RiteWeight: Randomized Iterative Trajectory Reweighting for Steady-State Distributions Without Discretization Error

Sagar Kania[1], David Aristoff[2], and Daniel M. Zuckerman[*1]

[1]Department of Biomedical Engineering, Oregon Health and Science University, Portland, OR, 97239, USA
[2]Department of Mathematics, Colorado State University, Fort Collins, CO, 80523, USA

January 12, 2024

1 Abstract

Molecular dynamics (MD) and enhanced sampling simulations have become fundamental tools for studying biomolecular events. A significant challenge in these simulations is ensuring that sampled configurations and transitions converge to the stationary distribution of interest, whether equilibrium or nonequilibrium. Lack of convergence constrains the estimation of mechanisms, free energy, and rates of complex molecular events. Here, we introduce the "Randomized Iterative Trajectory Reweighting" (RiteWeight) algorithm to estimate a stationary distribution from unconverged simulation data. This method iteratively reweights trajectories in a self-consistent way by solving for the stationary distribution using a discrete-state transition matrix, employing a new random clustering in each iteration. The iterative random clustering mitigates the phase-space discretization error inherent in existing trajectory reweighting techniques based on one-shot clustering and ultimately yields numerically unbiased, quasi-continuous configuration-space distributions and estimates of observables. We demonstrate the efficacy of RiteWeight using Trp-Cage synthetic MD trajectories.

2 Introduction

Despite numerous advances in enhanced equilibrium and path sampling methodologies in recent years [1, 2, 3, 4, 5, 6, 7], the comprehensive study of complex biomolecular systems remains a resource-intensive endeavor, often surpassing the computational capabilities accessible to many researchers. A major challenge is the estimation of the steady-state distribution of sampled states and transitions [8, 9, 10]. An accurate estimation is crucial because the stationary distribution reveals key thermodynamic and kinetic properties of the

---

*Corresponding author. Email: zuckermd@ohsu.edu

system, underpinning the calculation of free energies, identification of mechanistic pathways, rate constants and the "committor" reaction coordinate [11, 12, 13].
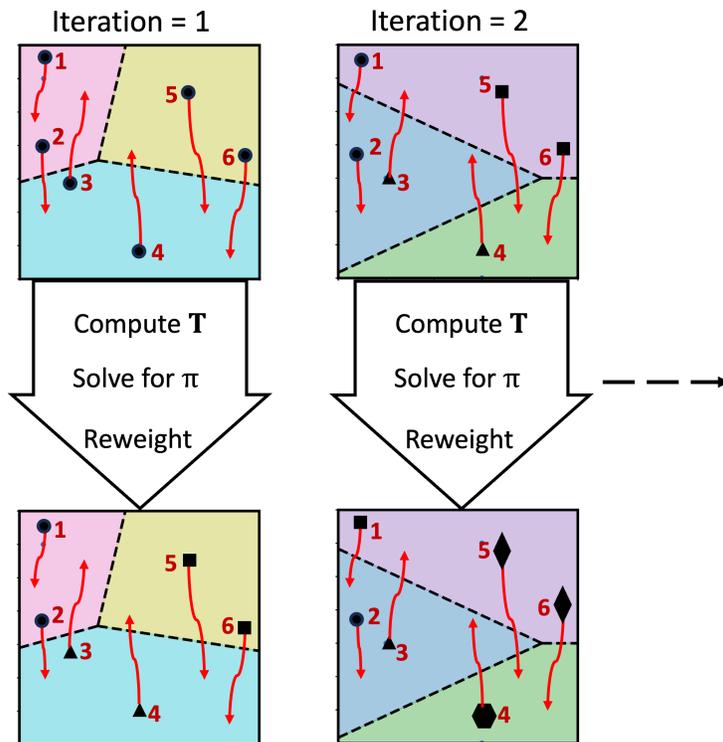
Figure 1: Schematic of the RiteWeight algorithm. In each iteration, a fixed set of trajectories (red arrows) is organized into clusters (colored regions). Based on the discrete clusters and current weights of the trajectories, the transition matrix **T** is computed and solved to yield the stationary distribution π for the given clusters. In the simplest RiteWeight algorithm, each trajectory is then assigned a new weight (schematized by filled black symbols) so that total cluster weights match π but the relative weights of trajectories within each cluster remain unchanged. Identical symbols indicate identical weight. In subsequent iterations, the process is repeated with new cluster boundaries, which enables changes in the relative weights of trajectories formerly in the same cluster, e.g., trajectories 1 and 2 in iteration 2. Ultimately all the initial weights are likely to be changed relative to one another. As described, the RiteWeight learning rate $r = 1$ has been assumed for reweighting via Eq (2)

One common approach to estimate the steady-state distribution from sampled transition data is the construction of a "Markov state model" (MSM), i.e., a discrete state transition matrix, which can easily be solved for the stationary condition [14, 15]. A further step can be taken, namely, the subsequent reweighting of the trajectories based on the matrix's stationary solution [12]. This method – referred to here as 'single-shot reweighting' – does not correct the weights of trajectories within the chosen discrete states, nor does it provide a set of weights consistent with a subsequently computed transition matrix[11, 16], potentially skewing the estimation of thermodynamic properties.

For unbiased estimation of the steady-state distribution and the accuracy of key observables using a standard MSM framework, trajectories within the states from which the transition matrix is calculated must themselves be locally in the stationary distribution [11, 12, 17]. This creates a "chicken and egg" problem for which an iterative solution has previously been proposed [11, 16] but which can require long trajectories. We note that other methods beyond traditional MSMs can be used to generate unbiased estimates of observables, e.g., [13]

As an approach for correcting standard MSMs using a single transition matrix, we introduce the "Randomized Iterative Trajectory Reweighting" (RiteWeight) algorithm. RiteWeight reweights trajectories, generated without biasing forces, into their correct stationary distribution when sufficient data is available, whether for equilibrium or nonequilibrium steady states. It naturally uses trajectories of any length, including a single-step, as it does not rely on dynamical relaxation. Thus it can employ data generated from standard molecular dynamics (MD) or path sampling approaches, so long as no biasing forces have been used.

Procedurally, as shown in Figure 1, the RiteWeight algorithm iteratively reweights trajectories using the stationary solution (π) of a discretized state transition matrix ($\mathbf{T}$), with each iteration featuring a new random clustering. The changing of cluster definitions is the key element, enabling RiteWeight to adjust phase-space weights with a precision beyond the resolution of the discrete state matrix: any given configuration finds itself in a new cluster in each iteration, implicitly providing information on its relative location within the clusters. The algorithm generates local stationarity within discrete states, consistent with the available trajectory data.

We validated the efficacy of the RiteWeight algorithm using synthetic molecular dynamics (SynMD) trajectories of the Trp-Cage miniprotein. The use of synMD enables comparison to exactly calculable reference distributions [18]. We compared the performance of RiteWeight and single-shot reweighting in estimating the stationary distribution from unconverged trajectory data, employing both finer and coarser discrete state transition matrices.

3   RiteWeight Algorithm

The RiteWeight algorithm estimates steady-state probability distributions by analyzing kinetic information extracted from numerous short MD trajectories or one long trajectory. This information consists of phase-space transition pairs, i.e., two time-sequential phase-space points, observed at a fixed lag time. In this preprint, we refer to these pairs as 'trajectory fragments'. Notably, the algorithm's ability to estimate stationary probabilities is not constrained by the choice of lag time [15, 14]; thus, lag time should be set to match the interval between available trajectory frames to maximize the use of the available data. Unlike single-shot reweighting, which requires within-cluster equilibrium at the discrete state resolution for unbiased results [11, 12, 17],

RiteWeight imposes no such constraints. Therefore, transition pair data can be collected right from the start of the trajectories, ensuring that no data is discarded and all available information is leveraged in the analysis.

To ensure rotational and translational invariance, transition pair data initially in x, y, z coordinates are transformed into a featurized space such as residue C_α pair-wise distances. Subsequent steps to estimate the steady-state probability of trajectory fragments are as follows:

1. Randomly select n configurations from trajectory fragments, thereby setting the number of states in the discrete-state transition matrix.

2. Use these n random configurations as Voronoi cluster centers and assign each trajectory fragment's initial and final configuration to these clusters, discretizing the transition pair data.

3. From the discretized transition pair data, compute the transition matrix $\mathbf{T}$ in a weighted fashion based on the *current* weights of the trajectory fragments as:

$$T_{IJ} = \frac{W_{IJ}}{W_I} \tag{1}$$

where W_{IJ} is the sum of all trajectory fragments' weights transitioning from I^{th} to J^{th} cluster and $W_I = \sum_J (W_{IJ})$.

4. Calculate the stationary probabilities π_I for the n discrete states as the steady-state solution of matrix $\mathbf{T}$, i.e., $\pi \mathbf{T} = \pi$.

5. Define a new weight for each trajectory fragment i based on the current iteration's definition of clusters I (from Step 2), using a linear combination of the fragment's current weight and the stationary solution π from Step 4. Specifically, for a fragment i in cluster I the new weight is given as:

$$w_{i \in I}^{\mathrm{new}} = (1 - r)\, w_i + r\, \frac{\pi_I}{W_I}\, w_i \tag{2}$$

$$W_I = \sum_{i \in I} w_i \tag{3}$$

where w_i^{new} is the new weight of the i^{th} trajectory fragment and w_i is the trajectory fragment weight before update. Here, r represents the learning rate, a hyper-parameter within the interval (0,1]. The notation "$i \in I$" refers to whether the initial configuration of trajectory fragment i is in the current coarse cluster I. Thus, W_I in (3) is defined to be the sum of the current iteration's weights (before updates) in the currently defined cluster I, and π_I is the current iteration's estimate of the stationary probability for the current cluster I.

6. Repeat steps 1-5 until the user-defined convergence criteria are met.

Several points are noteworthy. Random clusters are defined via steps 1 and 2, and thus a new clustering is performed every iteration. In step 5, the parameter r functions as the learning rate and dictates the proportion of the I^{th} cluster's stationary probability π_I to be considered when reweighting the trajectory fragments during the k^{th} iteration of RiteWeight. This parameter becomes particularly significant when dealing with noisy data, with a smaller r value recommended to mitigate the impact of noise. In contrast, setting $r = 1$ means that a trajectory is assigned a weight that is simply a fraction of the currently estimated stationary probability for the enclosing cluster, with the fraction determined by the previous iteration's weights.

For the initial assignment of weights to the trajectory fragments, two main approaches can be adopted. One option is to assign uniform weights to all fragments, which essentially

serves as an uninformative prior. Alternatively, if there is prior knowledge available about the stationary distribution of the system, this information can be used to assign the initial weights in a more informed manner.

4   Test System: Trp-Cage Synthetic Molecular Dynamics (SynMD) Trajectory

To demonstrate the efficacy of the RiteWeight method, we apply it to reweight unconverged synthetic molecular dynamics (SynMD) trajectories for estimating the stationary probability distribution [18]. Synthetic MD trajectories are preferred in this study rather than low-dimensional toy models for their computational efficiency and their ability to capture the complex dynamics inherent in real systems which are absent from oversimplified models like the 1-D double-well potential. These SynMD trajectories are generated using a Markov Chain model based on a transition matrix derived from a Markov State Model (MSM) analysis of an ultra-long MD trajectory [18]. The MSM analysis employed here uses a finer discretization of the state space than typical MSM methods, with each state mapped to specific atomistic configurations, thus permitting calculations analogous to those done with standard MD trajectories [18].

Our test system utilizes the Trp-Cage protein SynMD model, generated from MD trajectories spanning 208 microseconds. The creation of this model, which comprises 10,500 discrete microstates indexed from 0 to 10,499, is detailed in the Ref [18]. The MSM transition matrix for this model is constructed with a lag time of 1ns.

From this Trp-Cage SynMD model, numerous short SynMD trajectories were generated, yielding a distribution significantly different from equilibrium. The initial states for the trajectories were selected by partitioning the SynMD discrete phase space into 500 bins along the state indices, from which 20 microstates were randomly chosen—either unique or repetitive—from each bin. Each selected microstate was used to initiate a short trajectory of 5τ length, where τ represents the lag time of the SynMD MSM. Transition pairs with lag time τ were then recorded from these short trajectories and assigned uniform weights. The initial probability distribution function for the Trp-Cage SynMD states, based on these uniform weights, is presented in Figure 2. These uniform-weighted trajectory fragments were processed with the RiteWeight algorithm to assess the algorithm's capability to estimate the equilibrium probabilities of the SynMD states. The exact equilibrium probability distribution function for the Trp-Cage SynMD model's states, determined as the stationary solution of the SynMD transition matrix, is also depicted in Figure 2.

For clustering and solving stationarity during the RiteWeight process, atomistic configurations were featurized based on the distances between every C_α pair, with further dimensionality reduction performed using TICA to enhance the computational efficiency of the RiteWeight algorithm. We implemented RiteWeight with both 1000 and 10 clusters, incorporating random clustering in each iteration as outlined in the methodology and the learning rate r was set to 1 because of the large amount of data available for this proof-of-principle study.
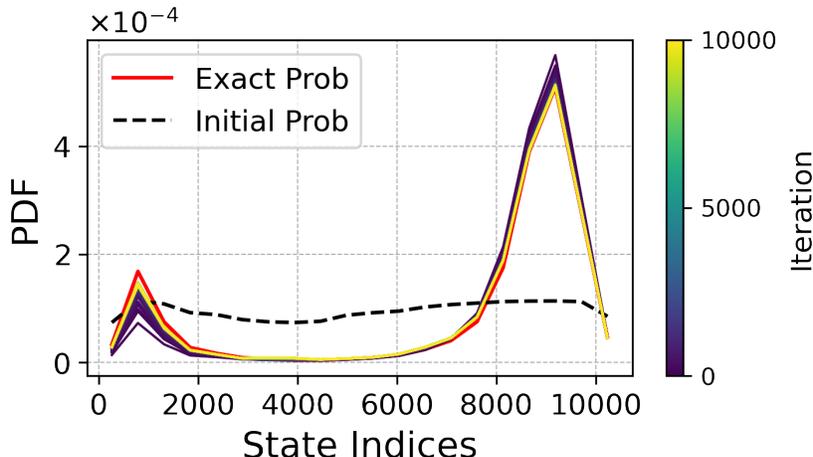
Figure 2: RiteWeight convergence to equilibrium in Trp-cage system. Plotted is the evolution of the Trp-Cage SynMD states' probability distribution function (PDF) across RiteWeight iterations (color bar at right). The black dashed curve represents the initial PDF based on uniformly weighted trajectory fragments, while the solid red curve indicates the exact PDF. Intermediate curves in various colors show the computed PDFs at every 100th RiteWeight iteration, illustrating the convergence toward the exact PDF. The microstate index shown horizontally is highly correlated with RMSD to the folded structure. The data shown used 1,000 clusters in the RiteWeight procedure.

5   Results and Discussion

The RiteWeight algorithm was used to estimate equilibrium probabilities in the Trp-Cage system. The algorithm was run until the probability distribution function (pdf) for the SynMD states converged. We visualized the pdf (Figure 2) by categorizing the trajectory fragments according to their initial state's indices into 20 uniform bins, then generating the pdf with these bins and the weights of the fragments.

We next examined whether the algorithm yields the correct stationary distribution when a very coarse (highly non-Markovian) clustering is used. In particular, we executed RiteWeight using only 10 clusters while keeping the input trajectory data and featurization identical to the 1000 clusters setup. For both cluster resolutions the final converged PDF, estimated by averaging PDFs over the last 1000 converged iterations of the RiteWeight algorithm, accurately matches the exact equilibrium probabilities, as shown in Figure 3.

In contrast to the converged RiteWeight distributions, the single-shot reweighting estimates – i.e., after the first iteration of the RiteWeight algorithm – show significant discrepancies due to discretization error. As demonstrated in Figure 3(b), the first-iteration distributions estimated with 10 clusters significantly diverges from the pdf estimated with 1000 clusters. In the 10-cluster scenario, single-shot reweighting is markedly different from the exact equilibrium probabilities. Even with a more detailed clustering of 1000 clusters, the single-shot reweighting does not accurately predict the true stationary probabilities of the Trp-Cage SynMD microstates. This indicates that an assumption of local equilibrium is not justified, even at the higher resolution of 1000 clusters.

We also monitored the "relaxation" over iterations of the RiteWeight pdfs. The KL-divergence between successive pdfs, at intervals of 100 iterations, is computed to check the convergence and shown in Figure 4. The convergence of the estimated pdf when utilizing 1000 clusters is much more rapid than with 10 clusters. Figure 4 indicates that For 10 clusters, the RiteWeight algorithm requires approximately two orders of magnitude more iterations to achieve convergence compared to the 1000 clusters. Nonetheless, both converge to the correct distribution as indicated in Figure 3(a).
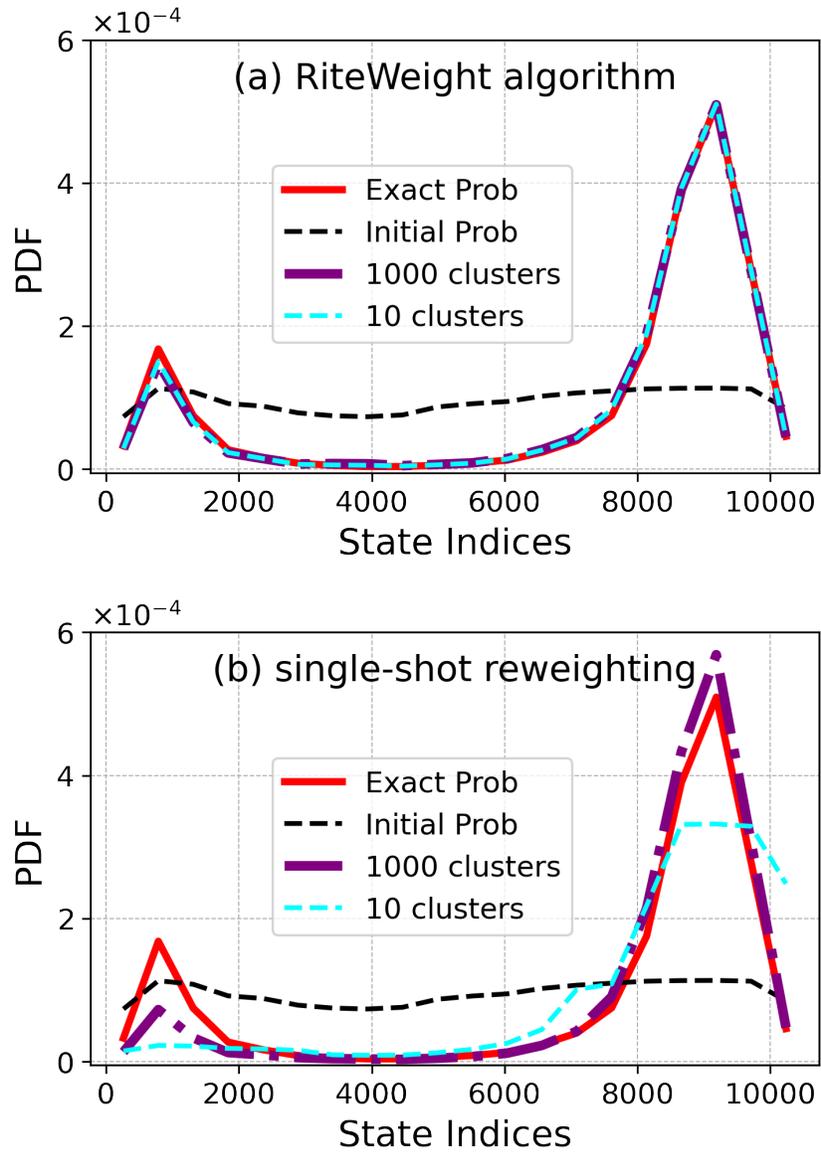
Figure 3: Comparison of distributions based on different numbers of clusters. Plotted are estimated stationary distributions for Trp cage using 10 clusters (cyan) and 1,000 clusters (magenta) for (a) the RiteWeight algorithm and (b) single-shot reweighting. For reference, the initial PDF (black dashed) and exact PDF (red) are also shown.
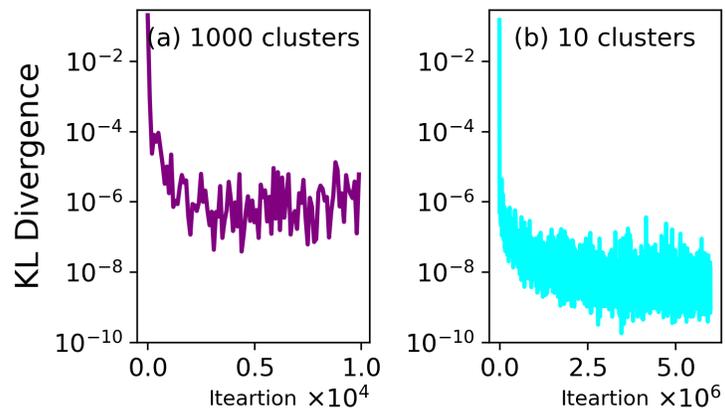
Figure 4: Convergence of the RiteWeight algorithm. The symmetric Kullback-Leibler (KL) divergence – comparing the estimated stationary distribution in the current iteration vs. the distribution of 100 iterations prior – is plotted against the number of iterations for two levels of clustering resolution. (a) The KL divergence for a finer resolution: 1000 clusters. (b) The KL divergence for a coarser resolution: 10 clusters.

6   Conclusions

We have described the RiteWeight iterative algorithm for reweighting trajectory segments based on stationary solutions of a coarse-grained transition matrix that uses randomized clusters newly defined at each iteration. Application of the method to the Trp-Cage SynMD model demonstrates its convergence to the stationary distribution. Further, the algorithm's performance in estimating the steady-state distribution is shown to be similar when either extremely coarse (10 clusters) or fine-grained (1,000 clusters) discretizations are used to define the transition matrix, suggesting mitigation of discretization error. In contrast, previously proposed single-shot and related reweighting methods are constrained by discrete state resolution and require locally equilibrated trajectory fragments [11, 12, 17].

RiteWeight corrects trajectory weights at a quasi-continuous configuration-space resolution, enabling estimation of steady state from trajectories of lengths as short as a single lag time. Thus, it can be considered a kind of "super resolution" Markov modeling.

The use of short lag times enables the characterization of shorter-timescale mechanistic features that get averaged out at longer lag times [19, 20].

Acknowledgements

References

[1] Peter G Bolhuis, David Chandler, Christoph Dellago, and Phillip L Geissler. Transition path sampling: Throwing ropes over rough mountain passes, in the dark. *Annual review of physical chemistry*, 53:291, 2002.

[2] Ron Elber. Milestoning: An efficient approach for atomically detailed simulations of kinetics in biophysics. *Annual review of biophysics*, 49:69, 2020.

[3] Frank Noé, Christof Schütte, Eric Vanden-Eijnden, Lothar Reich, and Thomas R Weikl. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proceedings of the National Academy of Sciences*, 106:19011, 2009.

[4] Bernd Ensing, Marco De Vivo, Zhiwei Liu, Preston Moore, and Michael L Klein. Metadynamics as a tool for exploring free energy landscapes of chemical reactions. *Accounts of chemical research*, 39:73, 2006.

[5] Diego E. Kleiman, Hassan Nadeem, and Diwakar Shukla. Adaptive sampling methods for molecular dynamics in the era of machine learning. *The Journal of Physical Chemistry B*, 127:10669, 2023.

[6] Daniel M Zuckerman and Lillian T Chong. Weighted ensemble simulation: review of methodology, applications, and software. *Annual review of biophysics*, 46:43, 2017.

[7] Jérôme Hénin, Tony Lelièvre, Michael R Shirts, Omar Valsson, and Lucie Delemotte. Enhanced sampling methods for molecular dynamics simulations. *arXiv preprint, arXiv:2202.04164*, 2022.

[8] Daniel M Zuckerman and Edward Lyman. A second look at canonical sampling of biomolecules using replica exchange simulation. *Journal of chemical theory and computation*, 2:1200, 2006.

[9] Wilfred F Van Gunsteren, Dirk Bakowies, Riccardo Baron, Indira Chandrasekhar, Markus Christen, Xavier Daura, Peter Gee, Daan P Geerke, Alice Glättli, Philippe H Hünenberger, et al. Biomolecular modeling: goals, problems, perspectives. *Angewandte Chemie International Edition*, 45:4064, 2006.

[10] PG Bolhuis and C Dellago. Practical and conceptual path sampling issues. *The European Physical Journal Special Topics*, 224:2409, 2015.

[11] John D Russo, Jeremy Copperman, David Aristoff, Gideon Simpson, and Daniel M Zuckerman. Unbiased estimation of equilibrium, rates, and committors from markov state model analysis. *arXiv preprint, arXiv:2105.13402*, 2021.

[12] Hongbin Wan and Vincent A Voelz. Adaptive markov state model estimation using short reseeding trajectories. *The Journal of chemical physics*, 152:24103, 2020.

[13] Feliks Nüske, Hao Wu, Jan-Hendrik Prinz, Christoph Wehmeyer, Cecilia Clementi, and Frank Noé. Markov state models from short non-equilibrium simulations—analysis and correction of estimation bias. *The Journal of Chemical Physics*, 146:94104, 2017.

[14] Brooke E Husic and Vijay S Pande. Markov state models: From an art to a science. *Journal of the American Chemical Society*, 140:2386, 2018.

[15] Jan-Hendrik Prinz, Hao Wu, Marco Sarich, Bettina Keller, Martin Senne, Martin Held, John D Chodera, Christof Schütte, and Frank Noé. Markov models of molecular kinetics: Generation and validation. *The Journal of chemical physics*, 134:174105, 2011.

[16] John D Russo, Jeremy Copperman, and Daniel M Zuckerman. Iterative trajectory reweighting for estimation of equilibrium and non-equilibrium observables. *arXiv preprint, arXiv:2006.09451*, 2020.

[17] Jeremy Copperman and Daniel M Zuckerman. Accelerated estimation of long-timescale kinetics from weighted ensemble simulation via non-markovian "microbin" analysis. *Journal of chemical theory and computation*, 16:6763, 2020.

[18] John D Russo and Daniel M Zuckerman. Simple synthetic molecular dynamics for efficient trajectory generation. *arXiv preprint, arXiv:2204.04343*, 2022.

[19] Ernesto Suárez, Rafal P Wiewiora, Chris Wehmeyer, Frank Noé, John D Chodera, and Daniel M Zuckerman. What markov state models can and cannot do: Correlation versus path-based observables in protein-folding models. *Journal of chemical theory and computation*, 17:3119, 2021.

[20] Ernesto Suárez and Daniel M Zuckerman. Pathway histogram analysis of trajectories: A general strategy for quantification of molecular mechanisms. *arXiv preprint, arXiv:1810.10514*, 2018.