xTRAM: Estimating Equilibrium Expectations from Time-Correlated Simulation Data at Multiple Thermodynamic States

Antonia S. J. S. Mey,[*] Hao Wu,[†] and Frank Noé[‡]

Freie Universität Berlin, Arnimallee 6, 14195 Berlin, Germany
(Received 22 January 2013; revised manuscript received 26 June 2014; published 29 October 2014)

Computing the equilibrium properties of complex systems, such as free energy differences, is often hampered by rare events in the dynamics. Enhanced sampling methods may be used in order to speed up sampling by, for example, using high temperatures, as in parallel tempering, or simulating with a biasing potential such as in the case of umbrella sampling. The equilibrium properties of the thermodynamic state of interest (e.g., lowest temperature or unbiased potential) can be computed using reweighting estimators such as the weighted histogram analysis method or the multistate Bennett acceptance ratio (MBAR). weighted histogram analysis method and MBAR produce unbiased estimates, the simulation samples from the global equilibria at their respective thermodynamic states—a requirement that can be prohibitively expensive for some simulations such as a large parallel tempering ensemble of an explicitly solvated biomolecule. Here, we introduce the transition-based reweighting analysis method (TRAM)—a class of estimators that exploit ideas from Markov modeling and only require the simulation data to be in local equilibrium within subsets of the configuration space. We formulate the expanded TRAM (xTRAM) estimator that is shown to be asymptotically unbiased and a generalization of MBAR. Using four exemplary systems of varying complexity, we demonstrate the improved convergence (ranging from a twofold improvement to several orders of magnitude) of xTRAM in comparison to a direct counting estimator and MBAR, with respect to the invested simulation effort. Lastly, we introduce a random-swapping simulation protocol that can be used with xTRAM, gaining orders-of-magnitude advantages over simulation protocols that require the constraint of sampling from a global equilibrium.

DOI: 10.1103/PhysRevX.4.041018

Subject Areas: Biological Physics,
Computational Physics,
Statistical Physics

I. INTRODUCTION

The successful prediction of equilibrium behavior of complex systems is of critical importance in computational physics. Often, rare events in the system's dynamics make such estimates through direct simulations impractical. For this reason, the past 20 years have seen vast progress in computational techniques used for efficient sampling of rare-event systems with complex energy landscapes. These developments were, in particular, driven by the study of systems such as spin glasses [1,2], quantum frustrated spin systems [3,4], QCD [5–7], and molecular dynamics (MD) of biomolecules [8,9].

Commonly used approaches include generalized ensemble methods such as simulated tempering (ST) [10], parallel tempering (PT) [8,9,11], or replica-exchange molecular dynamics (REMD) [12]. Generalized ensemble methods can greatly improve the convergence over direct simulation for systems with high energy barriers but relatively few degrees of freedom [13–15]. The speed of convergence depends on the overlap of energy distributions between adjacent temperatures, and thus efforts have been made in choosing optimal temperature distributions [8,16–18]. Unfortunately, the number of replicas needed to fill the relevant range of temperatures increases with the number of degrees of freedom of the system and produces expensive computational requirements for many-body systems such as explicitly solvated molecules.

Once a multi-ensemble simulation is generated, there are different estimator options that can be used for the analysis of the simulation data in order to extract information such as free-energy differences between conformations of interest. The simplest estimator is to bin the simulation data (in a single-order parameter or using clusters of a high-dimensional parameter space) at the temperature of interest and count the number of occurrences of each of the discrete states. We will refer to this estimation method as the direct counting estimator. An improvement over direct counting of single-temperature histograms is the weighted histogram

[*]antonia.mey@fu-berlin.de

[†]hwu@zedat.fu-berlin.de

[‡]To whom all correspondence should be addressed.
frank.noe@fu-berlin.de

analysis method (WHAM) [19,20]. WHAM makes use of data from all simulated temperatures by reweighting them to the target temperature via the Boltzmann density. The traditional WHAM formulation in multitemperature ensembles requires one to discretize the system's potential energy in order to formulate a reweighting factor for each energy bin [21]. A formulation of WHAM that avoids the potential-energy binning is given in Ref. [22]. References [23] generalizes this concept and derives the multistate Bennett acceptance ratio (MBAR) that produces statistically optimal estimates of equilibrium properties given a set of equilibrium data.

All of the above estimators assume that the data are sampled from global equilibrium; i.e., simulations at all temperatures have entirely decorrelated from their starting configurations. This often requires discarding large amounts of data for producing unbiased estimates. This can lead to very long simulation times in order to obtain an uncorrelated sample.

Here, we combine ideas from reweighting estimators [24] and Markov-model theory [25–30] in order to derive a transition-based reweighting analysis method (TRAM) for estimating equilibrium properties from simulation data at multiple thermodynamic states. TRAM differs from established reweighting estimators such as direct counting, WHAM, or MBAR, in that it uses conditional transition probabilities between the discrete states or bins and therefore does not require the underlying data to be in global equilibrium. Thus, TRAM can achieve unbiased equilibrium estimates for data that are not yet in global equilibrium, allowing accurate estimates to be obtained with sometimes orders-of-magnitude-smaller simulation times compared to established estimators.

Markov models are usually used for predicting long-term kinetics from short-time simulation data [31], requiring the use of sufficiently long lag times when computing the conditional transition probabilities [30]. Therefore, extracting kinetics from generalized ensemble simulations is difficult. If desired, one either has to limit the lag time to the short contiguous simulations times [32], or one has to reweight entire trajectories [33,34]. However, a transition matrix can be used to estimate the equilibrium distribution of a system without requiring long lag times [30]. At a given temperature T^I, the corresponding transition matrix $\mathbf{P}^I$ provides an estimate of the equilibrium distribution $\boldsymbol{\pi}^I$ as its dominant eigenvector. However, in order to exploit the existence of high temperatures in the simulation, an estimator must be constructed that connects the different temperatures in a rigorous way. This leads to the proposal of a TRAM.

TRAM can also be employed to get estimates from multiple biased simulations, such as umbrella sampling [35] or metadynamics [36], although here we will focus on applications using multitemperature ensembles. In general, by TRAM we refer to a class of estimators with the following behavior:

(1) Given simulations at different thermodynamic states $I = 1, \ldots, m$ (temperatures, bias potentials, …), and a configuration-space discretization into discrete states $i = 1, \ldots, n$ (binning along an order parameter or clustering of a high-dimensional space), harvest the following statistics:
   (a) At each thermodynamic state I, the number of transitions c_{ij}^I is observed between configuration states i, j at a time lag τ (here, usually the data-storage interval).
   (b) For each sample $\mathbf{x}$ along the trajectory, the probability weight $\mu^J(\mathbf{x})$, this configuration $\mathbf{x}$ would be attributed, in each thermodynamic state, $J = 1, \ldots, m$.
(2) Compute an optimal estimate of the equilibrium probability π_i^I for all configuration states i at all thermodynamic states I.

With the help of the equilibrium probabilities π_i^I, other equilibrium expectations can be computed. Because of property (1a), TRAM is a "transition-based" estimator rather than a histogram method. Because of property (1b), TRAM is also a "reweighting" estimator. TRAM estimators do not depend on actual temperature transitions in the generalized ensemble. Rather, all configurations visited during the simulation will be used to estimate transition probabilities between thermodynamic states.

Different implementations of TRAM estimators and formulations of their optimality may be possible; we therefore consider TRAM to be a class of estimators rather than a unique method. In this paper, we propose a TRAM estimator that formally constructs an expanded ($mn \times mn$) transition matrix in the joint space of all m thermodynamic states and n configuration states. Therefore, the present estimator is called expanded TRAM (xTRAM). The stationary eigenvector of the xTRAM transition matrix contains the equilibrium probabilities at all thermodynamic states.

While simulation protocols such as ST, PT, and REMD require a strong overlap of energy distributions between neighboring temperatures to be efficient for sampling, this is much less relevant for the usefulness of TRAM estimators, as the reweighting factors are useful information even when the transition probabilities between thermodynamic states are small. It is thus tempting to design new simulation protocols that achieve more efficient sampling by sacrificing the asymptotic global equilibrium property achieved by ST, PT, and REMD but can still yield unbiased estimates of equilibrium probabilities when used in conjunction with TRAM estimators. In this paper, we make a first attempt to this end and propose the random-swapping (RS) simulation method. RS achieves rapid mixing between a set of replicas that would be too sparsely distributed for ST, PT, or REMD because it exchanges without the Metropolis-Hastings acceptance step. The associated violation of global equilibrium can be approximately recovered by xTRAM

because local equilibrium is guaranteed by adjusting lag times τ during the estimation accordingly.

xTRAM is shown to be asymptotically unbiased. Moreover, we show that xTRAM is a generalization of MBAR, and they converge to identical estimators for the free-energy differences between thermodynamic states in the limit of global equilibrium and to identical estimators for general equilibrium expectations in the limit of global equilibrium and large statistics.

Using xTRAM and, in particular, the combination of xTRAM and RS, estimates of equilibrium properties of complex dynamical systems can be obtained with orders-of-magnitude-fewer simulation data than that required by conventional estimators. We illustrate the performance of TRAM, MBAR, and direct counting on two bistable model potentials and two explicitly solvated peptides simulated with molecular dynamics simulations.

II. THEORY AND METHODS

A. Scope

A configuration $\mathbf{x}$ is a point in configuration space Ω containing all quantities characterizing the instantaneous state of the system, such as particle positions or spins, system volume (in constant-pressure ensembles), and particle numbers (in ensembles of constant chemical potential).

We consider a set of simulation trajectories, each sampling from Ω, at a given thermodynamic state I. A thermodynamic state, denoted here as capital superscript letters I, J, and K, is characterized by its thermodynamic variables, such as temperature, pressure, or chemical potential. The dynamics are assumed to fulfill microscopic detailed balance at their respective thermodynamic states.

We consider Ω to be partitioned into subsets S_i such that $\Omega = \bigcup_{i=1}^{n} S_i$. We subsequently refer to subsets S_i as configuration states and index them by small subscript letters i, j, and k. This discretization serves to distinguish the states that are relevant to the analysis. As such, it may consist of a fine discretization of an order parameter of interest (e.g., magnetization in an Ising model) or a Voronoi partition obtained from clustering molecular dynamics data, as is frequently used for the construction of Markov models.

TRAM estimators will use statistics from transitions among configuration states but also exploit the fact that the statistical weight of a configuration $\mathbf{x}$ can be reweighted between thermodynamic states. Consider the following two examples: (1) In PT or REMD simulations, the weighting occurs through the different temperatures: Given a configuration $\mathbf{x}$ with potential energy $U(\mathbf{x})$, the statistical weight at temperature T^I is proportional to $e^{-u^I(\mathbf{x})}$ using the reduced potential energy

$$u^I(\mathbf{x}) = \frac{U(\mathbf{x})}{k_B T^I}, \tag{1}$$

with Boltzmann constant k_B. (2) When the simulation setup contains multiple biased simulations, such as in umbrella sampling or metadynamics, there is usually a unique temperature T, but different potentials $U^I(\mathbf{x}) = U(\mathbf{x}) + B^I(\mathbf{x})$ are employed where $B^I(\mathbf{x})$ is the Ith bias potential. Then, the statistical weights in each of these potentials is given by $e^{-u^I(\mathbf{x})}$, with

$$u^I(\mathbf{x}) = \frac{U(\mathbf{x}) + B^I(\mathbf{x})}{k_B T}. \tag{2}$$

We generalize this concept following the example of Ref. [23]. In a thermodynamic state I, defined by a particular combination of the potential energy function U^I, pressure p^I, chemical potentials μ_s^I of chemical species s, and temperature T^I, our system has a reduced potential defined by

$$u^I(\mathbf{x}) = \frac{U^I(\mathbf{x}) + p^I V(\mathbf{x}) + \sum_s \mu_s^I N_s(\mathbf{x})}{k_B T^I}. \tag{3}$$

Here, $V(\mathbf{x})$ is the volume of the system in configuration $\mathbf{x}$ and $N_s(\mathbf{x})$ counts the particle numbers of species s at configuration $\mathbf{x}$. The probability density of configuration $\mathbf{x}$ can, for any arbitrarily chosen thermodynamic state, be expressed as

$$\rho^I(\mathbf{x}) = \frac{1}{Z^I} e^{-u^I(\mathbf{x})}, \tag{4}$$

where Z^I is the partition function of thermodynamic state I:

$$Z^I = \int_\Omega d\mathbf{x} e^{-u^I(\mathbf{x})}. \tag{5}$$

The partition function of configuration state i at thermodynamic state I is

$$Z_i^I = \int_{S_i} d\mathbf{x} e^{-u^I(\mathbf{x})}. \tag{6}$$

B. Aims

Next, we define the quantities that we would like to estimate using TRAM. The reduced free energy of thermodynamic state I, f^I, and the reduced free energy of configuration state i at thermodynamic state I, f_i^I, here termed the configuration free energy, are defined as

$$f^I := -\ln Z^I, \tag{7}$$

$$f_i^I := -\ln Z_i^I. \tag{8}$$

The equilibrium probability of configuration state i, given that the system is at thermodynamic state I, is

$$\pi_i^I := \frac{Z_i^I}{Z^I} = e^{f^I - f_i^I}. \qquad (9)$$

Finally, we are interested in computing expectation values of arbitrary functions of configuration state $A(\mathbf{x})$:

$$\langle A \rangle_I = \int_\Omega d\mathbf{x} \rho^I(\mathbf{x}) A(\mathbf{x}). \qquad (10)$$

The multistate Bennett acceptance ratio estimator [23] can provide statistically optimal estimates for quantities in Eqs. (7)–(10) when all samples $\mathbf{x}$ used from the set of simulation data are independently drawn from the global equilibrium distributions at the respective thermodynamic states. This requirement can induce a large statistical inefficiency that we will attempt to avoid by deriving an estimator that does not rely on global equilibrium.

C. xTRAM

The expanded TRAM estimator is based on the idea of constructing a Markov-model-like transition process in an expanded space whose states are defined by combinations of configuration states and thermodynamic states. An expanded state is the pairing of thermodynamic state I and configuration state i. We use the convention of ordering all expanded vectors and matrices first in blocks of an equal thermodynamic state and, within each such block, by a configuration state.

At a given thermodynamic state I, the matrix $\mathbf{C}^I = (c_{ij}^I)$ contains the number of transitions that have been observed in the data between pairs of configuration states i and j. The diagonal matrix $\mathbf{B}^{IJ} = \text{diag}(b_i^{IJ})$ contains transition counts for each configuration state i from thermodynamic state I to J that have not been observed but are constructed so as to obey the correct reweighting between thermodynamic states. The expanded transition count matrix $\tilde{\mathbf{N}} \in \mathbb{R}^{nm \times nm}$ is given by

$$\tilde{\mathbf{N}} = \begin{pmatrix} \mathbf{C}^1 + \mathbf{B}^{1,1} & \mathbf{B}^{1,2} & \cdots & \mathbf{B}^{1,m} \\ \mathbf{B}^{2,1} & \mathbf{C}^2 + \mathbf{B}^{2,2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{B}^{m-1,m} \\ \mathbf{B}^{m,1} & \cdots & \mathbf{B}^{m,m-1} & \mathbf{C}^m + \mathbf{B}^{m,m} \end{pmatrix}. \qquad (11)$$

$\tilde{\mathbf{N}}$ has a sparse structure given by diagonal blocks and off-diagonal bands, as indicated below:

$$\underbrace{\begin{pmatrix} \blacksquare & & & \\ & \ddots & & \\ & & \ddots & \\ & & & \blacksquare \end{pmatrix}}_{\mathbf{C}^1, \dots, \mathbf{C}^m} + \underbrace{\begin{pmatrix} \diagdown & \diagdown & \cdots & \diagdown \\ \diagdown & \diagdown & \ddots & \vdots \\ \vdots & \ddots & \ddots & \diagdown \\ \diagdown & \cdots & \diagdown & \diagdown \end{pmatrix}}_{\mathbf{B}^{1,1}, \dots, \mathbf{B}^{m,m}} = \underbrace{\begin{pmatrix} \blacksquare & \diagdown & \cdots & \diagdown \\ \diagdown & \blacksquare & \ddots & \vdots \\ \vdots & \ddots & \ddots & \diagdown \\ \diagdown & \cdots & \diagdown & \blacksquare \end{pmatrix}}_{\tilde{\mathbf{N}}}. \qquad (12)$$

The expanded transition matrix $\tilde{\mathbf{P}}$ is defined as the maximum likelihood reversible transition matrix given $\tilde{\mathbf{N}}$. The key step in xTRAM is to estimate $\tilde{\mathbf{P}}$ from $\tilde{\mathbf{N}}$ so as to compute the expanded stationary distribution $\tilde{\boldsymbol{\pi}}^\top = \tilde{\boldsymbol{\pi}}^\top \tilde{\mathbf{P}}$, which has the structure

$$\tilde{\boldsymbol{\pi}}^\top = (w^1 \boldsymbol{\pi}^1, \dots, w^m \boldsymbol{\pi}^m), \qquad (13)$$

consisting of subvectors, $\boldsymbol{\pi}^I = (\pi_1^I, \dots, \pi_n^I)$, each containing the normalized equilibrium probabilities of configuration states i given that the system is at thermodynamic state I. The weights w^I, normalized to $\sum_I w^I = 1$, scale all subvectors such that the expanded equilibrium vector is also normalized, $\sum_i \sum_I w^I \pi_i^I = 1$. Figure 1 illustrates how different thermodynamic and configuration states are related by xTRAM quantities.

Data preparation and configuration-state transition counts.—We process all trajectory data as follows. Each sample $\mathbf{x}$ occurring in a trajectory at time t is selected when a successor sample $\mathbf{y}$ at time $t + \tau$ exists such that the trajectory fragment $\mathbf{x} \rightarrow \mathbf{y}$ is generated using the same dynamics (i.e., without intermediate changes of the thermodynamic state).

All such samples $\mathbf{x}$ are sorted into sets S_i^I according to their configuration state i and thermodynamic state I. The configuration-state transition counts
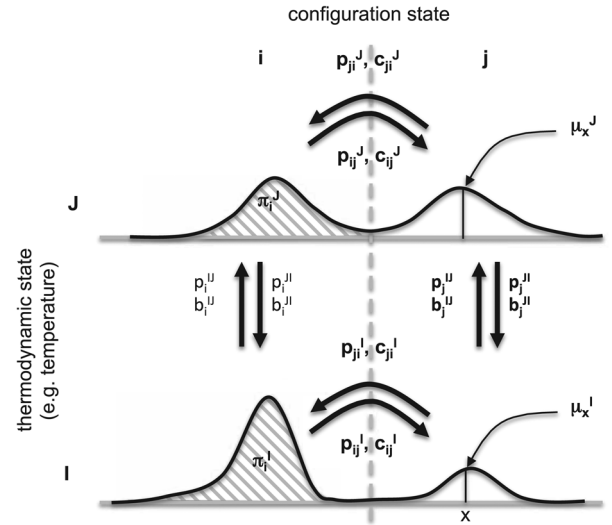

configuration state

FIG. 1. Illustration of the expanded transition process in xTRAM and the symbols used: c_{ij}^I and p_{ij}^I are transition counts and probabilities between different configuration states i, j at the same thermodynamic state I. b_i^{IJ} and p_i^{IJ} are transition counts and probabilities between different thermodynamic states I, J at the same configuration state i. The latter are constructed such that reweighting to the equilibrium densities $\mu_\mathbf{x}^I$, $\mu_\mathbf{x}^J$ of configurations $\mathbf{x}$ at different thermodynamic states I, J, occurs.

$$c_{ij}^I = |\{(\mathbf{x} \in S_i^I, \mathbf{y} \in S_j)\}|$$

count the number of transitions from all samples $\mathbf{x}$ in S_i^I to target configuration state j ($\mathbf{y}$ itself can be at any thermodynamic state as long as the dynamics to reach $\mathbf{y}$ from $\mathbf{x}$ is realized at thermodynamic state I). These counts yield m count matrices $\mathbf{C}^1, \dots, \mathbf{C}^m$.

We define the total counts N, total counts at thermodynamic state I, N^I, and total counts at thermodynamic state I and configuration state i:

$$N_i^I := \sum_j c_{ij}^I, \tag{14}$$

$$N^I := \sum_i N_i^I, \tag{15}$$

$$N := \sum_I N^I. \tag{16}$$

Note that the xTRAM estimations will not depend on the choice of τ, provided that the count matrices $\mathbf{C}^I$ are obtained from simulations that ensure a local equilibrium within each starting state S_i^I. If this is the case, τ can be chosen arbitrarily short, e.g., equal to the interval at which the sampled configuration is saved. In practice, deviations from local equilibrium can be significant for certain simulation setups and for poor choices of the discretization, but they can be compensated by using longer lag times (see random-swapping results below). When local equilibrium is not ensured by the simulation setup, TRAM estimates should be performed for a series of τ values, and then we choose the smallest τ value for which converged estimates are obtained.

Thermodynamic-state transition counts.—The elements I, J of the thermodynamic-state count matrix at configuration state i are constructed by attributing to each sample $\mathbf{x}$ at thermodynamic state I a single count that is split to all target thermodynamic states J proportional to the respective probability $p^{IJ}(\mathbf{x})$:

$$b_i^{IJ} = \sum_{\mathbf{x} \in S_i^I} p^{IJ}(\mathbf{x}), \tag{17}$$

where $p^{IJ}(\mathbf{x})$ is the transition probability to thermodynamic state J given that the system is at configuration $\mathbf{x}$ and thermodynamic state I. In the example of a multitemperature simulation, $p^{IJ}(\mathbf{x})$ can be interpreted as the probability for which a hypothetical simulated tempering (ST) trial from temperature I to J would be accepted at configuration $\mathbf{x}$. In a more general setting, the transition probabilities between thermodynamic states can be derived from Bennett's acceptance ratio [24]. From $\sum_J p^{IJ}(\mathbf{x}) = 1$,

it directly follows that $N_i^I = \sum_J b_i^{IJ}$. Different choices for $p^{IJ}(\mathbf{x})$ are possible as long as they respect detailed balance.

The statistical weights w^I of thermodynamic states in the expanded ensemble are chosen as the fraction of samples seen at each thermodynamic state I:

$$w^I := \frac{N^I}{N}. \tag{18}$$

We will show later that this choice leads to a statistically optimal estimator. As an example, consider a replica-exchange simulation; all replicas I are propagated in parallel and therefore $N^1 = \dots = N^m$, resulting in equal weights $w_I = 1/m$. When the input data stem from simulated tempering simulations between different temperatures I, the fraction of time spent at each temperature depends on the choice of the simulated tempering weights [10] and could therefore be different. With choices given by Eq. (18), the absolute probability of configuration $\mathbf{x}$ in the expanded ensemble is

$$\mu^I(\mathbf{x}) = w^I \rho^I(\mathbf{x}) = \frac{N^I}{N} e^{f^I - u^I(\mathbf{x})}. \tag{19}$$

In order for the thermodynamic-state transition process to sample from the correct statistical weights, it must fulfill the detailed balance equations:

$$\mu^I(\mathbf{x}) p^{IJ}(\mathbf{x}) = \mu^J(\mathbf{x}) p^{JI}(\mathbf{x}). \tag{20}$$

Various choices for $p^{IJ}(\mathbf{x})$ can be made that meet these constraints. It turns out that the statistically optimal choice is to a thermodynamic state J according to its equilibrium probability of that state in the expanded ensemble, i.e.,

$$p^{IJ}(\mathbf{x}) = \frac{\mu^J(\mathbf{x})}{\sum_K \mu^K(\mathbf{x})} = \frac{N^J e^{f^J - u^J(\mathbf{x})}}{\sum_K N^K e^{f^K - u^K(\mathbf{x})}}. \tag{21}$$

The choices (19) and (21) obviously fulfill the detailed balance equations (20). Using Eq. (21) in an implementation requires shifting the absolute energy value in order to avoid numerical overflows when evaluating the exponential (see the Appendix of Ref. [23] for a discussion).

An alternative choice for the thermodynamic-state transition process is the Metropolis rule, which is easier to implement and produces indistinguishable results compared to choice (21) in our applications:

$$p^{IJ}(\mathbf{x}) = \frac{N^I}{N} \min \left\{ 1, \frac{N^J e^{f^J - u^J(\mathbf{x})}}{N^I e^{f^I - u^I(\mathbf{x})}} \right\} \quad I \neq J$$

$$p^{II}(\mathbf{x}) = 1 - \sum_{J \neq I} p^{IJ}(\mathbf{x}). \tag{22}$$

Free energies.—In order to compute $p^{IJ}(\mathbf{x})$ in Eq. (17) using Eq. (21) or (22), an estimate of the free energies f^I is needed. At initialization, the f^I's are estimated using Bennett's acceptance ratio [24]. To this end, the thermodynamic states at which simulations have been conducted are sorted in a sequence $(1, \ldots, I, I+1, \ldots, m)$, e.g., ascending temperatures. The free energies are then initially set to

$$f^1 := 0, \tag{23}$$

$$f^{I+1} := f^I - \ln \frac{\frac{1}{N^I} \sum_{\mathbf{x} \in S^I} \min\{1, e^{u^I(\mathbf{x}) - u^{I+1}(\mathbf{x})}\}}{\frac{1}{N^{I+1}} \sum_{\mathbf{x} \in S^{I+1}} \min\{1, e^{u^{I+1}(\mathbf{x}) - u^I(\mathbf{x})}\}}, \tag{24}$$

where the sample averages are taken over all samples in a given thermodynamic state, as denoted by $\mathbf{x} \in S^I$. In subsequent iterations, we can update the free energies using

$$f^{I,\text{new}} := f^I - \ln \frac{N}{N^I} \sum_i \tilde{\pi}_i^I. \tag{25}$$

The iteration is converged when in the expanded equilibrium distribution, $\tilde{\pi}$ has weights that are equal according to the target values w^I: $\sum_i \tilde{\pi}_i^I = (N^I/N)$ for all I. Equation (25) will adapt f^I until this equilibrium is achieved (see Supplemental Material [37] for details).

Estimation of the equilibrium distribution.—In every iteration, we obtain a transition-count matrix possessing the sparsity structures sketched in Eq. (12). Because the theory is based on a transition matrix fulfilling detailed balance, we can estimate $\tilde{\mathbf{P}}$ using the reversible transition matrix estimator described in Ref. [30] which also provides the expanded equilibrium distribution $\tilde{\boldsymbol{\pi}}$ as a by-product.

However, we can derive a simple direct estimator for $\tilde{\boldsymbol{\pi}}$ without going through $\tilde{\mathbf{P}}$ (see Supplemental Material [37]). Let x_i^I be variables that are iterated to approximate π_i^I. Iteratively updating according to (26), x_i^I converges towards the maximum likelihood estimate of π_i^I:

$$x_i^{I,\text{new}} = \frac{1}{2} \sum_{j=1}^n \frac{c_{ij}^I + c_{ji}^I}{\frac{N_i^I}{w^I \pi_i^I} + \frac{N_j^I}{w^I \pi_j^I}} + \frac{1}{2} \sum_{J=1}^m \frac{b_i^{IJ} + b_i^{JI}}{\frac{N_i^I}{w^I \pi_i^I} + \frac{N_i^J}{w^J \pi_i^J}}, \tag{26}$$

$$\pi_i^{I,\text{new}} = \frac{x_i^I}{\sum_{j=1}^n x_j^I}. \tag{27}$$

The xTRAM estimator is summarized in algorithm 1.

Algorithm 1: xTRAM Algorithm for estimating the free energies f^I and equilibrium probabilities π_i^I.

---

1. Compute the largest connected set from the projection of the multitemperature trajectory ensembles onto states. All vectors and matrices are defined on that connected set. For all other states, π_i^I is set to 0.

2. Initial guess of free energies: Set $f_1 := 0$ and for $I = 1, \ldots, m-1$ set

$$f^{I+1} := f^I - \ln \frac{\frac{1}{N^I} \sum_{\mathbf{x} \in S^I} \min\{1, e^{u^I(\mathbf{x}) - u^{I+1}(\mathbf{x})}\}}{\frac{1}{N^{I+1}} \sum_{\mathbf{x} \in S^{I+1}} \min\{1, e^{u^{I+1}(\mathbf{x}) - u^I(\mathbf{x})}\}}.$$

3. Compute configuration-state counts $\mathbf{C}^I = (c_{ij}^I)$. c_{ij}^I is the number of times a trajectory simulated at thermodynamic state I was found to be at configuration state i at time t, and at state j at time $t + \tau$. Define $N_i^I := \sum_j c_{ij}^I$, $N^I := \sum_i N_i^I$, $N := \sum_I N^I$.

4. Initial guess of equilibrium probabilities:

$$\pi_i^I := \frac{N_i^I}{N^I}.$$

5. Iterate to convergence of f^I:
   (a) Compute thermodynamic-state counts by

$$b_i^{IJ} := \sum_{x \in S_i^I} p^{IJ}(\mathbf{x}),$$

with $p^{IJ}(\mathbf{x})$ from Eq. (21) or (22).
   (b) Iterate to convergence of π_i^I using $w^I := N^I/N$:

$$x_i^{I,\text{new}} := \sum_j \frac{c_{ij}^I + c_{ji}^I}{\frac{N_i^I}{w^I \pi_i^I} + \frac{N_j^I}{w^I \pi_j^I}} + \sum_J \frac{b_i^{IJ} + b_i^{JI}}{\frac{N_i^I}{w^I \pi_i^I} + \frac{N_i^J}{w^J \pi_i^J}}$$

$$\pi_i^{I,\text{new}} := \frac{x_i^I}{\sum_j x_j^I}.$$

   (c) Update free energies:

$$f^I := f^I - \ln \frac{N}{N^I} \sum_i \pi_i^I$$

Estimation of arbitrary expectation functions.—Now we can derive an efficient estimator of the equilibrium expectation values $\langle A \rangle$ of an arbitrary function $A(\mathbf{x})$, as defined by Eq. (10), at an arbitrary thermodynamic state (at which it is possibly not simulated). For this, we employ Eqs. (14) and (15) in Ref. [23], treating every configuration state at every thermodynamic state as a separate MBAR thermodynamic state. We define the weights

$$g(\mathbf{x}) = \frac{e^{-u(\mathbf{x})}}{\sum_K \sum_i N_i^K e^{f_i^K - u^K(\mathbf{x})}}, \qquad (28)$$

where the configuration free energies can be computed as $f_i^I = f^I - \ln \pi_i^I$. As shown in the Supplemental Material [37], the expectation values of an arbitrary function of configuration $\langle A \rangle$ can thus be estimated as

$$\langle A \rangle = \frac{\sum_{\mathbf{x}} g(\mathbf{x}) A(\mathbf{x})}{\sum_{\mathbf{x}} g(\mathbf{x})}, \qquad (29)$$

where $\sum_{\mathbf{x}}$ runs over all samples in the data.

D. Asymptotic correctness of the xTRAM estimator

The exact transition probability between sets S_i and S_j at thermodynamic state I is given by

$$p_{ij}^I = \frac{1}{\pi_i^I} \int_{S_i} d\mathbf{x} \mu^I(\mathbf{x}) \int_{S_j} d\mathbf{y} p^I(\mathbf{x}, \mathbf{y}; \tau), \qquad (30)$$

where $p^I(\mathbf{x}, \mathbf{y}; \tau)$ is the probability density of the system to be in configuration $\mathbf{y}$ at time $t + \tau$ given that it is in configuration $\mathbf{x}$ at thermodynamic state I at time t. By definition, microscopic detailed balance holds for the dynamics at thermodynamic state I: $\mu^I(\mathbf{x}) p^I(\mathbf{x}, \mathbf{y}; \tau) = \mu^I(\mathbf{y}) p^I(\mathbf{y}, \mathbf{x}; \tau)$. Using detailed balance in Eq. (30) directly leads to

$$\pi_i^I p_{ij}^I = \pi_j^I p_{ji}^I. \qquad (31)$$

The exact thermodynamic state transition probability from thermodynamic state I to J at configuration state i is given by

$$p_i^{IJ} = \frac{1}{w^I \pi_i^I} \int_{S_i} d\mathbf{x} \mu_{\mathbf{x}}^I p_{\mathbf{x}}^{IJ}. \qquad (32)$$

Together with Eq. (20), we have detailed balance in discrete states:

$$w^I \pi_i^I p_i^{IJ} = w^J \pi_i^J p_i^{JI}. \qquad (33)$$

In the statistical limit $N \to \infty$, which can be either realized by trajectories of great length or by a large number of short trajectories, our expected transition counts converge to the following limits:

$$\hat{c}_{ij}^I = \lim_{N \to \infty} c_{ij}^I = N_i^I p_{ij}^I, \qquad (34)$$

$$\hat{b}_{ij}^I = \lim_{N \to \infty} b_{ij}^I = N_i^I p_i^{IJ}. \qquad (35)$$

Plugging these counts and the reversibility conditions (31) and (33) into the estimator of equilibrium probabilities (26), we obtain the accurate result

$$x_i^I = w^I \pi_i^I. \qquad (36)$$

Furthermore, in the statistical limit, the Bennett acceptance ratio initialization (algorithm 1, step 2.) is exact. With result (36), this estimate is not changed in algorithm 1, step 5c. Thus, the xTRAM estimator converges to unbiased estimates of all equilibrium properties (7)–(10) in the statistical limit.

E. Special cases

With one thermodynamic state, xTRAM is a Markov model.—Consider the situation in which simulations were conducted at a single thermodynamic state, such as unbiased molecular dynamics simulations of a macromolecule at a fixed temperature I. The xTRAM count matrix is now an $n \times n$ configuration-state count matrix $\mathbf{C} = (c_{ij})$.

We only have one free energy $f^I = 0$. Using Eq. (9), the configuration free energies are given by $f_i^I = -\ln \pi_i$, where π_i are the estimated equilibrium probabilities of discrete configuration states i. These equilibrium probabilities can be obtained by the special case of Eqs. (26) and (27):

$$x_i^{\text{new}} = \sum_{j=1}^{n} \frac{c_{ij} + c_{ji}}{\frac{N_i}{\pi_i} + \frac{N_j}{\pi_j}}, \qquad (37)$$

$$\pi_i^{\text{new}} = \frac{x_i}{\sum_{j=1}^{n} x_j}. \qquad (38)$$

Equations (37) and (38) are the equilibrium probability of the maximum likelihood $n \times n$ reversible transition matrix, given count matrix $\mathbf{C}$. Therefore, in the single-thermodynamic state case, our estimates are identical to those of a reversible Markov model.

Standard methods can be used to compute the maximum likelihood reversible transition matrix $\mathbf{P}$ [30,38]. However, if we wish to use $\mathbf{P}$ to extract not only stationary but kinetic information, the lag time τ used to obtain the count matrix $\mathbf{C}$ must be chosen sufficiently large in order to obtain an accurate estimate [30].

When all thermodynamic states are in global equilibrium, xTRAM is identical to MBAR in the estimation of f^I.—In order to show the relationship between TRAM and MBAR, we use the TRAM equations (25), (26), and (27), and specialize them using the MBAR assumption that each thermodynamic state is sampled from global equilibrium. This assumption can be modeled by merging all configuration states to one state. When converged, the TRAM quantities then fulfill the equations

$$\pi^I p^{IJ} = \frac{b^{IJ} + b^{JI}}{\frac{N^I}{\pi^I} + \frac{N^J}{\pi^J}}, \qquad (39)$$

$$0 = -\ln \frac{N}{N^I} \pi^I. \qquad (40)$$

By combining these equations with Eqs. (17) and (21) (see Supplemental Material [37] for details), we obtain

$$f^I = -\ln \sum_{\text{all } \mathbf{x}} \frac{e^{-u^I(\mathbf{x})}}{\sum_K N^K e^{f^K - u^K(\mathbf{x})}}, \qquad (41)$$

which is identical to the MBAR estimator for the reduced free energy of thermodynamic state I [see Eq. (11) in Ref. [23]].

The MBAR and xTRAM estimators of π_i^I are consistent.— Again using the condition that all simulations are in their respective global equilibria and tending to the statistical limit $N \to \infty$ (see Supplemental Material [37] for details), we can show that the xTRAM estimate for the equilibrium probabilities π_i^I can be written as

$$\pi_i^I = \frac{\sum_{\mathbf{x} \in S_i} \frac{e^{-u^I(\mathbf{x})}}{\sum_K N^K e^{f^K - u^K(\mathbf{x})}}}{\sum_{\text{all } \mathbf{x}} \frac{e^{-u^I(\mathbf{x})}}{\sum_K N^K e^{f^K - u^K(\mathbf{x})}}}, \qquad (42)$$

which is identical to the MBAR expectation value for π_i^I [to obtain this result, use Eqs. (14) and (15) in Ref. [23] with the indicator function on set S_i).

F. Random-swapping simulations

PT, ST, and REMD simulation protocols are constructed such that they sample from global equilibrium at all temperatures after a sufficiently long burn-in phase. Global equilibrium is ensured by constructing appropriate Metropolis acceptance criteria for temperature swaps. The disadvantage is that to ensure a good mixing rate between replicas, dense replica spacing is required—a problem that becomes increasingly difficult for systems with a large number of degrees of freedom, as is the case of biomolecular simulations of 10^5 or more atoms.

However, because of the use of transition matrices, xTRAM only requires local equilibrium within the discrete configurational states rather than global equilibrium—a much weaker requirement. It is thus tempting to consider using a simulation protocol that is much more efficient than PT, ST, and REMD while sacrificing the property that it samples from global equilibrium at all temperatures. Such a protocol would be useful if it is still possible to recover the correct stationary probabilities using xTRAM. One can consider the simple RS protocol, in which the replica makes a random walk in a predefined set of temperatures $T^1, \ldots, T^m$. Every so many MD/MC simulation steps, the replica jumps up or down in temperature with equal probability. The temperature move is always accepted, unlike in ST. In this way, temperature and configuration space can be efficiently sampled with very widely spaced replicas, providing a good set of input trajectories for xTRAM.

Because there is no Metropolis-Hastings acceptance criterion involved, the initial samples after each temperature swap are definitely out of global but also out of local equilibrium at the new temperature. While discarding an initial fragment of the data would seem to be a viable option, it turns out that, instead, using larger lag times τ appears to work much better in correcting the estimates, as established for Markov-model construction [30]. However, a solid theory for this observation has yet to be found and is beyond the scope of the current paper.

III. RESULTS

To demonstrate the validity and resulting advantage of the proposed estimator, two Langevin processes in model potentials and two explicitly solvated molecular dynamics processes are considered. In all cases, we compare three different estimators, which are the newly proposed xTRAM estimator, MBAR, and histogram counting (direct counting estimate), each applied to the same sets of data. Both accuracy and precision of all methods will be studied by evaluating the systematic and statistical errors for representative discrete states and temperatures of interest.

A. Two-well potential with solvent degrees of freedom

As a first example, we consider Langevin dynamics in an asymmetric double-well potential [Fig. 2(a)] with the corresponding stationary (Boltzmann) distribution $P(x)$ shown in Fig. 2(b) for the reduced temperature $k_B T = 1$. In order to make the system more complex, we add a set of N solvent particles. Each solvent coordinate i is subject to a harmonic potential $U(y_i) = y_i^2$, where y_i is the particle's position.

The state space is discretized into two states, corresponding to the two potential basins. We aim to estimate the equilibrium distribution of these two states from a set of different multitemperature simulation protocols in combination with any of the estimators considered (xTRAM, MBAR, and direct counting). All simulations are initiated from a local stationary distribution in state S_1, and the three different simulation protocols chosen are PT, ST, and RS simulations. With each simulation protocol, 100 independent realizations were generated, and their results are shown in Fig. 2. For all three simulation protocols, a temperature space needs to be defined, consisting of four exponentially spaced temperatures between $k_B T = 1$ and $k_B T = 10$ in reduced units, for Figs. 2(c)–2(f) and six exponentially spaced temperatures between $k_B T = 1$ and $k_B T = 15$ in reduced units for Figs. 2(g) and 2(h). The temperatures are chosen in such a way that barrier crossings at the lowest temperature are very rare events. For more details on the simulation protocols and setup, see the Supplemental Material [37].

Figures 2(c,d) and 2(e,f) show the results of ST and PT simulations with two solvent particles, respectively.
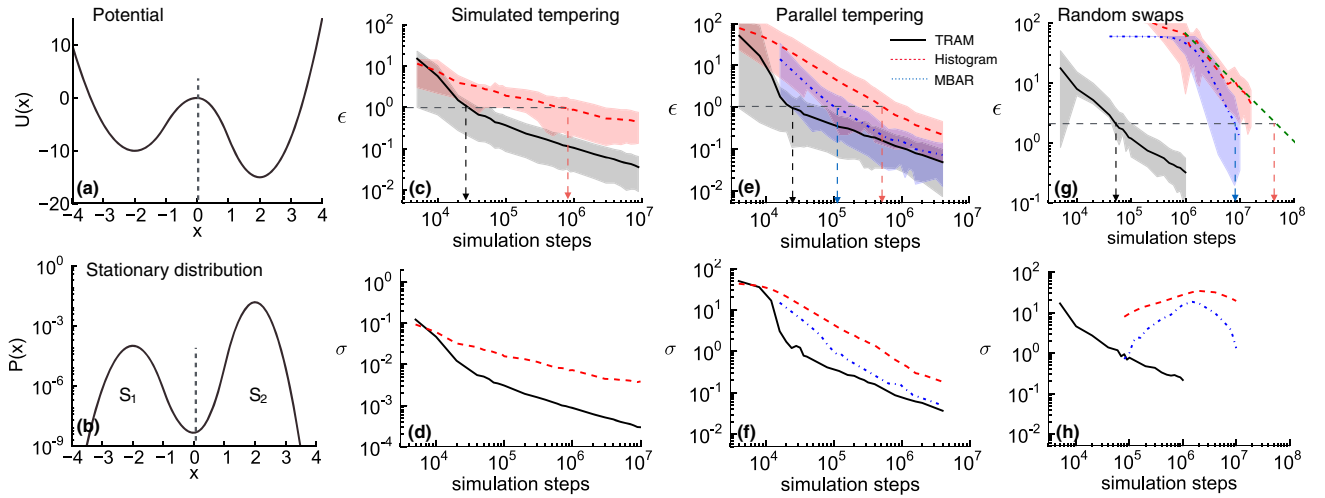
FIG. 2.   (a) Double-well potential $U(x)$. (b) Corresponding stationary distribution at a reduced temperature $k_B T = 1$. (c, d) Results of the ST simulation. (c) The relative error of π_1 with respect to (w.r.t.) the number of simulation steps for simulations with two solvent particles and four temperatures. The direct histogram estimate (red, dashed line) and xTRAM (black, continuous line) are evaluated on the same data set. (d) Standard deviation (σ) from 100 realizations w.r.t. simulation steps taken for direct count histogram (red, dashed line) and xTRAM (black, continuous line). (e, f) Results of the PT simulation. (e) Relative error π_1 w.r.t. to simulation steps from PT simulations over 100 realizations for direct counts (red, dashed line), MBAR (blue, dashed-dotted line), and xTRAM (black, continuous line). (f) Standard deviation from data in (e). (g, h) Results of a system with $N = 50$ solvent particles using PT simulations for the MBAR and a direct counting estimate and RS simulations for the xTRAM estimate. (g) Relative error as a function of the total number of simulation steps for comparing a 20 temperature replica PT simulation, analyzed with direct count histograms (red, dashed line), and MBAR (blue, dashed-dotted line), with the xTRAM (black, continuous line) analysis of a six-temperature RS simulation. The dashed green line shows a fit of $t^{-0.5}$ to the tail of the relative error. (h) Standard deviations of data in (g).

The results are displayed in the form of a log-log plot of the relative error of the estimate of π_1 and its convergence behavior with increased simulation time. The relative error is given by

$$\epsilon = \left| \frac{\pi(S_1)_{\text{exact}} - \pi(S_1)_{\text{estimate}}}{\pi(S_1)_{\text{exact}}} \right|. \qquad (43)$$

The stationary distribution, at the lowest reduced temperature of $k_B T = 1$, is obtained using all three estimators: (1) direct counting, (2) MBAR, and (3) xTRAM. Figures 2(c), 2(e), and 2(g) report averages and confidence intervals of the time-dependent relative errors computed from 100 realizations of each simulation. Figures 2(d), 2(f), and 2(h) report standard deviations (σ) of the simulation data from 100 independent realizations and their time dependence. In panels (c) and (e), the tail of the mean error for all three methods decays approximately equal to $b/\sqrt{t}$, where b is a constant related to the decorrelation time t_{corr} required to generate an uncorrelated configuration at the temperature analyzed. Arrows in panels (c) and (e) indicate a relative error of 1. In the case of the ST simulation, xTRAM outperforms direct counting by a factor of 40, and in (e), for the parallel tempering simulation, xTRAM has a gain of a factor of 5 over MBAR estimates and a factor of 25 over direct counting estimates. This means that the xTRAM estimates converge up to 40-fold faster in comparison to direct counts and at least

fivefold faster in comparison to MBAR, indicating that the decorrelation time with xTRAM can be much shorter. Consequently, less simulation time needs to be invested when the data are analyzed with xTRAM. Second, the standard deviation as seen in (d) and (f) is consistently lower for xTRAM, meaning that over independent realizations, the accuracy of the estimate is better in comparison to MBAR and direct counting.

Additional efficiency can be gained when the simple RS simulation protocol is employed instead of ST or PT simulations because then the number of replicas can be reduced such that xTRAM gives good results, while ST or PT replicas would not mix well in temperature space. In Fig. 2(g), the results of a simulation with 50 solvent particles are depicted. In order to achieve a good mixing in a PT simulation, 20 temperatures exponentially spaced in the range of 1 to 15 in reduced units need to be used, which is compared to a six-replica RS simulation (see Supplemental Material [37] for more details). The lag time τ for the evaluation could actually be chosen as small as the saving interval of the simulation in this case, resulting in the same convergence as using larger lag times. Looking at a relative error of $\epsilon = 2$, an extrapolation needs to be made to compute how many simulation steps are needed for the PT direct counting estimate. From the extrapolated convergence behavior, it is found to be around 1×10^8 simulation steps. Despite the fact that the RS protocol itself is not in equilibrium, the correct equilibrium probabilities can be

recovered when used in conjunction with xTRAM. In this case, a relative error $\epsilon = 2$ is achieved at around 1×10^5 simulation steps, indicating an efficiency gain of around 3 orders of magnitude for RS/TRAM and more than 2 orders of magnitude over MBAR with the same PT simulation data as for the direct counts. It should be stressed again that for MBAR and direct counting, only simulation sampling from a global equilibrium can be used; therefore, these estimators are not suitable to be used in conjunction with a random-swapping simulation. Figure 2(h) shows the standard deviation of the relative error from 100 realizations. Initially, the standard deviation is deceptively small for direct counts and MBAR because many of the 100 simulations have only seen state 1. The standard deviation increases, as the second state is discovered, reaching a peak, from which σ decreases again.

B. Simple protein-folding model

In order to illustrate the limitations of the method, we now discuss an example where xTRAM offers no significant advantage over the established MBAR estimator. We consider an idealized folding potential with an energetically stabilized native state and an entropically stabilized denatured state. The state space is spanned by a vector in five dimensions, such that $x \in \mathbb{R}^5$ and $r = |\mathbf{x}|$ is the distance from zero. The potential is defined as

$$U(r) = \begin{cases} -2.5(r-3)^2 & \text{if } r < 3 \\ 0.5(r-3)^3 - (r-3)^2 & \text{if } r \geq 3 \end{cases} \quad (44)$$
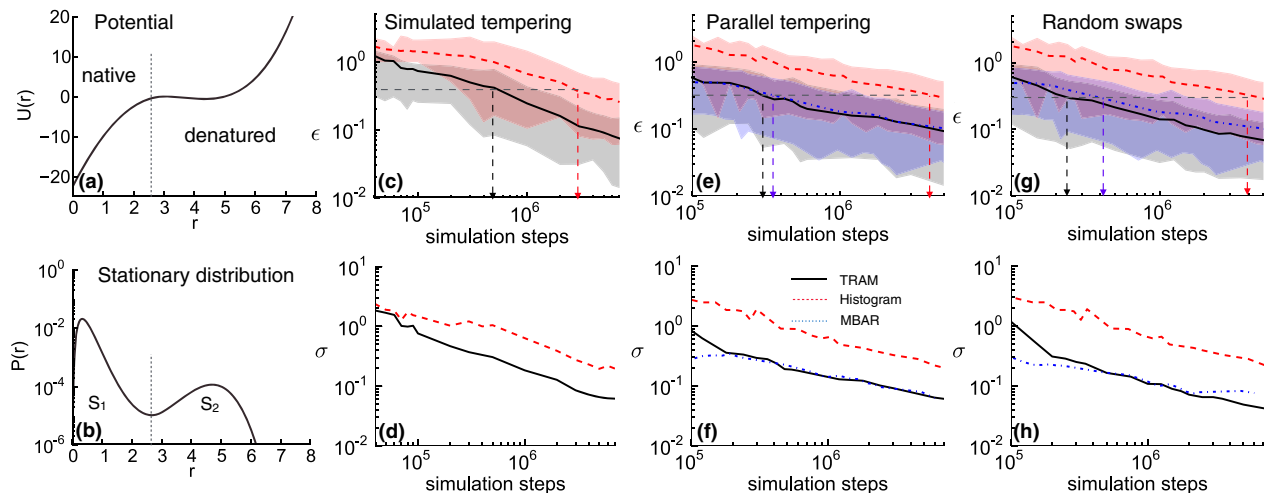
and depicted in Fig. 3(a). Again, a Langevin dynamics simulation was carried out, with more details provided in the Supplemental Material [37]. The system is discretized into two states: native and denatured, with a state boundary at $r = 2.7$, representing the distance around which the lowest probability density is observed. All simulations are initiated in the native state.

The potential and the exact stationary density $\pi(r)$ at $(1/\beta) = 1.1 k_B T$ are shown in Figs. 3(a) and 3(b), respectively. Note that the denatured state is stabilized by entropy, as more microstates are available for $r > 2.7$ than for $r < 2.7$. The convergence of the estimate of the relative error, Eq. (43), of the unfolded state is measured for a set of ST, PT, and RS simulations. Results are taken from 100 different realizations and six exponentially spaced temperatures between $(1/\beta) = 1.1 k_B T$ and $(1/\beta) = 1.7 k_B T$. Panels (c) and (d) of Fig. 3 show the results of the ST simulation, comparing the convergence of direct counting against xTRAM of the relative error of being in the denatured state. As before, xTRAM is shown by the black, continuous line, and direct counting by the dashed, red line. Arrows indicating an error level of $\epsilon = 0.3$ are used as guidance for the comparison of the convergence. Using xTRAM as the estimator for the analysis of the simulation results in a ninefold gain over direct counting. Shaded areas indicate confidence intervals. Figure 3(d) shows the convergence of the standard deviation obtained from 100 independent realizations of the ST simulation from (c). The standard deviation of the xTRAM estimate is consistently lower than for the direct counting estimate. Figures 3(e) and 3(f) summarize the results of the parallel

FIG. 3.   (a) Potential $U(r)$. (b) Corresponding stationary distribution at $(1/\beta) = 1.1 k_B T$. Results are shown in terms of the relative error ϵ of the probability of being in the "denatured" state versus the number of simulation steps required. (c, d) Results of the ST simulation. (c) xTRAM estimate (black, continuous line) and direct counting estimate (red, dashed line). Arrows indicate an $\epsilon = 0.3$ error level (d) standard deviation of the 100 realizations (c) for both estimators. (e, f) Results of the parallel tempering simulation. (e) MBAR estimate (blue, dashed-dotted line), with arrows indicating the $\epsilon = 0.2$ error level. (f) Standard deviation of data in (e) from 100 realizations. (g, h) Comparison of the RS + xTRAM estimate to the PT estimate using MBAR and direct counts in (e). Arrows indicate the $\epsilon = 0.2$ error level. (h) Standard deviation of the data in (g) over 100 realizations.

tempering simulations. Here, an 11-fold gain is observed when using xTRAM over direct counting, but MBAR and xTRAM perform almost equally as well, indicated by the arrows shown at an error level of $\epsilon = 0.2$. This behavior suggests that in this model, samples from the local and global equilibria are generated at the same rate. Figure 3(f) shows the standard deviation of the data in (e) from the 100 independently generated simulations.

Figures 3(g) and 3(h) compare the direct counting and MBAR estimate of (e) with an RS simulation using only a single replica and four exponentially spaced temperatures in the range $k_BT = [1.1-1.7]$. Now, xTRAM has a twofold gain over MBAR.

As xTRAM is a local-equilibrium generalization of MBAR, it is guaranteed to have equal or better estimation accuracy. However, the results above indicate that the accuracy gain of xTRAM over MBAR can be small in some systems. In the folding potential, this is presumably due to the fact that the different temperatures not only help in barrier crossing but also give rise to vastly different equilibrium probabilities of the folded state (stable at low temperatures) and the unfolded state (stable at high temperatures). Thus, even at short simulation times, both the folded and unfolded states are present in the replica ensemble and can successfully be reweighted using MBAR without relying on too many actual configuration-state transitions. xTRAM will gain efficiency in situations where the state space exploration is slowed down by higher friction or additional barriers, as often found in macromolecules.

C. Alanine dipeptide

In order to test the xTRAM estimator on a system with many degrees of freedom, we turn to MD simulations using an explicit solvent model. To this end, we study solvated alanine dipeptide, a small and well-studied peptide with multiple metastable states [27,39–41] and around 6000 degrees of freedom in the case of the system setup used here. Alanine dipeptide was prepared using an explicit water model and simulated in the MD software package OpenMM [42]. All the necessary simulation details are provided in the Supplemental Material [37].

The dominant conformations of this system are the different rotamers set by the dihedral angles ψ and ϕ. This means we are interested in estimating the free-energy surface in ϕ/ψ space at a low temperature of interest.

Again, we are interested in extracting the stationary probabilities of metastable basins at different temperatures. However, the system at $T = 300$ K is not very metastable; thus, we introduce an artificial metastability to the torsional angles. For this purpose, we add a potential to the minima of the ϕ and ψ dihedral angles in order to extend the time the system stays in a particular angular configuration. The ideal choices for such an additional potential are periodic von Mises potentials of the form

$$U(\delta) = \epsilon \exp\left(\frac{\kappa \cos(\delta - \delta_0)}{2\pi I_0}\right), \qquad (45)$$

where I_0 is a zeroth-order Bessel function and δ is the angle to which the additional potential is added. For more details, see the Supplemental Material [37]. We use two different types of multitemperature simulations: a set of ten independent realizations of a 33-temperature REMD simulation with temperatures exponentially spaced in the range of $T = [300$ K–600 K$]$ and a set of independent realizations of a 13-temperature RS simulation, where only every third temperature out of the REMD multitemperature ensemble was used. From the ten REMD trajectories, the last 1 ns of each was used to estimate the free-energy surface in dihedral ϕ and ψ space, as shown in Fig. 4(a). From the free-energy surface, four discrete states could be defined, numbered, and highlighted by the white boxes. All simulations were initiated in state IV and 5ns long trajectories at each temperature were generated. Discretizing simulation trajectories then allows for a stationary estimate through direct counts of the frequency of each state visited along the trajectory (in the case of the REMD simulation). The same discrete trajectories are also used for xTRAM and MBAR estimation.

For the RS simulation, the total simulation time was less, as only 13 instead of the 33 parallel replicas were simulated. Confidence intervals are indicated by the shaded regions calculated over the independent realizations of every simulation type. In order for the RS simulation to produce valid results, the lag time at which transition counts are evaluated needs to be adjusted. Here, $\tau = 1$ ps, is used, while temperature switches were carried out every 10 ps; for more details on the RS simulation, refer to the Supplemental Material [37].

Figures 4(b), 4(d), 4(f), and 4(h) show the convergence results of the REMD simulation. While all estimators yield similar (and inaccurate) estimates for very short simulation times, xTRAM exhibits considerable advantages over MBAR and direct counting after 10 ns of simulation time. The fastest-converging estimator (see below) produces stable equilibrium probabilities of about $(0.57, 0.25, 0.13, 0.1)$ for states I–IV at 100 ns of accumulated simulation time. Using REMD data, xTRAM converges to within about 10% of these values with roughly 1 order of magnitude simulation data compared to MBAR and direct counting (20 ns versus 200 ns).

Figures 4(c), 4(e), 4(g), and 4(i) compare the performance of the RS simulations when analyzed with xTRAM, with respect to the standard REMD simulations. As a result of the smaller number of replicas required and the enhanced mixing properties, another order of magnitude is gained with the RS protocol when compared to the xTRAM estimate of the REMD simulations. Since xTRAM is currently the only available estimator to unbias RS simulations, the advantage of xTRAM over MBAR and direct
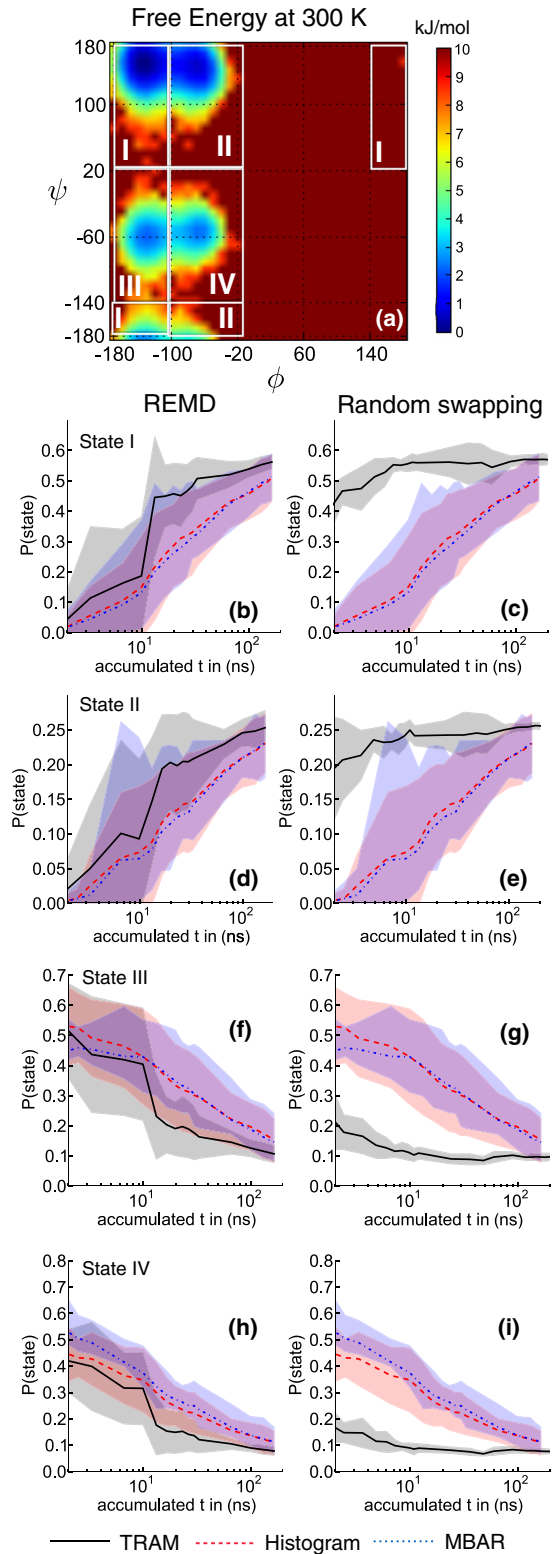
FIG. 4. (a) Free energy with a 10 kJ/mol cutoff, indicating discrete states. (b, d, f, h) REMD simulation convergence of state probability over ten realizations. Estimates are obtained from xTRAM (black, continuous line), MBAR (blue, dashed-dotted line), and direct counting (red, dashed line). (c, e, g, i) RS simulation convergence of state probabilities over five realizations, compared to REMD simulation estimates using MBAR and direct counts.

counting amounts to 2 orders of magnitude (xTRAM with RS versus MBAR with REMD). This advantage of xTRAM in conjunction with RS can be much larger for systems with many degrees of freedom, where a REMD simulation would need many closely spaced replicas, thus resulting in vast computational effort and slow exchange dynamics.

D. Deca-alanine

Finally, we consider the ten amino-acid long deca-alanine (Ala$_{10}$). This peptide is known to undergo a helix-coil transition, which has been studied extensively [43–46]. Ten independent runs of all-atom replica-exchange simulations were conducted with the GROMACS software MD package, each using 24 exponentially spaced temperatures ranging from $T = 290$ K to 400 K [47]. We ran the simulation for 40 ns total simulation time per replica and conducted ten independent realizations of these. For a more detailed description of the simulation details, see the Supplemental Material [37].

In this larger molecular system, the discretization of configuration space is no longer trivial. For this purpose, we use time-lagged independent component analysis (TICA) on the replica trajectories of the set of C$_\alpha$ distances, omitting nearest-neighbor distances along the peptide chain [48]. TICA is useful to identify the subspace in which the slowest transitions occur. Here, we chose three leading TICA coordinates and used a regular spatial clustering on these with a minimum distance cutoff of three, yielding 44 discrete clusters. This analysis was carried out using the Markov-model package EMMA [49]. See Supplemental Material [37] for more details.

xTRAM, MBAR, and histogram counting were used in order to estimate the equilibrium probabilities on the 44 discrete configuration states. In order to obtain a simple representation of the results, the equilibrium probabilities summed over all α-helical states are shown in Fig. 5(b). As before, we are interested in the analysis at the lowest simulation temperature ($T = 290$ K).

As seen in Fig. 5(a), the advantage gained from TRAM in comparison to MBAR and direct counting in this case is only about twofold. However, this can be attributed to the fact that the system does not display a very strong metastability, and the slowest time scale, computed independently with a Markov-state model on direct 290-K trajectories, is only 14 ns. Moreover, like the simple folding model above, Ala$_{10}$ has the same property that the temperatures stabilize the folded and unfolded states quite differently, leading to simultaneous observation of folded and unfolded states at early stages of the replica simulation, and also to the fact that significantly many transitions between folded and unfolded states occur only in a very small part of the replica ensemble (in the replicas around the melting point). As a result, the advantage of taking configuration-state transitions into account is smaller in this case
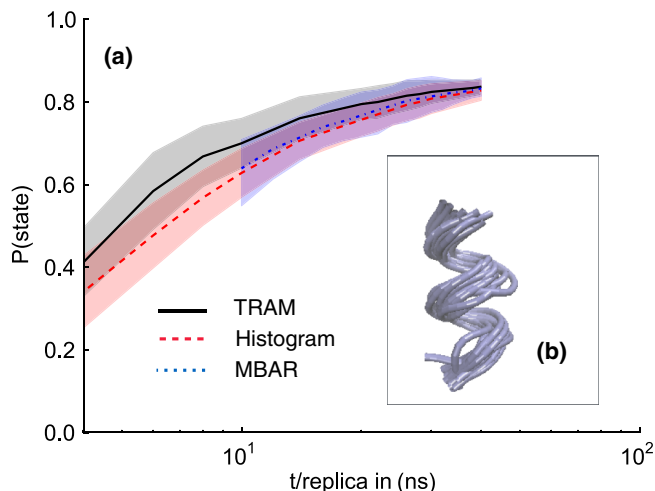
FIG. 5.   (a) Convergence of the probability of the system being in a helical state with respect to the total simulation time per replica. The xTRAM estimate is given by the continuous black line, the histogram count by the dashed red line, and the MBAR estimate by the dashed-dotted blue line. Confidence intervals are obtained from ten independent REMD simulations. (b) Structure averages of the helix state whose convergence is shown in (a).

compared to systems at which the different metastable states are present at a larger range of temperatures. As demonstrated for the simple folding model above, larger gains of computational efficiency can still be obtained by reducing the number of replicas, e.g., by employing the RS protocol in conjunction with the xTRAM estimator.

IV. DISCUSSION AND CONCLUSION

Expanded TRAM can be used to obtain estimates of equilibrium properties from any set of simulations that were conducted at different thermodynamic states, such as multiple temperatures, Hamiltonians, or bias potentials, which we demonstrated here for multiple temperature simulations. It is therefore applicable to generalized ensemble simulations such as replica-exchange and parallel or simulated tempering, as well as umbrella sampling or metadynamics. The quantities estimated can include free-energy differences, equilibrium probabilities, or equilibrium expectations of functions of a configuration state. As such, xTRAM has the same application scope as existing reweighting estimators (e.g., WHAM and MBAR).

In contrast to WHAM and MBAR, xTRAM does not assume that simulation data are generated from global equilibrium. Rather, xTRAM combines ideas from MBAR and Markov modeling into an estimator that makes use of both, Boltzmann reweighting of sampled configurations between thermodynamic states and transition count statistics between different configuration states generated by contiguous trajectory segments. Compared to MBAR, estimates obtained from xTRAM can be more accurate, as they suffer less from data that have not yet decorrelated from the

initial configurations, as well as more precise, as the statistics in the simulation data can be used more efficiently when every conditional independent transition count, rather than only every globally independent count, is useful.

xTRAM is an asymptotically correct estimator; i.e., its estimates converge to the exact values in the limit of long or many simulation trajectories. We have also shown that in the special case when simulation data *are* at global equilibrium, we can derive the MBAR equations from the expectation values of the xTRAM equations. MBAR is thus a special case of xTRAM, suggesting that the accuracy of xTRAM estimates should be at least equally good as those of MBAR estimates but may be significantly better when parts of the simulation data are not in global equilibrium. The applications shown here confirm this result, sometimes exhibiting order-of-magnitude more accurate estimates when xTRAM is employed and therefore allowing the use of shorter simulation times while maintaining the same accuracy in the estimate.

While MBAR provides statistically optimal (i.e., minimum-variance) estimates under the condition that data are in global equilibrium, it is currently not known whether xTRAM is also statistically optimal. However, the applications in this paper suggest that the variances of xTRAM estimates are, in most cases, significantly better than those of MBAR or direct counting.

An interesting aspect of xTRAM is the fact that it does not rely on the data being at global equilibrium, thus opening the door to consider new simulation methods that deliberately sacrifice global equilibrium for enhanced sampling. This feature reflects the Markov-model nature of the configuration-state statistics in xTRAM—Markov models also allow us to obtain unbiased estimates from short trajectories that are not sampling from global equilibrium, as long as they sample from local equilibrium within each configuration state. We have demonstrated this ability by using xTRAM to unbias simple random-swapping simulations that exchange temperature replicas in complete ignorance of the Metropolis acceptance probability. The hope is that with such a setup, large systems can be simulated with very few widely spaced replicas, which would be inappropriate for a PT or ST simulation that needs energy overlap between adjacent replicas. It was shown that with a sufficiently large lag time τ for evaluating the transition counts, xTRAM provided accurate estimates with such a protocol, while achieving a sampling efficiency that is orders of magnitude more efficient than classical replica-exchange or parallel tempering simulations. In the future, it will be necessary to develop a theory that quantifies the local equilibrium error made by brute-force sampling protocols such as random swapping in order to put them to use on solid ground.

An implementation of xTRAM will be available for download through the python free energy analysis toolkit (pyfeat) at [50].

ACKNOWLEDGMENTS

————————————

[1] H. G. Katzberger, M. Palassini, and A. P. Young, *Monte Carlo Simulations of Spin Glasses at Low Temperatures,* Phys. Rev. B **63,** 184422 (2001).

[2] E. Bittner and W. Janke, *Free-Energy Barriers in the Sherrington-Kirkpatrick Model,* Europhys. Lett. **74,** 195 (2006).

[3] R. G. Melko, *Simulations of Quantum XXZ Models on Two-Dimensional Frustrated Lattices,* J. Phys. Condens. Matter **19,** 145203 (2007).

[4] A. de Candia and A. Coniglio, *Spin and Density Overlaps in the Frustrated Ising Lattice Gas,* Phys. Rev. E **65,** 016132 (2001).

[5] E.-M. Ilgenfritz, W. Kerler, M. Müller-Preussker, and H. Stüben, *Parallel Tempering in Full QCD with Wilson Fermions,* Phys. Rev. D **65,** 094506 (2002).

[6] G. Burgio, M. Fuhrmann, W. Kerler, and M. Müller-Preussker, *Modified SO(3) Lattice Gauge Theory at $T \neq 0$ with Parallel Tempering: Monopole and Vortex Condensation,* Phys. Rev. D **75,** 014504 (2007).

[7] B. Joó, B. Pendleton, S. M. Pickles, Z. Sroczynski, A. C. Irving, and J. C. Sexton (UKQCD Collaboration) , *Parallel Tempering in Lattice QCD with O(a)-improved Wilson Fermions,* Phys. Rev. D **59,** 114501 (1999).

[8] U. H. E. Hansmann, *Parallel Tempering Algorithm for Conformational Studies of Biological Molecules,* Chem. Phys. Lett. **281,** 140 (1997).

[9] Y. Sugita and Y. Okamoto, *Replica-Exchange Molecular Dynamics Method for Protein Folding,* Chem. Phys. Lett. **314,** 141 (1999).

[10] E. Marinari and G. Parisi, *Simulated Tempering: A New Monte Carlo Scheme,* Europhys. Lett. **19,** 451 (1992).

[11] K. Hukushima and K. Nemoto, *Exchange Monte Carlo Method and Application to Spin Glass Simulations,* J. Phys. Soc. Jpn. **65,** 1604 (1996).

[12] C. G. Greyer, in *Computing Science and Statistics, The 23rd Symposium of the Interface* (Interface Foundation, Fairfax, 1991), pp. 156–163.

[13] D. J. Sindhikara, D. J. Emerson, and A. E. Roitberg, *Exchange Often and Properly in Replica Exchange Molecular Dynamics,* J. Chem. Theory Comput. **6,** 2804 (2010).

[14] S. Park, *Comparison of the Serial and Parallel Algorithms of Generalized Ensemble Simulations: An Analytical Approach,* Phys. Rev. E **77,** 016709 (2008).

[15] M. J. Abraham and J. E. Gready, *Ensuring Mixing Efficiency of Replica-Exchange Molecular Dynamics Simulations,* J. Chem. Theory Comput. **4,** 1119 (2008).

[16] W. Zheng, M. Andrec, E. Gallicchio, and R. M. Levy, *Simulating Replica Exchange Simulations of Protein Folding with a Kinetic Network Model,* Proc. Natl. Acad. Sci. USA **104,** 15340 (2007).

[17] E. Rosta and G. Hummer, *Error and Efficiency of Replica Exchange Molecular Dynamics Simulations,* J. Chem. Phys. **131,** 165102 (2009).

[18] E. Rosta and G. Hummer, *Error and Efficiency of Simulated Tempering Simulations,* J. Chem. Phys. **132,** 034102 (2010).

[19] A. M. Ferrenberg and R. H. Swendsen, *Optimized Monte Carlo Data Analysis,* Phys. Rev. Lett. **63,** 1195 (1989).

[20] S. Kumar, D. Bouzida, R. H. Swendsen, P. A. Kollman, and J. M. Rosenberg, *The Weighted Histogram Analysis Method for Free-Energy Calculations on Biomolecules. 1. The Method,* J. Comput. Chem. **13,** 1011 (1992).

[21] E. Gallicchio, M. Andrec, A. K. Felts, and R. M. Levy, *Temperature Weighted Histogram Analysis Method, Replica Exchange, and Transition Paths,* J. Phys. Chem. B **109,** 6722 (2005).

[22] M. Souaille and B. Roux, *Extension to the Weighted Histogram Analysis Method: Combining Umbrella Sampling with Free Energy Calculations,* Comput. Phys. Commun. **135,** 40 (2001).

[23] M. R. Shirts and J. D. Chodera, *Statistically Optimal Analysis of Samples from Multiple Equilibrium States,* J. Chem. Phys. **129,** 124105 (2008).

[24] C. H. Bennett, *Efficient Estimation of Free Energy Differences from Monte Carlo Data,* J. Comput. Phys. **22,** 245 (1976).

[25] C. Schütte, A. Fischer, W. Huisinga, and P. Deuflhard, *A Direct Approach to Conformational Dynamics Based on Hybrid Monte Carlo,* J. Comput. Phys. **151,** 146 (1999).

[26] N. Singhal and V. S. Pande, *Error Analysis and Efficient Sampling in Markovian State Models for Molecular Dynamics,* J. Chem. Phys. **123,** 204909 (2005).

[27] J. D. Chodera, W. C. Swope, J. W. Pitera, and K. A. Dill, *Long-Time Protein Folding Dynamics from Short-Time Molecular Dynamics Simulations,* Multiscale Model. Simul. **5,** 1214 (2006).

[28] F. Noé, I. Horenko, C. Schütte, and J. C. Smith, *Hierarchical Analysis of Conformational Dynamics in Biomolecules: Transition Networks of Metastable States,* J. Chem. Phys. **126,** 155102 (2007).

[29] J. D. Chodera, N. Singhal, V. S. Pande, K. A. Dill, and W. C. Swope, *Automatic Discovery of Metastable States for the Construction of Markov Models of Macromolecular Conformational Dynamics,* J. Chem. Phys. **126,** 155101 (2007).

[30] J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. D. Chodera, C. Schütte, and F. Noé, *Markov Models of Molecular Kinetics: Generation and Validation,* J. Chem. Phys. **134,** 174105 (2011).

[31] F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl, *Constructing the Full Ensemble of Folding Pathways from Short Off-Equilibrium Simulations,* Proc. Natl. Acad. Sci. USA **106,** 19011 (2009).

[32] S. Sriraman, I. G. Kevrekidis, and G. Hummer, *Coarse Master Equation from Bayesian Analysis of Replica Molecular Dynamics Simulations* J. Phys. Chem. B **109,** 6479 (2005).

[33] J. D. Chodera, W. C. Swope, F. Noé, J.-H. Prinz, M. R. Shirts, and V. S. Pande, *Dynamical Reweighting: Improved Estimates of Dynamical Properties from Simulations at Multiple Temperatures*, J. Chem. Phys. **134,** 244107 (2011).

[34] J.-H. Prinz, J. D. Chodera, V. S. Pande, W. C. Swope, J. C. Smith, and F. Noé, *Optimal Use of Data in Parallel Tempering Simulations for the Construction of Discrete-State Markov Models of Biomolecular Dynamics*, J. Chem. Phys. **134,** 244108 (2011).

[35] G. M. Torrie and J. P. Valleau, *Nonphysical Sampling Distributions in Monte Carlo Free-Energy Estimation: Umbrella Sampling*, J. Comput. Phys. **23,** 187 (1977).

[36] A. Laio and M. Parrinello, *Escaping Free Energy Minima*, Proc Natl. Acad. Sci. USA **99,** 12562 (2002).

[37] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevX.4.041018 for derivations and proofs of equations and algorithms as well as details of simulation setups.

[38] G. R. Bowman, K. A. Beauchamp, G. Boxer, and V. S. Pande, *Progress and Challenges in the Automated Construction of Markov State Models for Full Protein Systems*, J. Chem. Phys. **131,** 124101 (2009).

[39] D. S. Chekmarev, T. Ishida, and R. M. Levy, *Long-Time Conformational Transitions of Alanine Dipeptide in Aqueous Solution: Continuous and Discrete-State Kinetic Models*, J. Phys. Chem. B **108,** 19487 (2004).

[40] P. E. Smith, *The Alanine Dipeptide Free Energy Surface in Solution*, J. Chem. Phys. **111,** 5568 (1999).

[41] W.-N. Du, K. A. Marino, and P. G. Bolhuis, *Multiple State Transition Interface Sampling of Alanine Dipeptide in Explicit Solvent*, J. Chem. Phys. **135,** 145102 (2011).

[42] P. Eastman *et al.*, *OpenMM 4: A Reusable, Extensible, Hardware Independent Library for High Performance Molecular Simulation*, J. Chem. Theory Comput. **9,** 461 (2013).

[43] R. B. Best and G. Hummer, *Optimized Molecular Dynamics Force Fields Applied to the Helix/Coil Transition of Polypeptides*, J. Phys. Chem. B **113,** 9004 (2009).

[44] N. Michaud-Agrawal, E. J. Denning, T. B. Woolf, and O. Beckstein, *MDAnalysis: A Toolkit for the Analysis of Molecular Dynamics Simulations*, J. Comput. Chem. **32,** 2319 (2011).

[45] E. J. Sorin and V. S. Pande, *Exploring the Helix-Coil Transition via All-Atom Equilibrium Ensemble Simulations*, Biophys. J. **88,** 2472 (2005).

[46] C.-Y. Huang, J. W. Klemke, Z. Getahun, W. F. DeGrado, and F. Gai, *Temperature-Dependent Helix-Coil Transition of an Alanine Based Peptide*, J. Am. Chem. Soc. **123,** 9235 (2001).

[47] B. Hess, C. Kutzner, D. van der Spoel, and E. Lindahl, *GROMACS 4: Algorithms for Highly Efficient, Load-Balanced, and Scalable Molecular Simulation*, J. Chem. Theory Comput. **4,** 435 (2008).

[48] G. Pérez-Hernández, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, *Identification of Slow Molecular Order Parameters for Markov Model Construction*, J. Chem. Phys. **139,** 015102 (2013).

[49] M. Senne, B. Trendelkamp-Schroer, A. S. J. S. Mey, C. Schütte, and F. Noé, *EMMA—A Software Package for Markov Model Building and Analysis*, J. Chem. Theory Comput. **8,** 2223 (2012).

[50] https://github.com/cmb-fu/pyfeat; or via the pypi python package index https://pypi.python.org/pypi/pyfeat.