes (15) over the set of transition matrices.

1. Non-reversible estimation

It is well known that the non-reversible MLE for the transition probability from state i to state j is simply given by the ratio of observed counts from i to j divided by the total number of outgoing transitions from state i ,⁵⁷

$$\hat{P}_{ij}^{\text{nonrev}} = \frac{c_{ij}}{\sum_k c_{ik}}. \quad (18)$$

We use the hat in order to denote an estimator. The term non-reversible implies that reversibility has not been used as a constraint in the estimation of \hat{P}^{nonrev} . Of course \hat{P}^{nonrev} can be coincidentally reversible and will be reversible if the count matrix C is symmetric. For this reason, some early contributions in the field forced symmetry in C by counting S forward and backward. This practice is strongly discouraged as it will create a large bias unless the trajectories used are very long compared to the slowest time scales of the molecule.

2. Reversible estimation

Now we consider the problem of finding the reversible MLE \hat{P}^{rev} by enforcing detailed balance (4) with respect to an unknown equilibrium distribution (π_i) as constraint in the estimation procedure. Note that the count matrix used for this approach is not modified, i.e., it comes from a forward-only or nonreversible counting and is generally not symmetric. Constraints (4) can be more conveniently handled by defining the new set of variables

$$x_{ij} = \pi_i p_{ij}. \quad (19)$$

Note that

$$x_i = \sum_j x_{ij} = \pi_i. \quad (20)$$

We can thus recover the transition matrix from $X = (x_{ij})$ by

$$p_{ij} = \frac{x_{ij}}{x_i}. \quad (21)$$

Inserting (21) into (16) and adding constraints for detailed balance and stochasticity leads to the reversible maximum likelihood problem,

$$\begin{aligned} & \underset{X}{\text{maximize}} && \sum_{i,j} c_{ij} \log \frac{x_{ij}}{\sum_k x_{ik}} \\ & \text{subject to} && X_{ij} = X_{ji}, \\ & && \sum_k X_{ik} > 0, \\ & && X_{ij} \geq 0. \end{aligned} \quad (22)$$

Ignoring the inequality constraints, the optimality conditions are

$$\frac{c_{ij} + c_{ji}}{x_{ij}} - \frac{c_i}{x_i} - \frac{c_j}{x_j} = 0, \quad (23)$$

with $c_i = \sum_j c_{ij}$ and $x_i = \sum_j x_{ij}$. There is no closed form solution when including the detailed balance constraint so that (22) has to be solved numerically. One option is to directly solve (23) for x_{ij} and turn it into a fixed-point iteration, as first proposed in Ref. 37,

$$x_{ij}^{(k+1)} = \frac{c_{ij} + c_{ji}}{\frac{c_i}{x_i^{(k)}} + \frac{c_j}{x_j^{(k)}}}, \quad (24)$$

where k counts the iteration number in the algorithm. For a starting iterate $x_{ij}^{(0)}$ fulfilling the constraints in (22), for example, $x_{ij}^{(0)} = (c_{ij} + c_{ji}) / \sum_{i,j} (c_{ij} + c_{ji})$, the iterates will be symmetric and fulfill the inequality constraints for all $k > 0$.

If we sum over j on both sides of (24) and use (20), we can instead reduce the problem to iterative estimation of the equilibrium distribution,

$$\pi_i^{(k+1)} = \sum_{j=1}^n \frac{c_{ij} + c_{ji}}{\frac{c_i}{\pi_i^{(k)}} + \frac{c_j}{\pi_j^{(k)}}}. \quad (25)$$

The iteration is terminated when $\|\pi^{(k+1)} - \pi^{(k)}\| < \epsilon$. The final estimate $\hat{\pi}$ is then inserted into (23) to recover the reversible transition matrix estimate,

$$\hat{p}_{ij}^{\text{rev}} = \frac{(c_{ij} + c_{ji})\pi_j}{c_i\pi_j + c_j\pi_i}. \quad (26)$$

Note that both the optimum sought by Eqs. (25) and (26) exhibits $\hat{p}_{ij} = 0$ if $c_{ij} + c_{ji} = 0$. Thus, in both optimization algorithms, the sparsity structure of the matrix $\mathbf{C} + \mathbf{C}^T$ can be used in order to restrict all iterations to the elements that will result in a nonzero element $\hat{p}_{ij} > 0$.


Furthermore, note that (25) and (26) are special cases of the transition-based reweighing analysis (TRAM) method — see Ref. 16, Eqs. (29) and (30) — for the special case of a single thermodynamic state. An example for the progress of the self-consistent iteration using an alanine dipeptide simulation is shown in Fig. 3.

A different method of iterative solution presented in Ref. 7 updates x_{ij} with the exact solution to the quadratic problem arising from (23) while holding all other variables x_{kl} fixed. As shown in Fig. 3, this approach can exhibit faster convergence properties than fixed-point iteration (25).

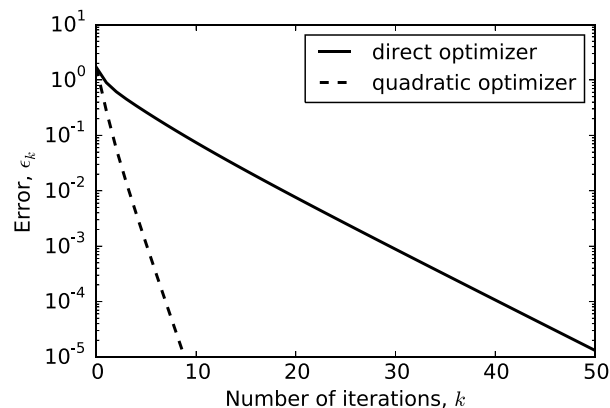
Uniqueness of the estimator: Optimization problem (22) can be equivalently transformed into a convex optimization problem by replacing the decision variables x_{ij} with $z_{ij} = \log(x_{ij})$ (see Ref. 17 for details), which implies the uniqueness of the maximum likelihood estimator.

3. Reversible estimation for given stationary vector

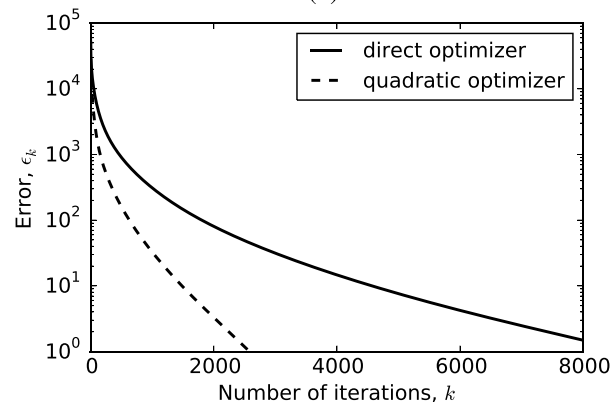
While unbiased MD simulations are useful to estimate state-to-state transition probabilities p_{ij} , enhanced sampling algorithms such as umbrella sampling and replica-exchange MD can be much more efficient in order to gain insight of the equilibrium distribution π . Reference 45 demonstrates how an uncertain estimate of π can be combined with unbiased “downhill” trajectories in order to estimate rare event kinetics. A key in such a procedure is a way to estimate a reversible Markov model that is most likely given transition counts



(a)



(b)



(c)

FIG. 3. Performance of algorithms for reversible maximum likelihood estimation. (a) Reversible transition matrix estimated using fixed-point iteration (25) from an $n = 228$ state count-matrix obtained from alanine-dipeptide simulation data. Convergence is shown for different total simulation lengths T . (b) Performance comparison of direct fixed-point iteration (25) and the quadratic optimizer described in Ref. 7 for reversible transition matrix estimation given the count matrix $\mathbf{C} = ((5, 2, 0), (1, 1, 1), (2, 5, 20))^T$. Shown is the difference of the current likelihood to the optimal likelihood. (c) Same as b, but using the 1734×1734 count matrix from Pin WW folding simulations used in Ref. 20.

observed from MD simulations, but at the same time has a fixed equilibrium distribution π . Here we derive a new, efficient estimation algorithm for this task.

Enforcing reversibility with respect to a given stationary vector results in the following constrained optimization problem:

$$\underset{P}{\text{maximize}} \quad \sum_{i,j} c_{ij} \log p_{ij} \quad (27)$$

$$\text{subject to } \sum_j p_{ij} = 1, \quad (28)$$

$$p_{ij} \geq 0, \quad (29)$$

$$\pi_i p_{ij} = \pi_j p_{ji}. \quad (30)$$

π can only be the unique stationary distribution of P if P is irreducible. To ensure irreducibility, we restrict the state space to the largest (weakly) connected set of the undirected graph that is defined by the adjacency matrix $C + C^T$. For a system with n states, Eqs. (27)–(30) are a convex minimization problem in $O(n^2)$ unknowns with $O(n^2)$ equality and inequality constraints. Solving this with a standard interior-point method requires the solution of a linear system with $O(n^2)$ unknowns to compute the search direction at each step. The resulting computational effort of $O(n^6)$ operations for solving the linear system quickly becomes unfeasible for increasing n . Therefore, we will propose a fixed-point iteration that is also feasible for large values of n .

To solve the maximization problem, we ignore inequality constraint (29) at first. Row-stochasticity constraint (28) is enforced by introducing Lagrange multipliers λ_i and adding penalty terms $\lambda_i (\sum_j p_{ij} - 1)$ for all $i = 1, \dots, n$ to the objective function. Detailed balance constraint (30) is included into the likelihood explicitly by the change of variables,

$$p'_{ij} = \begin{cases} p_{ij} & \text{if } i \leq j \\ \frac{\pi_j}{\pi_i} p_{ji} & \text{else} \end{cases}.$$

These substitutions result in the Lagrange function

$$F = \sum_i c_{ii} \log p'_{ii} + \sum_{i < j} (c_{ij} + c_{ji}) \log p'_{ij} - \sum_{i < j} p'_{ij} \left(\lambda_i + \lambda_j \frac{\pi_i}{\pi_j} \right) - \sum_i \lambda_i p'_{ii} + \sum_i \lambda_i + \text{const} \quad (31)$$

that we seek to maximize. By setting the gradient of F with respect to all p'_{ij} to zero and subsequently reversing the change of variables, we find the following expression for the maximum likelihood estimate:

$$\hat{p}_{ij} = \frac{(c_{ij} + c_{ji})\pi_j}{\lambda_i \pi_j + \lambda_j \pi_i}. \quad (32)$$

Note the similarity of this equation with maximum likelihood result (26) where π has been self-consistently computed from the counts. The row counts c_i are here replaced by the yet unknown Lagrange multipliers λ_i . In order to find the Lagrange multipliers, we sum Eq. (32) over j ,

$$\sum_j \frac{(c_{ij} + c_{ji})\pi_j}{\lambda_i \pi_j + \lambda_j \pi_i} = 1. \quad (33)$$

This does not give a closed-form expression for λ_i . However, based on this equation, we propose the following fixed-point iteration for the Lagrange multipliers:

$$\lambda_i^{(n+1)} = \sum_{j, c_{ij} + c_{ji} > 0} \frac{(c_{ij} + c_{ji})\lambda_j^{(n)}\pi_j}{\lambda_j^{(n)}\pi_i + \lambda_i^{(n)}\pi_j}. \quad (34)$$

Motivated by the analogy between Lagrange multipliers and row counts described above, we set the starting point to




FIG. 4. Convergence of the reversible maximum likelihood estimation with fixed stationary vector. The transition matrix is estimated from an $n = 228$ state count-matrix obtained from alanine-dipeptide simulation data. Convergence is shown for different total simulation lengths. The stationary distribution was obtained using simple counting estimate (83).

$$\lambda_i^{(0)} = \frac{1}{2} \sum_j (c_{ij} + c_{ji}). \quad (35)$$

Taking the limit $\lambda_i \rightarrow 0^+$ in (34) still leads to a consistent solution. Choosing strictly positive starting parameters according to (35) results in valid iterates from (34). In analogy to the reversible case, we iterate (34) and (35) until $\|\lambda^{(k+1)} - \lambda^{(k)}\| < \epsilon$. An example for the progress of the self-consistent iteration using alanine dipeptide simulation data is shown in Fig. 4. Note that the convergence is nearly three orders of magnitude faster compared to the estimation with unknown equilibrium distribution (Fig. 3).

Given converged Lagrange multipliers, we can exploit (32) to find the maximum likelihood transition matrix \hat{P} . For this algorithm, inequality constraints (29) are automatically fulfilled when $c_{ij} \geq 0$ for all i, j . Care must be taken in two situations: (i) $\lambda_i = \lambda_j = 0$ — one can show that the simultaneous limit $\lambda_i \rightarrow 0^+$ and $\lambda_j \rightarrow 0^+$ can only occur for $c_{ij} + c_{ji} = 0$, but then we know that $\hat{p}_{ij} = 0$. (ii) A diagonal element c_{ii} is zero. Depending on the values of π , the solution λ_i may take the value of zero such that Equation (32) for $i = j$ becomes $\hat{p}_{ii} = c_{ii}/\lambda_i = 0/0$ which is meaningless and is not the correct limit of p_{ii} as c_{ii} goes to zero. However, this can be fixed easily by using $\hat{p}_{ii} = 1 - \sum_{j \neq i} \hat{p}_{ij}$ for the diagonal elements of P . In summary, we use the following equation for computing \hat{P} from converged Lagrangian multipliers:

$$\hat{p}_{ij} = \begin{cases} \pi_j \frac{c_{ij} + c_{ji}}{\lambda_i \pi_j + \lambda_j \pi_i} & i \neq j, \lambda_i + \lambda_j \neq 0 \\ 0 & i \neq j, \lambda_i + \lambda_j = 0 \\ 1 - \sum_{j \neq i} \hat{p}_{ij} & i = j \end{cases}. \quad (36)$$

Uniqueness of the estimator: Since the above estimation algorithm is iterative, it is fair to ask whether the estimator \hat{P} it converges to is unique, or whether there might be multiple local maxima that we could get stuck in. In this case, it is easy to show that the estimator is unique: Let P^* be an optimal transition matrix. $p_{ij}^* = 0$ exactly if $c_{ij} + c_{ji} = 0$. Let $\Omega = \{p_{ij} | c_{ij} + c_{ji} > 0\}$. Then, the function $f(P) = \sum_{i,j} c_{ij} \log p_{ij}$ is strictly convex on Ω and the constraints restrict the solution on a

convex subset $\tilde{\Omega} \subset \Omega$. The minimization of a strictly convex function over a convex set has a unique solution.

IV. BAYESIAN ESTIMATION

We introduce new algorithms for sampling the full posterior probability distribution of Markov models, and, in particular, for estimating uncertainties of quantities of interest, such as relaxation time scales or mean first passage times. A key in these algorithms is the choice of a suitable prior which enforces the sampled matrices to have the same sparsity pattern as the transition count matrix, as this allows the credible intervals to lie around the true value even for large transition matrices. The relevance of the prior is first demonstrated for nonreversible Markov models, for which an efficient sampling algorithm is known. We then introduce new Gibbs sampling algorithms for reversible Markov models with and without constraints on the equilibrium distribution that vastly outperform previous algorithms for sampling reversible Markov models.

A. Bayes' theorem and Monte Carlo sampling

Bayes' formula relates the likelihood of an observed effective count matrix C given a probability model P to the posterior probability of the model given the observation

$$\underbrace{\mathbb{P}(P|C)}_{\text{posterior}} \propto \underbrace{\mathbb{P}(P)}_{\text{prior}} \underbrace{\mathbb{P}(C|P)}_{\text{likelihood}}. \quad (37)$$

The posterior accounts for the uncertainty coming from a finite observation. It incorporates *a priori* knowledge about the quantity of interest using the prior probability $\mathbb{P}(P)$. We will see that a suitable choice of the prior is essential for the success of a Bayesian description for high-dimensional systems.

In general, if we are interested in an observable that is a function of a transition matrix, $f(P)$, we would like to compute its posterior moments, such as the mean and the variance,

$$\langle f \rangle = \int dP \mathbb{P}(P|C) f(P), \quad (38)$$

$$\text{Var}(f) = \int dP \mathbb{P}(P|C) (f(P) - \langle f \rangle)^2. \quad (39)$$

While we usually use the maximum likelihood transition matrix \hat{P} to provide "best" estimates, $f(\hat{P})$, the above integrals are of interest because $\sigma(f) = \sqrt{\text{Var}(f)}$ gives us an estimate of the statistical uncertainty of f . Alternatively, we might be interested in the credible intervals which encompass the true value of f with some probability, such as 0.683 (1σ intervals) or 0.95 (2σ intervals). As integrals (38) and (39) are high-dimensional, we need to use Monte Carlo methods to approximate them.

In Monte Carlo methods, we generate a sample of transition matrices $\{P^{(k)}\}_{k=1}^N$ distributed according to the posterior and evaluate f at each element $P^{(k)}$ in the ensemble. We then approximate posterior expectation value (38) and posterior variance (39) by

$$m[f] = \frac{1}{N} \sum_{k=1}^N f(P^{(k)}), \quad (40)$$

$$s^2[f] = \frac{1}{N-1} \sum_{k=1}^N (f(P^{(k)}) - m[f])^2. \quad (41)$$

Obtaining good and reliable samples of the posterior $\mathbb{P}(P|C)$ is very difficult. Previous approaches have suffered from some or all of the following difficulties, that are addressed here:

- Choice of the prior:** Given n Markov states (typically 100 s to 1000 s), transition matrices have on the order of n^2 elements and are thus extremely high dimensional. Most priors used in the past allow to populate all these elements p_{ij} , including those for which no transition has been observed. Although the effect of the prior can be overcome by enough simulation data, for any practical amount of simulation data, such priors will lead to posterior distributions whose probability mass is far away from the true model. This problem has been first addressed in Ref. 20 by designing a prior that equates mean and MLE for nonreversible transition matrices, leading to credible intervals that nicely envelop the true value. Here, we design corresponding priors for reversible Markov models.
- Uncorrelated transition counts C :** As discussed in Sec. III B, the likelihood, and thus the posterior, depends on how transition counts are harvested from the discrete trajectories which are generally time-correlated and not exactly Markovian at any particular lag time τ . While the MLE is often not or little affected by the exact way of counting C , the uncertainties will be dramatically different if C is, e.g., counted in a sliding window mode (using transitions starting at all times $t = 0, 1, 2, \dots$), or by subsampling the trajectory (using transitions starting at all times $t = 0, \tau, 2\tau$). Whereas the first approach underestimates the uncertainties, the second approach often overestimates them and is often not practical for large lagtimes τ . Here, we suggest to use the effective number of uncorrelated transition counts, $C = C^{\text{eff}}$, and a first approach to compute them is described in Ref. 56.
- Efficiency of the sampler:** Finally, given a choice of prior and C , a sampling algorithm needs to explore the high-dimensional space of transition matrices in a reasonable time. This is especially problematic for reversible Markov models. The first Monte Carlo algorithm for sampling the reversible posterior, described in Ref. 39, suffers from poor mixing due to small acceptance probabilities of the individual steps. In Ref. 42, an improved sampler was proposed. Here, we propose sampling algorithms for reversible Markov models with and without fixed equilibrium distribution whose efficiencies go far beyond previous approaches.

B. Non-reversible sampling

Let us first illustrate the effect of prior choice on Bayesian estimation of nonreversible Markov models. A convenient functional form for the prior is the Dirichlet prior

$$\mathbb{P}(P) \propto \prod_i \prod_j p_{ij}^{b_{ij}}, \quad (42)$$

where $B = (b_{ij})$ is a matrix of prior-counts. For this choice, the posterior is given by

$$\mathbb{P}(P|C) \propto \prod_i \left(\prod_j p_{ij}^{z_{ij}} \right). \quad (43)$$

balance amounts to sampling symmetric matrices, $X = X^T$,

$$\mathbb{P}(X) \propto \prod_{i \geq j} x_{ij}^{b_{ij}}. \quad (47)$$

The posterior for reversible sampling is then given by

$$\mathbb{P}(X|C) \propto \prod_{i \geq j} x_{ij}^{b_{ij}} \prod_{i,j} \left(\frac{x_{ij}}{\sum_k x_{ik}} \right)^{c_{ij}}. \quad (48)$$

Below we will first consider how to sample from (48) using general prior counts b_{ij} . Then, we will consider the specific choice $b_{ij} = -1$ for all $i \geq j$ and show that this choice has similar properties as the sparse prior in the nonreversible case.

D. Sampling reversible transition matrices

There is no known method to generate independent samples from the posterior under the reversibility requirement. Instead we will use a Markov chain Monte Carlo (MCMC) method to generate samples from the posterior ensuring that each sampled transition matrix fulfills detailed balance condition (4). Our Markov chain will generate the ensemble $\{X^{(k)}\}_{k=1}^N$ via a set of updates advancing the chain from $X^{(k)} \rightarrow X^{(k+1)}$ starting from a valid initial state $X^{(0)}$. We can do a simple row-normalization of the X matrices to obtain the desired ensemble $\{P^{(k)}\}_{k=1}^N$. Expectation values and variances will again be estimated using (40) and (41).

Similarly as in Refs. 39 and 42, we will construct our Markov chain using a Gibbs sampling procedure, where we sample a single element of X in each step while leaving the other elements unchanged. We repeat this sampling procedure for every element of X , thus completing a Gibbs sweep. As detailed in the Appendix, we can use the following general Gibbs step to sample posterior (48):

1. Select an arbitrary element x_{kl} . Propose a new (unscaled) matrix $X \rightarrow X'$ by sampling this element from the proposal density $q(x'_{kl}|X)$,

$$x'_{ij} = \begin{cases} \sim q(x'_{kl}|X) & (i,j) = (k,l) \\ x_{ij} & \text{else} \end{cases}. \quad (49)$$

Here, $q(x'_{kl}|X)$ is an arbitrary, scale-invariant density. Scale-invariance means that $q(x'_{kl}|X) \propto q(cx'_{kl}|cX)$ for any positive constant c .

2. Accept \bar{X}' as a new step in our Markov chain with probability $\min\{1, p_{acc}\}$ where

$$p_{acc} = (1 - x_{kl} + x'_{kl})^{-\frac{n(n+1)}{2} - b_0} \frac{q(x_{kl}|X') \gamma(x'_{kl}|X)}{q(x'_{kl}|X) \gamma(x_{kl}|X')}, \quad (50)$$

where $b_0 = \sum_{k \geq l} b_{kl}$ and γ is the marginal density,

$$\gamma(x'_{kl}|X) \propto \begin{cases} \frac{(x'_{kk})^{c_{kk}+b_{kk}}}{(x_k - x_{kk} + x'_{kk})^{c_k}}, & k = l \\ \frac{(x'_{kl})^{c_{kl}+c_{lk}+b_{kl}}}{(x_k - x_{kl} + x'_{kl})^{c_k} (x_l - x_{kl} + x'_{kl})^{c_l}}, & k \neq l \end{cases}.$$

3. Renormalize the matrix $X' \rightarrow \bar{X}'$ such that it fulfills (46),

$$\bar{x}'_{ij} = \frac{x'_{ij}}{1 - x_{kl} + x'_{kl}}. \quad (51)$$

While this approach will work for any choice of prior counts, we will now use the sparse prior $b_{ij} = -1$ for all i, j with the hope to achieve similarly good results as in the nonreversible case. For this choice, $\gamma(x'_{kl}|X)$ is scale-invariant, i.e., $\gamma(x'_{kl}|X) = \gamma(cx'_{kl}|cX)$, and the Jacobian pre-factor in (50) is one. Thus, we have

$$p_{acc} = \frac{\gamma(x'_{kl}|X) q(x_{kl}|X')}{\gamma(x_{kl}|X') q(x'_{kl}|X)}, \quad (52)$$

and the ideal choice of the proposal density is $q \equiv \gamma$, which would guarantee that the acceptance probability is always 1. This proposal density degenerates to a point probability at zero if $c_{kl} + c_{lk} = 0$, which implies $b_{ij} = -1$ encodes *a priori* belief that any transition for which neither the forward direction nor the backward direction has ever been observed in the data has zero probability in the posterior ensemble. Thus, this prior enforces P to have the same sparsity structure as the count matrix, like the choice $b_{ij} = -1$ for nonreversible sampling. Note that the reversible MLE has the same sparsity structure as can be seen from update rule (24).

We will choose proposal densities $q(x'_{kl}|X)$ that are also scale-invariant. In this case, normalization step 3 above ($X' \rightarrow \bar{X}'$) can be omitted, i.e., if we accept X' , we can directly set it as our new sample $X^{(k)}$ and obtain $P^{(k)}$ by row normalization. We will now outline how to design the proposal density γ such that the acceptance probability is 1 or nearly 1. For $k = l$, sampling $x'_{kk} \sim \gamma(x'_{kk}|X)$ is equivalent to sampling the following transformed variable (see the Appendix):

$$s' = \frac{x'_{kk}}{x_k - x_{kk} + x'_{kk}} \sim \text{Beta}(c_{kk}, c_k - c_{kk}). \quad (53)$$

So we can simply define $q(x'_{kk}|X) \equiv \gamma(x'_{kk}|X)$ and generate x'_{kk} by

$$s' \sim \text{Beta}(c_{kk}, c_k - c_{kk}), \quad (54)$$

$$x'_{kk} = (x_k - x_{kk}) \frac{s'}{1 - s'}.$$

For $k \neq l$, there is no known way to draw independent samples, but $\gamma(x'_{kl}|X)$ can be well approximated by a gamma distribution by matching the maximum point and the second derivative at the maximum. A gamma distribution can be efficiently sampled and we use it as proposal density and accept the resulting x'_{kl} with probability $\min\{1, p_{acc}\}$. Specifically, our proposal step is

$$x'_{kl} \sim q(x'_{kl}|X) = \text{Gamma}(x'_{kl}|\alpha, \beta) \quad (55)$$

with the parameters

$$\alpha = -h\bar{v}, \quad (56)$$

$$\beta = -h\bar{v}^2 \quad (57)$$

using

$$\bar{v} = \frac{-b + \sqrt{b^2 - 4ac}}{2a}, \quad (58)$$

ALGORITHM 1. Reversible sampling algorithm.

Input: $C, X^{(j)}$
Output: $X^{(j+1)}$
for all indexes (k, l) **with** $k \geq l$ **and** $c_{kl} + c_{lk} > 0$ **do**
 if $k = l$ **then**
 Sample $x_{kk}^{(j+1)}$ from (54)
 end
 else
 Calculate α and β by (56), (57), and sample x'_{kl} from $\text{Gamma}(\alpha, \beta)$.
 $P_{acc} = \frac{\gamma(x'_{kl}|X) \text{Gamma}(x_{kl}|\alpha, \beta)}{\gamma(x_{kl}|X) \text{Gamma}(x'_{kl}|\alpha, \beta)}$
 Accept x'_{kl} as $x_{kl}^{(j+1)}$ with probability $\min\{1, P_{acc}\}$
 Sample x'_{kl} by $\log x'_{kl} \sim \mathcal{N}(\log x_{kl}, 1)$.
 $P_{acc} = \frac{\gamma(x'_{kl}|X) x'_{kl}}{\gamma(x_{kl}|X) x_{kl}}$
 Accept x'_{kl} as $x_{kl}^{(j+1)}$ with probability $\min\{1, P_{acc}\}$
 end
end

$$h = \frac{c_k}{(\bar{v} + x_k - x_{kl})^2} + \frac{c_l}{(\bar{v} + x_l - x_{kl})^2} - \frac{c_{kl} + c_{lk}}{\bar{v}^2}, \quad (59)$$

$$a = c_k + c_l - c_{kl} - c_{lk}, \quad (60)$$

$$b = (c_k - c_{kl} - c_{lk})(x_l - x_{kl}) + (c_l - c_{kl} - c_{lk})(x_k - x_{kl}), \quad (61)$$

$$c = -(c_{kl} + c_{lk})(x_k - x_{kl})(x_l - x_{kl}), \quad (62)$$

which matches the value and the first two derivatives of the true marginal density at the maximum (see the [Appendix](#) for derivation) and leads to acceptance probabilities close to one for most values of x_{kl} . However, if the current value of x_{kl} is in one of the heavy tails of the distribution $\gamma(x'_{kl}|X)$, the acceptance probability can be much less than 1 and the Markov chain can get stuck. In order to avoid this problem, we utilize a second step to generate x'_{kl} : After we sample x'_{kl} from proposal density (55), we sample x'_{kl} according to

$$\log x'_{kl} \sim \mathcal{N}(\log x'_{kl} | \log x_{kl}, 1), \quad (63)$$

where $\mathcal{N}(x|m, s)$ denotes the normal distribution of x with mean m and standard deviation s . The proposal density defined by the above update is

$$\tilde{q}(x'_{kl}|X) = \frac{1}{x_{kl}} \mathcal{N}(\log x'_{kl} - \log x_{kl} | 0, 1), \quad (64)$$

and the corresponding acceptance probability is

$$P_{acc} = \frac{\gamma(x'_{kl}|X) x'_{kl}}{\gamma(x_{kl}|X) x_{kl}}. \quad (65)$$

In summary, the proposed Algorithm 1 is a Metropolis within Gibbs MCMC algorithm with adapted proposal probabilities for each Gibbs sampling step. For efficiency reasons, transition matrix elements (i, j) for which no forward or backward transition counts have been observed can be neglected in the sampling algorithm, in order to account for the effect of the sparse prior.

E. A prior for reversible Markov models with fixed equilibrium distribution

As before we will be working with variables $x_{ij} = \pi_i p_{ij}$ related to transition matrix entries p_{ij} via (45). In contrast to the previous algorithm, π is not a function of P but fixed. Thus, single normalization condition (46) is replaced by a condition for each row,

$$\sum_j x_{ij} = \pi_i, \quad (66)$$

in order to ensure reversibility with respect to the given π .

All x_{kl} in the lower triangle ($k > l$) are used as independent variables. Given a valid X matrix, an update that respects the constraints is given by

$$x_{kl} \rightarrow x'_{kl}, \quad (67a)$$

$$x_{kk} \rightarrow x_{kk} + (x'_{kl} - x_{kl}), \quad (67b)$$

$$x_{lk} \rightarrow x'_{kl}, \quad (67c)$$

$$x_{ll} \rightarrow x_{ll} + (x'_{kl} - x_{lk}). \quad (67d)$$

(67b) and (67d) ensure that normalization condition (66) holds for the new sample and will thus keep π constant, while (67c) restores the symmetry and thus ensures reversibility of P . We will again use prior (47) defined on the set of X matrices and sample from posterior (48). The ideal proposal density of x'_{kl} is

$$\gamma(x'_{kl} | X) \propto (x'_{kl})^{c_{kl} + c_{lk} + b_{kl}} (x_{kk} + x_{kl} - x'_{kl})^{c_{kk} + b_{kk}} (x_{ll} + x_{kl} - x'_{kl})^{c_{ll} + b_{ll}} \quad (67e)$$

which is the conditional distribution density for given all off-diagonal elements of X (except x_{kl}), π , and the counts C .

We have seen that a correct choice of prior parameters was essential in order to successfully apply the posterior sampling for meta-stable systems. As in the reversible case, we will use $b_{kl} = -1$ for $k > l$ to enforce $x_{kl} = 0$ whenever $c_{kl} + c_{lk} = 0$.

However, the choice of prior counts for the diagonal elements b_{kk} is less straightforward. According to our experience, a good choice is to determine the value of b_{kk} based on the maximum likelihood estimate \hat{p}_{kk} of the k th diagonal element as

$$b_{kk} = \begin{cases} 0 & \hat{p}_{kk} > 0, c_{kk} = 0 \\ -1 & \text{else} \end{cases} \quad (68)$$

which ensures that the posterior expectation of p_{kk} is zero if and only if $\hat{p}_{kk} = 0$, and the conditional expectation of (67e),

$$\mathbb{E}(x'_{kl} | X) = \frac{c_{kl} + c_{lk}}{c_{kl} + c_{lk} + c_{ll}} (x_{ll} + x_{kl}), \quad (69)$$

matches the MLE of the one-dimensional likelihood function for x_{kl} given X if $\hat{p}_{kk} > 0$ and $c_{kk} = 0$. (Note that for the MLE of P , there is at most one k which satisfies $\hat{p}_{kk} > 0$ and $c_{kk} = 0$ — see proof in the [Appendix](#).)

However, in the case that $\hat{p}_{kk} = 0$, conditional (67e) would then degenerate so that $x'_{kk} = 0$ with probability one, and the k th row and column of X would remain fixed in the sampling process. This effect can break ergodicity in the sampled Markov chain and therefore prevent convergence of the algorithm. This problem is avoided by regularizing the prior choosing the prior parameter as $b_{kk} = -1 + \epsilon$ for $\hat{p}_{kk} = 0$ such that (67e)

does not degenerate, where $\epsilon > 0$ is a small number. In addition, we need to ensure that the Markov chain is started from an initial state $X^{(0)}$ with $x_{kk}^{(0)} > 0$. In summary, we select the prior of X for reversible sampling with fixed π as

$$b_{kk} = \begin{cases} 0 & \hat{p}_{kk} > 0, c_{kk} = 0 \\ -1 + \epsilon & \hat{p}_{kk} = 0, c_{kk} = 0 \\ -1 & c_{kk} > 0 \end{cases}. \quad (70)$$

This choice of prior will again ensure that $c_{kl} + c_{lk} = 0$ results in $p_{kl} = 0$ and $p_{lk} = 0$ for all $k < l$ and for all posterior samples, a property shared by the reversible MLE with fixed stationary vector. This ensures that the posterior mass is located around the maximum likelihood estimate \hat{P} and again prevents the occurrence of artificial kinetic pathways in the posterior ensemble.

F. Sampling reversible Markov models with fixed equilibrium distribution

We now investigate how to efficiently sample conditional (67e). Here, we assume without loss of generality that $x_{kk} < x_{ll}$ and transform $x'_{kl} \in (0, x_{kk} + x_{kl})$ into a new variable $v' \in (0, +\infty)$ via

$$v' = \frac{x'_{kl}}{x_{kk} + x_{kl} - x'_{kl}}. \quad (71)$$

The ideal proposal density of v' is then

$$\begin{aligned} \gamma_{v'}(v'|X) &\propto \left| \frac{\partial x'_{kl}}{\partial v'} \right| \gamma(x'_{kl}|X) \\ &= (v')^{c_{kl}+c_{lk}+b_{kl}} \left(\frac{s}{s-1} + v' \right)^{c_{ll}+b_{ll}} \\ &\quad \cdot (1+v')^{-(c_{kl}+c_{lk}+c_{kk}+c_{ll}+b_{kl}+b_{kk}+b_{ll}+2)} \\ &= (v')^{a_1} \left(\frac{s}{s-1} + v' \right)^{a_3} (1+v')^{-(a_1+a_2+a_3)}, \end{aligned} \quad (72)$$

with

$$\begin{aligned} s &= \frac{x_{ll} + x_{kl}}{x_{kk} + x_{kl}}, \\ a_1 &= c_{kl} + c_{lk} + b_{kl}, \\ a_2 &= c_{kk} + b_{kk}, \\ a_3 &= c_{ll} + b_{ll}. \end{aligned}$$

Like in the previous algorithm, we can approximate the conditional of v by a gamma distribution as

$$\gamma_{v'}(v'|X) \approx \text{Gamma}(v'|\alpha, \beta), \quad (73)$$

with

$$\alpha = -h\bar{v}, \quad (74)$$

$$\beta = -h\bar{v}^2, \quad (75)$$

and

$$\bar{v} = \frac{-b + \sqrt{b^2 - 4ac}}{2a}, \quad (76)$$

$$h = -\frac{a_1}{\bar{v}^2} - \frac{a_3}{\left(\frac{s}{s-1} + \bar{v}\right)^2} + \frac{a_2}{(1+\bar{v})^2}, \quad (77)$$

$$a = a_2 + 1, \quad (78)$$

$$b = a_2 - a_1 + \frac{a_2 + a_3 + 1}{s-1}, \quad (79)$$

ALGORITHM 2. Reversible sampling algorithm with fixed stationary vector.

Input: $C, \pi, X^{(j)}$

Output: $X^{(j+1)}$

for all indexes (k, l) **with** $k > l$ **and** $c_{kl} + c_{lk} > 0$ **do**

if $x_{kk} < x_{ll}$ **then**

$$v = \frac{x_{kl}}{x_{kk}}$$

$$s = \frac{x_{ll} + x_{kl}}{x_{kk} + x_{kl}}$$

$$a_1 = c_{kl} + c_{lk} + b_{kl}$$

$$a_2 = c_{kk} + b_{kk}$$

$$a_3 = c_{ll} + b_{ll}$$

end

else

$$v = \frac{x_{kl}}{x_{ll}}$$

$$s = \frac{x_{kk} + x_{kl}}{x_{ll} + x_{kl}}$$

$$a_1 = c_{kl} + c_{lk} + b_{kl}$$

$$a_2 = c_{ll} + b_{ll}$$

$$a_3 = c_{kk} + b_{kk}$$

end

Calculate α and β by (74) and (75)

Sample v' from $\text{Gamma}(\alpha, \beta)$.

Let $x'_{kl} = \min\{x_{kk} + x_{kl}, x_{ll} + x_{kl}\} \cdot \frac{v'}{1+v'}$.

$P_{acc} = \frac{\gamma_{v'}(v'|X)\text{Gamma}(v|\alpha, \beta)}{\gamma_v(v|X)\text{Gamma}(v'|\alpha, \beta)}$ using (67a)–(67d)

Accept x_{kl} as $x_{kl}^{(j+1)}$ with probability $\min\{1, P_{acc}\}$

Sample v' by $\log v' \sim \mathcal{N}(\log v, 1)$.

$P_{acc} = \frac{\gamma_{v'}(v'|X)v'}{\gamma_v(v|X)v}$

Let $x'_{kl} = \min\{x_{kk} + x_{kl}, x_{ll} + x_{kl}\} \cdot \frac{v'}{1+v'}$.

Accept x'_{kl} as $x_{kl}^{(j+1)}$ with probability $\min\{1, P_{acc}\}$

end

$$c = \frac{s(a_1 + 1)}{1 - s}. \quad (80)$$

See the Appendix for derivation. Then, v' can be sampled by a Metropolis sampling step with the proposal density $\text{Gamma}(v'|\alpha, \beta)$ and the acceptance ratio

$$\min\{1, P_{acc}\} = \min\left\{1, \frac{\gamma_{v'}(v'|X)\text{Gamma}(v|\alpha, \beta)}{\gamma_v(v|X)\text{Gamma}(v'|\alpha, \beta)}\right\},$$

where $v = x_{kl}/x_{kk}$ denotes the original value of v . In addition, in order to avoid the sampler from getting stuck at an extremely small or large value of v , we also utilize the same strategy as in Section IV D to generate v' by

$$\log v' \sim \mathcal{N}(\log v' | \log v, 1).$$

The proposed Algorithm 2 for sampling of reversible transition matrices with fixed stationary vector can again be characterized as a Metropolis within Gibbs MCMC algorithm with adapted proposal probabilities.

V. RESULTS

A. Validation

We first demonstrate the validity of the reversible sampling algorithm for the following 2×2 count-matrix:

$$C = \begin{pmatrix} 5 & 2 \\ 3 & 10 \end{pmatrix}. \quad (81)$$




FIG. 6. Sampled histogram frequency (a) and analytical probability density (b) of reversible posterior for 2×2 count matrix. Sampled frequencies are in agreement with the analytical probabilities.

In Fig. 6, we compare the sampled histograms using Algorithm 1 with analytical values for the non-reversible posterior with Dirichlet-prior-counts $b_{ij} = -1$. Any 2×2 transition matrix automatically fulfills detailed balance, and therefore the analytical and sampled densities are expected to be equal. The histogram for samples of the reversible algorithm is indeed in agreement with the analytical posterior.

In Fig. 7, the sampled histogram for count matrix (81) with fixed stationary distribution $\pi = (0.25, 0.75)^\top$ using Algorithm 2 is compared with the exact posterior distribution. Detailed balance relation (4) with fixed stationary vector enforces a linear dependency between the transition matrix element p_{12} and p_{21} . The resulting posterior is, therefore, restricted to the line $\pi_1 p_{12} = \pi_2 p_{21}$ such that the projection on p_{12} in Fig. 7 already contains the full information about the one-dimensional posterior. A comparison between histogram frequency and analytical density demonstrates the validity of the algorithm.

B. Applications

To demonstrate the usefulness of the proposed algorithms, we apply them to molecular dynamics simulation data. Here, two systems are chosen to illustrate our methods: (1) the alanine-dipeptide molecule and (2) the bovine pancreatic trypsin inhibitor (BPTI) molecule.

We start by discussing the alanine dipeptide results. The system was simulated on GPU-hardware using the OpenMM simulation package⁵⁹ using the *amber99sb-ildn* forcefield⁶⁰

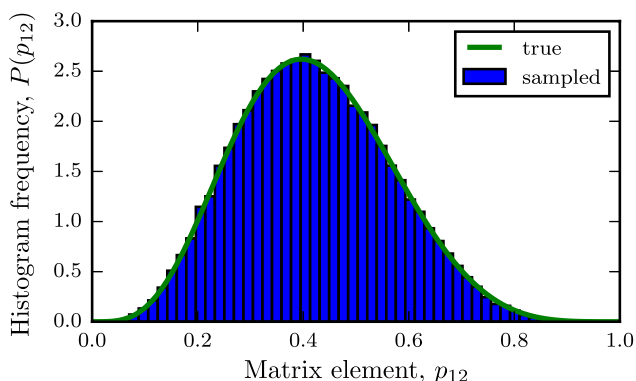


FIG. 7. Sampled histograms and analytical probability density for reversible posterior with fixed stationary vector for 2×2 count-matrix.

and the *tip3p* water model.⁶¹ The cubic box of length 2.7 nm contained a total of 652 solvent molecules. We used Langevin equations at $T = 300$ K with a time step of 2 fs. A total of 10 μ s of simulation data was generated. The ϕ and ψ dihedral angles were discretized using a 20×20 regular grid to obtain a matrix of transition counts $C = (c_{ij})$, here by sampling one count per lag time τ . Below we will show histograms for two important observables, *largest implied time scales* t_i and *expected hitting times*, $\tau(A \rightarrow B)$, for pairs A, B of meta-stable sets. We compute the posterior sample-mean and 90% credible intervals for 1 μ s of simulation data and show that the credible intervals nicely envelop a reference value obtained from the MLE transition matrix for the total simulation data, supporting the proposed prior as a “good” choice for reversible sampling in meta-stable systems.

1. Alanine dipeptide, reversible sampling

In Fig. 8, we show histograms of implied time scales computed from a reversible posterior sample. The mean values estimated from the posterior sample are in good agreement with the reference values. Tables IV and V compare the reference values with the sample mean μ and sample standard deviation σ for each observable.

In order to gain a first impression of the efficiency of the sampling and the quality of our estimates, we compute the integrated autocorrelation time t_{corr} for each quantity sampled (here implied time scales and hitting times). The error of the sample mean $m[f]$ compared to the true mean $\langle f \rangle$ can then be estimated as

$$\epsilon = \mathbb{E}[(m[f] - \langle f \rangle)^2] = \frac{\sigma^2[f]}{N_{\text{eff}}}, \quad (82)$$

where $N_{\text{eff}} = N/(1 + 2t_{corr})$ is the effective number of samples with N the total number of samples. See Ref. 62 for a thorough discussion. t_{corr} and ϵ are also reported in Table IV.

Figs. 8(d)–8(f) show histograms for expected hitting times for the three transitions $C_5 \rightarrow C_7^{\alpha x}$, $C_5 \rightarrow \alpha_L$, and $C_5 \rightarrow \alpha_R$ between meta-stable sets in the ϕ and ψ dihedral angle plane. Again, mean values are in good agreement with the corresponding reference values. Table V summarizes the computed results. The table columns again contain the reference value $\hat{\tau}$, the mean value μ , the standard deviation σ , the estimated correlation time t_{corr} , and the error of the mean value ϵ .

2. Alanine dipeptide, reversible sampling with fixed equilibrium distribution

Below we report results for sampling with fixed stationary distribution. The stationary distribution π_i was, for sake of simplicity, computed using the relative frequencies of state occurrences,

$$\pi_i = \frac{\sum_k c_{ik}}{\sum_{j,k} c_{jk}}. \quad (83)$$

It should be noted that a more useful and independent source of π are enhanced sampling simulations targeted at rapidly generating a good estimate of the equilibrium probabilities alone. See Ref. 45 for methods and applications that combine




FIG. 8. (a)-(c) Implied time scales, t_i . Histograms obtained from reversible posterior sampling. Dashed lines indicate the reference value, \hat{t}_i , dotted lines indicate the posterior sample mean, $\mu(t_i)$. MLE and posterior mean are in very good agreement for the proposed choice of prior. The 90% credible intervals are the shaded regions in gray. Expected hitting times, τ , (d)-(f). Histograms obtained from reversible posterior sampling. Dashed lines indicate the reference value, $\hat{\tau}$, dotted lines indicate the posterior sample mean, $\mu(\tau)$. In all cases, the reference value obtained from a long simulation is clearly compatible with the posterior sample (credible interval).

TABLE IV. Comparison of reference implied time scales (\hat{t}_i) with mean μ and standard deviation σ from the reversible posterior using $N = 10^5$ samples. ϵ is the estimated error of the mean μ and t_{corr} is the autocorrelation time of the sampled quantity.

	\hat{t}_i (ps)	μ (ps)	σ (ps)	ϵ (ps)	t_{corr}
t_2	1462	1556	303	19.00	197
t_3	71	73	1	0.01	10
t_4	36	43	5	0.06	7

TABLE V. Expected hitting times computed from reversible posterior using $N = 10^5$ samples. See Table IV for definition of other symbols.

	$\hat{\tau}$ (ns)	μ (ns)	σ (ns)	ϵ (ns)	t_{corr}
$\tau(C_5 \rightarrow C_7^{\alpha x})$	60.4	56.7	12.0	0.77	202
$\tau(C_5 \rightarrow \alpha_L)$	43.6	40.6	8.3	0.53	206
$\tau(C_5 \rightarrow \alpha_R)$	0.253	0.250	0.005	0.0004	218

MD simulations and enhanced sampling simulations in order to efficiently compute rare-event kinetics.

Results are shown in Table VI and Fig. 9. The sample mean is again in good agreement with the reference value. For

the computation of the reference value, we use the MLE transition matrix of the full simulation data reversible with respect to the input stationary distribution for the posterior sampling. The additional constraint imposed by fixing the stationary distribution is clearly reflected in smaller standard deviations for all shown observables compared to the reversible case.

Histograms for expected hitting times between metastable sets are shown in Figs. 9(d)–9(f). The sample mean is again in good agreement with the reference value. Again, we summarize our results, cf. Table VII.

TABLE VI. Comparison of reference implied time scales, (\hat{t}_i), with mean μ and standard deviation σ from the reversible posterior using a fixed equilibrium distribution and $N = 10^5$ samples. ϵ is the estimated error of the mean μ and t_{corr} is the autocorrelation time of the sampled quantity.

	\hat{t}_i (ps)	μ (ps)	σ (ps)	ϵ (ps)	t_{corr}
t_2	1594	1520	196	0.6	1
t_3	72	73	1	0.003	1
t_4	38	41	3	0.01	1




FIG. 9. (a)-(c) Implied time scales, t_i . Histograms obtained from reversible posterior sampling with fixed stationary vector. Dashed lines indicate the reference value, \hat{t}_i , dotted lines indicate the posterior sample mean, $\mu(t_i)$. The 90% credible intervals are the shaded regions in gray. (d)-(f) Expected hitting times τ . Histograms obtained from reversible posterior sampling with fixed stationary vector. Dashed lines indicate the reference value, $\hat{\tau}$, dotted lines indicate the posterior sample mean, $\mu(\tau)$. The reference value obtained from a long simulation is clearly compatible with the posterior sample (credible interval).

3. Bovine pancreatic trypsin inhibitor, reversible sampling

For BPTI, we used the 1 ms simulation generated on the Anton supercomputer.⁶³ Please refer to that paper for system setup and simulation details. We prepared data as follows: C_α atom positions were oriented to the mean structure and saved every 10 ns, resulting in about 100 000 configurations with 174 dimensions. Time-lagged independent component analysis (TICA)^{24,32} was applied to reduce this 174-dimensional space to the two dominant ICs as a spectral gap was found after the second nontrivial eigenvalue. k -means clustering with $k = 100$ was used to discretize this space.

Effective count matrices were obtained using the method described in Ref. 56 at a range of lag times up to 2 μ s.

TABLE VII. Expected hitting times computed from reversible sampling with fixed stationary vector using $N = 10^5$ samples. See Table VI for the definition of symbols.

	$\hat{\tau}$ (ns)	μ (ns)	σ (ns)	t_{corr}	ϵ (ns)
$\tau(C_5 \rightarrow C_7^{ax})$	56.0	54.5	5.0	1	0.02
$\tau(C_5 \rightarrow \alpha_L)$	41.5	39.5	5.1	1	0.02
$\tau(C_5 \rightarrow \alpha_R)$	0.249	0.251	0.003	1	$9.7 \cdot 10^{-6}$

Fig. 10 shows the implied relaxation time scales obtained from a maximum likelihood estimate with values comparable to the Hidden Markov model analysis in Ref. 10. The figure also shows uncertainties computed from a reversible transition matrix sampling as described above with $N = 1000$ samples. Only every 20th transition matrix sample was used to compute time scales in order to reduce the computational effort to 50 eigenvalue decompositions. It is seen that the MLE is nicely contained in the 2σ (95%) credible interval. The entire transition matrix sampling for Fig. 10 took about 12.5 s on a 1.7 GHz Intel Core i7. Given that 8 lag times were used for 1000 samples of 100×100 matrices that contained about 40% non-zeros, about 2.56×10^6 elements are sampled per second, and about 640 full transition matrix samples are generated per second. Below, a more systematic analysis of the computational efficiency is made.

C. Efficiency

We compute acceptance probabilities of the Metropolis-Hastings steps and compare the statistical efficiency of the proposed sampling algorithm with the algorithm proposed in Ref. 39 that uses uniform proposal densities. Efficiency is measured in terms of achieved autocorrelation times for sampling of transition matrices with different sizes. As a




FIG. 10. Implied time scales for a Markov model of bovine pancreatic trypsin inhibitor (BPTI). The error bars are 95% confidence intervals estimated using the reversible transition matrix sampling algorithm described here using transition counts as described in Ref. 56.

representative observable, we choose the largest relaxation time scale t_2 for the alanine dipeptide molecule and compute autocorrelation functions and autocorrelation times from a large sample of size $N = 10^6$. Two differently fine discretiza-

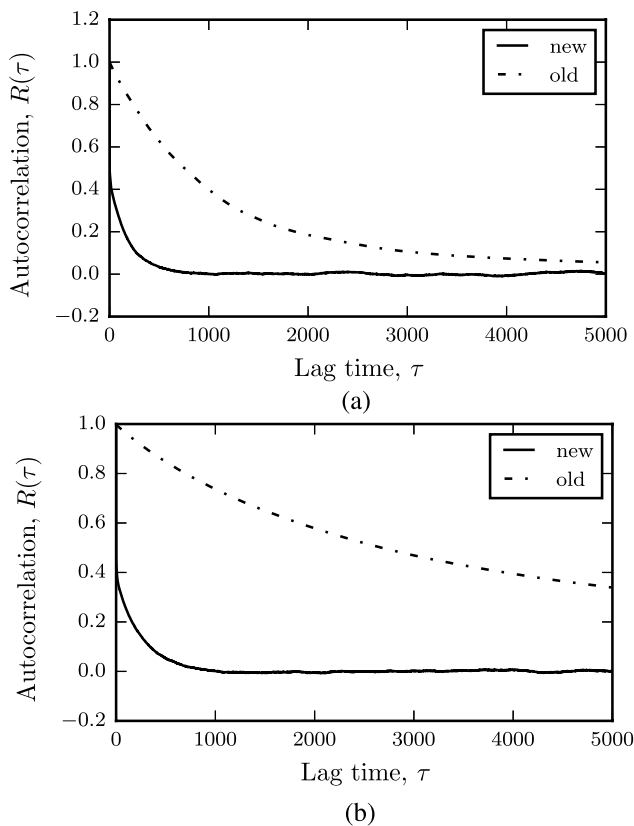


FIG. 11. Autocorrelation function for reversible sampling. Sampling of transition matrices with $n = 233$ states (a) and (b) with $n = 1108$ states. Increasing the dimension of the state space has only a small effect on the new sampling algorithm.

TABLE VIII. Acceptance probability and autocorrelation times for old vs. new reversible sampling algorithm. n : number of states; p_{offdiag} , p_{diag} : acceptance probabilities for off-diagonal and diagonal elements, respectively. t_{corr} : autocorrelation time for the sampling of the slowest relaxation time scale, in number of transition matrix sampling steps.

Algorithm	n	p_{offdiag}	p_{diag}	t_{corr}
Old	233	0.216	0.011	1088.1
	1108	0.271	0.005	3241.9
New	233	0.994	1.0	194.7
	1108	0.995	1.0	242.6

tions were used, resulting in $n \times n$ -shaped transition matrices with $n = 258$ and $n = 1108$.

Fig. 11 shows autocorrelation functions for the second largest relaxation time scale, t_2 , for a reversible posterior ensemble. The autocorrelation function for the reversible sampling algorithm with posterior adapted proposals shows a much faster decay than the autocorrelation function for the algorithm in Ref. 39. Table VIII compares acceptance probabilities and autocorrelation times. The present proposal steps lead to very high acceptance probabilities, $p > 0.99$, for the sampling off-diagonal entries. The main advance, however, comes from the fact that the step for sampling diagonal transition matrix elements in Ref. 39 has suffered from a very poor acceptance probability. As that step was the only step that modified the equilibrium distribution, the sampler




FIG. 12. Autocorrelation function for reversible sampling with fixed stationary distribution. Sampling of transition matrices with $n = 233$ states (a) and (b) with $n = 1108$ states. Increasing the dimension of the state space has only a small effect on the new sampling algorithm.

TABLE IX. Acceptance probability and autocorrelation times for old vs. new reversible sampling algorithm with fixed stationary distribution. Symbols as in Table VIII.

Algorithm	n	P_{offdiag}	t_{corr}
Old	233	0.175	72.096
	1108	0.230	>1000 ^a
New	233	0.752	2.893
	1108	0.706	3.157

^a Autocorrelation function not converged.

in Ref. 39 has very poor mixing properties. In contrast, our current algorithm generates independent samples for the diagonal elements, resulting in an acceptance probability of $p = 1.0$.

The autocorrelation times for the new sampler are more than a factor 5 shorter for the small (233 state) matrix and more than a factor 13 shorter for the large (1108 state) matrix, indicating a much higher efficiency of the new approach. The autocorrelation time of the new algorithm increases only mildly for matrices of increased dimension, indicating that the present algorithm will be useful for very large Markov models.

Fig. 12 shows autocorrelation functions for reversible sampling with fixed stationary vector. The posterior adapted proposals in reversible sampling algorithm with fixed stationary distribution again result in a much faster decay of the autocorrelation function than the uniform proposals of the algorithm in Ref. 39. Table IX compares acceptance probabilities and autocorrelation times for the two algorithms. For sampling with fixed stationary vector, there is no sampling step for the diagonal elements. Although the average acceptance probabilities are only a factor of 3-4 better for our new algorithm, the autocorrelation times are decreased by a factor 35 for the small system (233 states) and over a factor 300 for the large system (1108 states). Again, there is only a mild increase in autocorrelation time when the dimension of the sampled space is increased.

VI. CONCLUSION

In this work, we have described and significantly extended the state of the art in reversible Markov model estimation. Reversible Markov models are expected to naturally arise from molecular dynamics implementations that fulfill microscopic reversibility. Reversibility is an essential property in order to analyze the equilibrium kinetics of a molecule. However, in order to have reversibility in a Markov model, it needs to be enforced in the estimation procedure. When done correctly, reversible estimation does not bias the model but rather reduces statistical errors as a result of a smaller number of degrees of freedom.

We have presented minor improvements to an existing self-consistent estimation algorithm for reversible Markov models. Then, we have presented a new and efficient algorithm to estimate reversible maximum likelihood Markov models given a fixed equilibrium distribution.

The main part of the presented work focuses on the long-standing problem of Bayesian estimation of the posterior

ensemble of reversible transition matrices. Although several algorithms to sample reversible Markov models have been presented in the past, they have been hampered by three fundamental problems, two of which are addressed here: (i) Which prior should be chosen such that the posterior is located around the true value rather than completely elsewhere? (ii) How should transition counts be obtained when time series are correlated and not really Markovian at any given lag time τ ? (iii) How can the sampling algorithm be made efficient such that also large transition matrices can be sampled in reasonable time?

To address problem (i), we develop priors that ensure that reference values and sample mean are similar. The key property of these priors is that they make the *a priori* assumption that transitions between states that have not been sampled in the trajectory in either direction have zero probability. This is a sparse prior, i.e., an improper prior enforcing that the sampled transition matrix has the same sparsity structure as the maximum likelihood estimate and as induced by the observation. In contrast to most other priors that have been previously suggested, these sparse priors achieve the desired property of creating error bars that nicely envelop the reference estimates.

For problem (ii), we have described the principles of how it can be addressed. We suggest that effective count matrices are obtained using the concept of statistical inefficiencies. A separate preprint⁵⁶ suggests an initial solution towards this aim that is successfully applied on simulation data of the bovine pancreatic trypsin inhibitor in the present paper. The solution of problem (ii) is still in its infancy and needs further investigation.

For problem (iii), we present highly efficient Gibbs sampling algorithms for reversible transition matrices and reversible transition matrices with fixed equilibrium distribution. Both methods are demonstrated to have acceptance probabilities close to 1 in their individual update steps. Autocorrelation times from samples of the slowest relaxation time scale are one or two orders of magnitude shorter than with a previous Gibbs sampling algorithm, indicating a high statistical efficiency of our sampler.

Implementations of all algorithms described here are available in PyEMMA⁶⁴ as of version 2.0 or later — www.pyemma.org.

ACKNOWLEDGMENTS

This work was funded by the Deutsche Forschungsgemeinschaft (DFG) Project Nos. SFB 1114/C03 (F.P.), SFB 1114/A04 (H.W.), Grant No. 825/3-1 (B.T.S.), and the European Research Council (ERC) starting grant “pcCell” (F.N., B.T.S., F.P.). The authors would like to thank the anonymous referees for their valuable comments.

APPENDIX: DETAILS FOR TRANSITION MATRIX SAMPLING

1. Reversible transition matrix sampling: Derivation of marginal densities

We first pick a single element (k, l) of X (diagonal or off-diagonal) and sample it from a proposal density $x'_{kl} \sim q(x'_{kl}|X)$

that is scale-invariant with $q(x'_{kl}|X) \propto q(cx'_{kl}|cX)$ for all $c > 0$,

$$x'_{ij} = \begin{cases} \sim q(x'_{kl}|X), & (i,j) = (k,l) \\ x_{ij}, & \text{else} \end{cases} \quad (\text{A1})$$

and then renormalize the matrix such that it retains an element sum of 1,

$$\bar{x}'_{ij} = \frac{x_{ij}}{1 - x_{kl} + x'_{kl}}. \quad (\text{A2})$$

Since $q(x'_{kl}|X)$ is a probability density function, we can obtain from its scale-invariance that

$$\int q(cx'_{kl}|cX)dx'_{kl} = \frac{1}{c} \int q(x'_{kl}|X)dx'_{kl} = \frac{1}{c} \quad (\text{A3})$$

and

$$q(cx'_{kl}|cX) = \frac{1}{c} q(x'_{kl}|X). \quad (\text{A4})$$

According to Theorem 13.1 in Ref. 65, the posterior distribution $\mathbb{P}(X|C)$ is the invariant distribution of the proposed update step if we accept \bar{X}' as the new sample with probability $\min\{1, p_{acc}\}$ and

$$p_{acc} = \frac{\mathbb{P}(\bar{X}')}{\mathbb{P}(X)} \cdot \frac{\mathbb{P}(C|\bar{X}')}{\mathbb{P}(C|X)} \cdot \frac{q_x(x_{kl}|\bar{X}')}{q_x(\bar{x}'_{kl}|X)} \cdot \prod_{(i,j) \neq (k,l), (i',j')} \frac{\partial \bar{x}'_{ij}}{\partial x_{ij}}, \quad (\text{A5})$$

where $q_x(\bar{x}'_{kl}|X)$ denotes the proposal density of \bar{x}'_{ij} given X . Note X only contains $n(n+1)/2 - 1$ free variables. So we select $\{x_{ij}|i \geq j, (i,j) \neq (i',j')\}$ as the free variable set of X , where $x_{i'j'}$ is an arbitrary element of X with $i' \geq j'$ and $(i',j') \neq (k,l)$. Let us consider each term on the right hand side of (A5).

From the definition of \bar{X}' , we have

$$\begin{aligned} \frac{\mathbb{P}(\bar{X}')}{\mathbb{P}(X)} &= \left(\frac{1}{1 - x_{kl} + x'_{kl}} \right)^{b_0} \left(\frac{y}{x_{kl}} \right)^{b_{kl}} \\ &= \left(\frac{1 - \bar{x}'_{kl}}{1 - x_{kl}} \right)^{b_0} \left(\frac{x'_{kl}}{x_{kl}} \right)^{b_{kl}} \end{aligned} \quad (\text{A6})$$

and

$$\frac{\mathbb{P}(C|\bar{X}')}{\mathbb{P}(C|X)} = \begin{cases} \frac{(x'_{kk})^{c_{kk}} (x_k - x_{kk} + x'_{kk})^{-c_k}}{x_{kk}^{c_{kk}} (x_k - x_{kk} + x_{kk})^{-c_k}}, & k = l \\ \frac{(x'_{kl})^{c_{kl} + c_{lk}} (x_k - x_{kl} + x'_{kl})^{-c_k}}{x_{kl}^{c_{kl} + c_{lk}} (x_k - x_{kl} + x_{kl})^{-c_k}} \\ \times \frac{(x_l - x_{kl} + x'_{kl})^{-c_l}}{(x_l - x_{kl} + x_{kl})^{-c_l}}, & k \neq l \end{cases}. \quad (\text{A7})$$

The proposal density of \bar{x}'_{kl} given X can be expressed as

$$\begin{aligned} q_x(\bar{x}'_{kl}|X) &= \left(\frac{\partial \bar{x}'_{kl}}{\partial x'_{kl}} \right)^{-1} q(x'_{kl}|X) \\ &= \frac{1 - x_{kl}}{(1 - \bar{x}'_{kl})^2} q(x'_{kl}|X) \\ &= \frac{1 - x_{kl}}{(1 - \bar{x}'_{kl})^2} q \left(\frac{1 - x_{kl}}{1 - \bar{x}'_{kl}} \cdot \bar{x}'_{kl} | X \right). \end{aligned} \quad (\text{A8})$$

Therefore,

$$\begin{aligned} q_x(x_{kl}|\bar{X}') &= \frac{1 - \bar{x}'_{kl}}{(1 - x_{kl})^2} q \left(\frac{1 - \bar{x}'_{kl}}{1 - x_{kl}} \cdot x_{kl} | \bar{X}' \right) \\ &= \frac{1}{1 - x_{kl}} q(x_{kl}|X') \end{aligned} \quad (\text{A9})$$

and

$$\frac{q_x(x_{kl}|\bar{X}')}{q_x(\bar{x}'_{kl}|X)} = \left(\frac{1 - x_{kl}}{1 - \bar{x}'_{kl}} \right)^2 \frac{q(x'_{kl}|X)}{q(x_{kl}|X')}. \quad (\text{A10})$$

The partial derivative of \bar{x}'_{ij} with respect to x_{ij} for $(i,j) \neq (k,l)$ can be calculated according to (A2) as

$$\frac{\partial \bar{x}'_{ij}}{\partial x_{ij}} = \frac{1}{1 - x_{kl} + x'_{kl}} = \frac{1 - \bar{x}'_{kl}}{1 - x_{kl}}. \quad (\text{A11})$$

Combining all the above results leads to

$$p_{acc} = (1 - x_{kl} + x'_{kl})^{-\frac{n(n+1)}{2} - b_0} \frac{q(x_{kl}|X')}{q(x'_{kl}|X)} \frac{\gamma(x'_{kl}|X)}{\gamma(x_{kl}|X)}, \quad (\text{A12})$$

with

$$\gamma(x'_{kl}|X) \propto \begin{cases} \frac{(x'_{kk})^{c_{kk} + b_{kk}}}{(x_k - x_{kk} + x'_{kk})^{c_k}}, & k = l \\ \frac{(x'_{kl})^{c_{kl} + c_{lk} + b_{kl}}}{(x_k - x_{kl} + x'_{kl})^{c_k} (x_l - x_{kl} + x'_{kl})^{c_l}}, & k \neq l \end{cases}. \quad (\text{A13})$$

2. Reversible transition matrix sampling: Efficient proposal densities

a. Diagonals

Let us define variable $s' = \frac{x'_{kk}}{x_k - x_{kk} + x'_{kk}}$. If x'_{kk} is sampled from the proposal density $\gamma(x'_{kk}|X)$, the corresponding proposal density of s' can be expressed as

$$s' \sim \left| \frac{\partial x'_{kk}}{\partial s'} \right| \gamma(x'_{kk}|X). \quad (\text{A14})$$

Note that

$$x'_{kk} = (x_k - x_{kk}) \frac{s'}{1 - s'}. \quad (\text{A15})$$

Therefore,

$$\frac{\partial x'_{kk}}{\partial s'} = (x_k - x_{kk}) (1 - s')^{-2} \quad (\text{A16})$$

and

$$\begin{aligned} s' &\sim (x_k - x_{kk}) (1 - s')^{-2} \cdot (x'_{kk})^{c_{kk} - 1} \cdot (x_k - x_{kk} + x'_{kk})^{-c_k} \\ &\propto (1 - s')^{-2} \cdot \left(\frac{s'}{1 - s'} \right)^{c_{kk} - 1} \cdot (1 - s')^{c_k} \\ &= (s')^{c_{kk} - 1} (1 - s')^{c_k - c_{kk} - 1}. \end{aligned} \quad (\text{A17})$$

The above equation implies that s' follows the beta distribution with parameters c_{kk} and $c_k - c_{kk}$. So we can sample $s' \sim \text{Beta}(c_{kk}, c_k - c_{kk})$ and get x'_{kk} by the back-transform.

b. Off-diagonals

We consider how to approximate $\gamma(x'_{kl}|X)$ with $k \neq l$ and $b_{kl} = -1$ by a gamma distribution density function. Define

$$\begin{aligned} & (x'_{kl})^{c_{kl}+c_{lk}-1} (x_k - x_{kl} + x'_{kl})^{-c_k} (x_l - x_{kl} + x'_{kl})^{-c_l} \\ & = (x'_{kl})^{-1} \exp f(x'_{kl}). \end{aligned} \quad (\text{A18})$$

The function $f(x'_{kl})$ is then given by

$$\begin{aligned} f(x'_{kl}) = & (c_{kl} + c_{lk}) \log x'_{kl} - c_k \log(x_k - x_{kl} + x'_{kl}) \\ & - c_l \log(x_l - x_{kl} + x'_{kl}). \end{aligned} \quad (\text{A19})$$

We approximate f using a three parameter family of functions

$$\hat{f}(x'_{kl}|\alpha, \beta, f_0) = \alpha \log x'_{kl} - \beta x'_{kl} + f_0 \quad (\text{A20})$$

so that the corresponding approximate $\gamma(x'_{kl}|X)$ is

$$\hat{\gamma}(x'_{kl}|X) \propto (x'_{kl})^{-1} \exp \hat{f}(x'_{kl}|\alpha, \beta, f_0). \quad (\text{A21})$$

(A21) is a gamma distribution with parameters α, β which can be efficiently sampled.⁶⁶

The three parameters α, β, f_0 are obtained matching f and \hat{f} up to second derivatives at the maximum point

$$\bar{v} = \arg \max f(x'_{kl}). \quad (\text{A22})$$

This leads to the following linear system:

$$\alpha \log \bar{v} + \bar{v} \beta + f_0 = f(\bar{v}), \quad (\text{A23a})$$

$$\frac{\alpha}{\bar{v}} - \beta = f'(\bar{v}) = 0, \quad (\text{A23b})$$

$$-\frac{\alpha}{\bar{v}^2} = f''(\bar{v}) = h, \quad (\text{A23c})$$

with solution

$$\alpha = -h\bar{v}^2, \quad (\text{A23d})$$

$$\beta = -h\bar{v}, \quad (\text{A23e})$$

$$f_0 = f(\bar{v}) + h\bar{v}^2(\log \bar{v} - 1), \quad (\text{A23f})$$

and

$$h = f''(\bar{v}) = \frac{c_k}{(x_k - x_{kl} + \bar{v})^2} + \frac{c_l}{(x_l - x_{kl} + \bar{v})^2} - \frac{c_{kl} + c_{lk}}{\bar{v}^2}.$$

The maximum point can be computed as the root of a quadratic equation in the usual way,

$$\bar{v} = \frac{-b + \sqrt{b^2 - 4ac}}{2a}, \quad (\text{A23g})$$

with parameters a, b, c given by

$$a = c_k + c_l - c_{kl} - c_{lk}, \quad (\text{A23h})$$

$$\begin{aligned} b = & (c_k - c_{kl} - c_{lk})(x_l - x_{kl}) \\ & + (c_l - c_{kl} - c_{lk})(x_k - x_{kl}), \end{aligned} \quad (\text{A23i})$$

$$c = -(c_{kl} + c_{lk})(x_k - x_{kl})(x_l - x_{kl}). \quad (\text{A23j})$$

The second solution corresponding to (A23g) with negative sign in front of the square root can be safely excluded since \bar{v} is required to be non-negative.

3. Reversible sampling with fixed stationary distribution: Efficient proposal densities

We first prove that there exists a maximum likelihood estimate $\hat{\mathbf{P}} = [\hat{p}_{ij}]$ of the transition matrix which satisfies $\{k|\hat{p}_{kk} > 0, c_{kk} = 0\} \leq 1$. Suppose that $\mathbf{X}' = [x'_{ij}]$ is a maximum likelihood estimate of \mathbf{X} and there are two different

indices k, l which satisfy that $x'_{ll} \geq x'_{kk} > 0$ and $c_{kk} = c_{ll} = 0$. We can then construct a new matrix $\mathbf{X}'' = [x''_{ij}]$ with

$$x''_{kl} = x'_{kl} + x'_{kk},$$

$$x''_{kk} = 0,$$

$$x''_{kl} = x'_{kl} + x'_{kk},$$

$$x''_{ll} = x'_{ll} + x'_{kl} - x'_{kk}.$$

It can be verified that the likelihood of \mathbf{X}'' is larger than \mathbf{X}' , which implies that \mathbf{X}'' is also a maximum likelihood estimate of \mathbf{X} . Repeat the above procedure, we can finally get an maximum likelihood estimate $\hat{\mathbf{X}}$ of \mathbf{X} which has at most one k with $\hat{x}_{kk} > 0$ and $c_{kk} = 0$, and the corresponding $\hat{\mathbf{P}}$ satisfies that $\{k|\hat{p}_{kk} > 0, c_{kk} = 0\} \leq 1$.

We will now investigate how to approximate the density

$$\gamma_v(v'|X) \propto (v')^{\alpha_1} \left(\frac{s}{s-1} + v'\right)^{\alpha_3} \left(1 + v'\right)^{-(\alpha_1 + \alpha_2 + \alpha_3 + 2)}, \quad (\text{A24})$$

with $s > 1$ by a gamma distribution density.

As in the reversible case, we will use the representation

$$\gamma_v(v'|X) = (v')^{-1} \exp f(v') \quad (\text{A25})$$

and approximate $f(v')$ using the three parameter family, $\hat{f}(v'|\alpha, \beta, f_0)$, given in (A20). The resulting approximate density, $\hat{\gamma}_v(v'|X)$, has the same nice properties as the one from (A21).

Parameters α, β , and f_0 are given by (A23d), (A23e), and (A23f). The maximum point \bar{v} is given by (A23g). The parameters a, b, c are

$$a = a_2 + 1, \quad (\text{A26a})$$

$$b = a_2 - a_1 + \frac{a_2 + a_3 + 1}{s-1}, \quad (\text{A26b})$$

$$c = \frac{s(a_1 + 1)}{1-s}. \quad (\text{A26c})$$

¹C. Schütte, A. Fischer, W. Huisinga, and P. Deuffhard, *J. Comput. Phys.* **151**, 146 (1999).

²W. C. Swope, J. W. Pitner, and F. Suits, *J. Phys. Chem. B* **108**, 6571 (2004).

³N. Singhal, C. D. Snow, and V. S. Pande, *J. Chem. Phys.* **121**, 415 (2004).

⁴V. Schultheis, T. Hirschberger, H. Carstens, and P. Tavan, *J. Chem. Theory Comput.* **1**, 515 (2005).

⁵F. Noe, I. Horenko, C. Schütte, and J. C. Smith, *J. Chem. Phys.* **126**, 155102 (2007).

⁶A. Pan and B. Roux, *J. Chem. Phys.* **129**, 064107 (2008).

⁷J. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. Chodera, C. Schütte, and F. Noé, *J. Chem. Phys.* **134**, 174105 (2011).

⁸P. Deuffhard and M. Weber, in *Linear Algebra and Its Applications*, edited by M. Dellnitz, S. Kirkland, M. Neumann, and C. Schütte (Elsevier, New York, 2005), Vol. 398C, pp. 161–184.

⁹N. Singhal and V. S. Pande, *J. Chem. Phys.* **123**, 204909 (2005).

¹⁰F. Noé, H. Wu, J.-H. Prinz, and N. Plattner, *J. Chem. Phys.* **139**, 184114 (2013).

¹¹F. Noé, S. Doose, I. Daidone, M. Löllmann, J. D. Chodera, M. Sauer, and J. C. Smith, *Proc. Natl. Acad. Sci. U. S. A.* **108**, 4822 (2011).

¹²W. Zhuang, R. Z. Cui, D.-A. Silva, and X. Huang, *J. Phys. Chem. B* **115**, 5415 (2011).

¹³B. G. Keller, J.-H. Prinz, and F. Noé, *Chem. Phys.* **396**, 92 (2012).

¹⁴B. Lindner, Z. Yi, J.-H. Prinz, J. C. Smith, and F. Noé, *J. Chem. Phys.* **139**, 175101 (2013).

¹⁵S. Sriraman, I. G. Kevrekidis, and G. Hummer, *J. Phys. Chem. B* **109**, 6479 (2005).

¹⁶H. Wu, A. S. J. S. Mey, E. Rosta, and F. Noé, *J. Chem. Phys.* **141**, 214106 (2014).

¹⁷H. Wu and F. Noé, *SIAM Multiscale Model. Simul.* **12**, 25 (2014).

- ¹⁸E. Rosta and G. Hummer, *J. Chem. Theory Comput.* **11**, 276 (2015).
- ¹⁹S. P. Elmer, S. Park, and V. S. Pande, *J. Chem. Phys.* **123**, 114903 (2005).
- ²⁰F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl, *Proc. Natl. Acad. Sci. U. S. A.* **106**, 19011 (2009).
- ²¹I. Buch, T. Giorgino, and G. De Fabritiis, *Proc. Natl. Acad. Sci. U. S. A.* **108**, 10184 (2011).
- ²²T. J. Lane, G. R. Bowman, K. Beauchamp, V. A. Voelz, and V. S. Pande, *J. Am. Chem. Soc.* **133**, 18413 (2011).
- ²³G. R. Bowman, V. A. Voelz, and V. S. Pande, *J. Am. Chem. Soc.* **133**, 664 (2011).
- ²⁴G. Perez-Hernandez, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, *J. Chem. Phys.* **139**, 015102 (2013).
- ²⁵N. Stanley, S. Esteban-Martin, and G. D. Fabritiis, *Nat. Commun.* **5**, 6272 (2014).
- ²⁶M. Held, P. Metzner, J. Prinz, and F. Noé, *Biophys. J.* **100**, 701 (2011).
- ²⁷D. Huang and A. Caflisch, *PLoS Comput. Biol.* **7**, e1002002 (2011).
- ²⁸D.-A. Silva, G. R. Bowman, A. Sosa-Peinado, and X. Huang, *PLoS Comput. Biol.* **7**, e1002054 (2011).
- ²⁹G. Bowman and P. Geissler, *Proc. Natl. Acad. Sci. U. S. A.* **109**, 11681 (2012).
- ³⁰N. Plattner and F. Noé, *Nat. Commun.* **6**, 7653 (2015).
- ³¹J. Chodera, N. Singhal, V. Pande, K. Dill, and W. Swope, *J. Chem. Phys.* **126**, 155101 (2007).
- ³²C. R. Schwantes and V. S. Pande, *J. Chem. Theory Comput.* **9**, 2000 (2013).
- ³³M. Sarich, F. Noé, and C. Schütte, *SIAM Multiscale Model. Simul.* **8**, 1154 (2010).
- ³⁴F. Noé and F. Nüske, *Multiscale Model. Simul.* **11**, 635 (2013).
- ³⁵F. Nüske, B. G. Keller, G. Pérez-Hernández, A. S. J. S. Mey, and F. Noé, *J. Chem. Theory Comput.* **10**, 1739 (2014).
- ³⁶N. V. Buchete and G. Hummer, *J. Phys. Chem. B* **112**, 6057 (2008).
- ³⁷G. R. Bowman, K. A. Beauchamp, G. Boxer, and V. S. Pande, *J. Chem. Phys.* **131**, 124101 (2009).
- ³⁸N. S. Hinrichs and V. S. Pande, *J. Chem. Phys.* **126**, 244101 (2007).
- ³⁹F. Noé, *J. Chem. Phys.* **128**, 244103 (2008).
- ⁴⁰J. D. Chodera and F. Noé, *J. Chem. Phys.* **133**, 105102 (2010).
- ⁴¹C. Schütte, F. Noé, J. Lu, M. Sarich, and E. Vanden-Eijnden, *J. Chem. Phys.* **134**, 204105 (2011).
- ⁴²B. Trendelkamp-Schroer and F. Noé, *J. Chem. Phys.* **138**, 164113 (2013).
- ⁴³W. F. Van Gunsteren and H. J. C. Berendsen, *Mol. Simul.* **1**, 173 (1988).
- ⁴⁴M. Tuckerman, B. J. Berne, and G. J. Martyna, *J. Chem. Phys.* **97**, 1990 (1992).
- ⁴⁵B. Trendelkamp-Schroer and F. Noé, e-print [arXiv:1409.6439](https://arxiv.org/abs/1409.6439) (2014).
- ⁴⁶P. Metzner, F. Noé, and C. Schütte, *Phys. Rev. E* **80**, 021106 (2009).
- ⁴⁷S. Bacallado, J. D. Chodera, and V. Pande, *J. Chem. Phys.* **131**, 045106 (2009).
- ⁴⁸J. Besag and D. Mondal, *Biometrics* **69**, 488 (2013).
- ⁴⁹A. Bittracher, P. Koltai, and O. Junge, "Pseudo generators of spatial transfer operators," *SIAM J. Appl. Dyn. Syst.* **14**(3), 1478 (2015).
- ⁵⁰B. N. Parlett, *The Symmetric Eigenvalue Problem*, Classics in Applied Mathematics (Society for Industrial and Applied Mathematics, Philadelphia, 1998).
- ⁵¹S. Röblitz and M. Weber, *Adv. Data Anal. Classif.* **7**, 147 (2013).
- ⁵²P. Metzner, C. Schütte, and E. Vanden-Eijnden, *Multiscale Model. Simul.* **7**, 1192 (2009).
- ⁵³A. Berezhkovskii, G. Hummer, and A. Szabo, *J. Chem. Phys.* **130**, 205102 (2009).
- ⁵⁴W. E and E. Vanden-Eijnden, *J. Stat. Phys.* **123**, 503 (2006).
- ⁵⁵W. Janke, in *Quantum Simulations of Complex Many-Body Systems: From Theory to Algorithms*, NIC Series, edited by J. Grotendorst, D. Marx, and A. Muramatsu (John von Neumann Institute for Computing, Jülich, 2002), Vol. 10, pp. 423–445.
- ⁵⁶F. Noé, preprint, <http://publications.mi.fu-berlin.de/1699/> (2015).
- ⁵⁷T. W. Anderson and L. A. Goodman, *Ann. Math. Stat.* **28**, 89 (1957).
- ⁵⁸S. Ji, Y. Xue, and L. Carin, *IEEE Trans. Signal Process.* **56**, 2346 (2008).
- ⁵⁹P. Eastman, M. S. Friedrichs, J. D. Chodera, R. J. Radmer, C. M. Bruns, J. P. Ku, K. A. Beauchamp, T. J. Lane, L.-P. Wang, D. Shukla, T. Tye, M. Houston, T. Stich, C. Klein, M. R. Shirts, and V. S. Pande, *J. Chem. Theory Comput.* **9**, 461 (2013).
- ⁶⁰K. Lindorff-Larsen, S. Piana, K. Palmo, P. Maragakis, J. L. Klepeis, R. O. Dror, and D. E. Shaw, *Proteins: Struct., Funct., Bioinf.* **78**, 1950 (2010).
- ⁶¹W. L. Jorgensen, J. Chandrasekhar, J. D. Madura, R. W. Impey, and M. L. Klein, *J. Chem. Phys.* **79**, 926 (1983).
- ⁶²C. J. Geyer, *Stat. Sci.* **7**, 473 (1992).
- ⁶³D. E. Shaw, P. Maragakis, K. Lindorff-Larsen, S. Piana, R. Dror, M. Eastwood, J. Bank, J. Jumper, J. Salmon, Y. Shan, and W. Wriggers, *Science* **330**, 341 (2010).
- ⁶⁴M. K. Scherer, B. Trendelkamp-Schroer, F. Paul, G. Perez-Hernandez, M. Hoffmann, N. Plattner, C. Wehmeyer, J.-H. Prinz, and F. Noé, "PyEMMA 2: A software package for estimation, validation, and analysis of Markov models," *J. Chem. Theory Comput.* (published online)
- ⁶⁵S. Sawyer, *The Metropolitan-Hastings Algorithm and Extensions* (Washington University, 2006).
- ⁶⁶L. Devroye, *Non-Uniform Random Variate Generation* (Springer, 1986).