Published in final edited form as:

J Chem Theory Comput. 2018 November 13; 14(11): 5459–5475. doi:10.1021/acs.jctc.8b00500.

Choice of adaptive sampling strategy impacts state discovery, transition probabilities, and the apparent mechanism of conformational changes

Maxwell I. Zimmerman¹, Justin R. Porter¹, Xianqiang Sun¹, Roseane R. Silva¹, and Gregory R. Bowman^{1,2,3,*}

¹Department of Biochemistry and Molecular Biophysics, Washington University School of Medicine, St. Louis, Missouri 63110, United States

²Department of Biomedical Engineering, Washington University in St. Louis, St. Louis, Missouri, 63110, United States

³Center for Biological Systems Engineering, Washington University in St. Louis, St. Louis, Missouri, 63110, United States

Abstract

Interest in atomically-detailed simulations has grown significantly with recent advances in computational hardware and Markov state modeling (MSM) methods, yet outstanding questions remain that hinder their widespread adoption. Namely, how do alternative sampling strategies explore conformational space and how might this influence predictions generated from the data? Here, we seek to answer these questions for four commonly used sampling methods: 1) a single long simulation, 2) many short simulations run in parallel, 3) adaptive sampling, and 4) our recently developed goal-oriented sampling algorithm, FAST. We first develop a theoretical framework for analytically calculating the probability of discovering select states on simple landscapes, where we uncover the drastic effects of varying the number and length of simulations. We then use kinetic Monte Carlo simulations on a variety of physically inspired landscapes to characterize the probability of discovering particular states and transition pathways for each of the four methods. Consistently, we find that FAST simulations discover each target state with the highest probability, while traversing realistic pathways. Furthermore, we uncover the potential pathology that short parallel simulations sometimes predict an incorrect transition pathway by crossing large energy barriers that long simulations would typically circumnavigate. We refer to this pathology as “pathway tunneling”. To protect against this phenomenon when using adaptive-sampling and FAST simulations, we introduce the FAST-string method. This method enhances sampling along the highest-flux transition paths to refine an MSMs transition probabilities and discriminate between competing pathways. Additionally, we compare the performance of a variety

*Corresponding Author: bowman@biochem.wustl.edu.

Supporting Information

This information is available free of charge via the Internet at <http://pubs.acs.org>.

Contents include an example of calculating the discover probabilities on a simple landscape, figures of discover probabilities, pathway probabilities, and relative entropies for various sampling methods on generated landscapes, predicted state populations from different estimators, λ -repressor folding pathways, and a table quantitatively detailing the performance of each sampling methods' performance on the various generated landscapes.

of MSM estimators in describing accurate thermodynamics and kinetics. For adaptive sampling, we recommend simply normalizing the transition counts out of each state after adding small pseudo-counts to avoid creating sources or sinks. Lastly, we evaluate whether our insights from simple landscapes hold for all-atom molecular dynamics simulations of the folding of the λ -repressor protein. Remarkably, we find that FAST-contacts predicts the same folding pathway as a set of long simulations but with orders of magnitude less simulation time.

Introduction

The use of all-atom molecular dynamics (MD) simulations for long time-scale phenomena are often thwarted by insufficient computational resources. Many interesting biological processes occur on the millisecond to second timescale, where a single simulation may take longer than a lifetime to gather. Notable attempts to alleviate hardware limitations are the purpose-built ANTON supercomputers.^{1,2} These supercomputers are an engineering feat, yet are still limited by sampling and not accessible to many researchers. Since increasing commodity hardware performance by many orders of magnitude is not likely in the immediate future, the observation of interesting biological phenomena requires the use of clever sampling techniques.

A common technique to increase the observation of long time-scale phenomena is to alter the underlying energy landscape. These methods aim to guide a simulation towards some end goal or toward the exploration of a set of order parameters. Some examples include, G models,^{3,4} replica-exchange,⁵⁻⁷ steered MD,^{8,9} accelerated MD,^{10,11} meta-dynamics,^{12,13} among others.¹⁴⁻¹⁶ Unfortunately, these methods do not capture proper kinetic information, and can traverse unrealistically high energy barriers. Here, we are particularly interested in sampling methods that access long time-scale phenomena without perturbing the underlying energy landscape, such that both thermodynamic and kinetic properties can be inferred.

As an alternative to a single long simulation, many independent simulations can be run in parallel. Combined, these parallel simulations tractably capture time-scales longer than any single simulation. To illustrate: if we assume that the transition between conformational states A and B follows a Poisson process, the probability of observing a transition to state B is dependent only on the aggregate simulation time from A, not the length of each simulation.¹⁷ Put another way, the probability of traversing a single energy barrier is based on the number of attempts to cross that barrier, regardless of whether they are in parallel or successive. Thus, parallel simulations may offer a significant enhancement in the observation of rare events, since it is usually easier to add more computational resources than to make them faster. This is the strategy of Folding@home, which takes advantage of around 100,000 personal computers, whose resources are donated for massively distributed MD simulations.¹⁸ Additionally, many parallel simulations may provide better estimates of transition rates for this single barrier, since there are more statistics on the outward transitions.

For large sets of independent simulations that are in local equilibrium (i.e. they sample from the underlying energy distribution), we can reconstruct both the proper thermodynamics and kinetics with the use of Markov state models (MSMs).¹⁹ An MSM is a network model that

describes a protein's energy landscape in terms of a set of structural states the protein tends to adopt and the probabilities of transitioning between neighboring states in a fixed time interval. The utility of an MSM depends on accurately estimating the conditional transition probabilities between conformational states, without requiring that any individual simulation achieve global equilibration. As a consequence, the number of times different states are sampled does not need to be Boltzmann distributed for an accurate description of their populations at equilibrium, provided that estimates of transition probabilities are accurate. MSMs have recently succeeded in guiding the design of new proteins^{20,21} and allosteric modulators, among many other applications.^{19,23–31}

MSMs' ability to integrate information from many parallel simulations whose starting states are not necessarily Boltzmann distributed opens the possibility of performing adaptive sampling. First developed for refining MSMs by identifying conformational states that contribute the most to statistical uncertainty,³² adaptive sampling schemes typically have the following steps: 1) run simulations, 2) build an MSM from simulations, 3) rank each state by some metric, 4) start new simulations from the highest ranked states, and 5) repeat steps 2-4 for some number of rounds or until a convergence criterion is met. The main difference between adaptive sampling algorithms is in the metric for ranking and selecting states for future sampling.^{32–41} Recently, we have developed the goal-oriented sampling algorithm, Fluctuation Amplification of Specific Traits (FAST), that ranks states on some structural metric in addition to a statistical metric.^{42,43} We have demonstrated that this method increases the rate of state exploration by at least an order of magnitude, and additionally, can capture thermodynamic and kinetic information that agrees with a multitude of experiments.^{21,44}

Each of the equilibrium-based sampling methods mentioned above (long-, parallel-, adaptive-, and FAST-simulations) should converge on identical MSMs, provided with near infinite sampling. Unfortunately, for most systems of interest, simulations are not able to reach global equilibrium, and are usually significantly under-sampled. It should be noted that FAST, and other adaptive sampling algorithms, do not *increase* the amount of sampling, but rather focus sampling efforts to specific regions of conformational space to make the most of limited computational resources. With that, the functional differences between methods are simply the rates at which specific sections of conformational-space are explored. However, it is not completely understood how each of these methods influences the probability of discovering states, nor how this influences the mechanism of conformational changes that is observed, especially when conformational sampling is far from global equilibrium.

In this work, we seek to assess the relative performance of different sampling strategies. We develop an analytical expression for the probability of discovering a conformational state for very simple landscapes. We find that state discovery is dependent on the number and length of simulations, in addition to the shape of the energy landscape. We then examine the performance of the four equilibrium-based sampling methods above in finding the highest-flux pathway between two states, for a variety of energy landscapes. These results are very informative for tuning the many hyperparameters in adaptive sampling, and even identify pitfalls that should be avoided. Lastly, we demonstrate that insights from our simple

landscapes are consistent with observations using all-atom MD simulations, by generating folding trajectories of a fast-folding version of the λ -repressor.

Theory

To understand how the probability of discovering a state on a particular landscape is dependent on sampling, we develop a mathematical formalism for describing the probability that a set of simulations will discover a particular conformational state. First, we consider sampling to occur on a discretized energy landscape with N conformational states, where the state index is represented as n_i , $i = 1, \dots, N$. Transitions between states are described by the $N \times N$ -transition probability matrix, T_{ij} , which is the probability of transitioning from state n_i to n_j at a specified lag-time, τ . A simulation on this landscape of K -steps is denoted with the symbol $\mathbf{X}$, where the conformation at the k -th time step is X_k , $k = 1, \dots, K$. For a dataset with M simulations, we denote the m -th simulation as $\mathbf{X}_m$, $m = 1, \dots, M$. For multiple simulations of various lengths (different number of time steps), we choose $\mathbf{K}$ to represent a vector of lengths, where K_m , $m = 1, \dots, M$, is the length of the m -th simulation.

Towards our goal of describing the probability of discovering a particular conformational state on an energy landscape given sampling parameters, we introduce the $N \times N$ -matrix, $D_{ij}^{\mathbf{K}, \mathbf{M}}$, which indicates if state n_j is ever discovered within the trajectories $\mathbf{X}_M$, started from state n_i with lengths described by $\mathbf{K}$. For example, if state n_j is a state within the trajectories, $D_{ij}^{\mathbf{K}, \mathbf{M}}$ is 1, otherwise it is 0. This can be represented with,

$$D_{ij}^{K, M} = \begin{cases} 1 & \text{if } n_j \in \mathbf{X}_M \\ 0 & \text{if } n_j \notin \mathbf{X}_M \end{cases} \quad [1]$$

While this can be determined for a set of trajectories, we wish to know the probability of having observed a state, *a priori*, or $P(D_{ij}^{\mathbf{K}, \mathbf{M}} = 1)$. This is the probability of discovering state n_j given the sampling parameters $\mathbf{K}$. For short hand, we call these probabilities the “discover probabilities”.

Before providing an expression for $P(D_{ij}^{\mathbf{K}, \mathbf{M}} = 1)$, we must first introduce another $N \times N$ -matrix, v_{ij}^k , which indicates if the conformation at the k -th step of a single trajectory, $\mathbf{X}$, belongs to the state n_j when started from state n_i .⁴⁵

$$v_{ij}^k = \begin{cases} 1 & \text{if } X_k = n_j \\ 0 & \text{if } X_k \neq n_j \end{cases} \quad [2]$$

Additionally, we are interested in the probability of this event occurring, denoted as $P(v_{ij}^k = 1)$. Since only one conformation at the k -th step can be observed, each row of $P(v_{ij}^k = 1)$ is a normalized probability vector indicating the state index at time k . For the

trivial case of the 0th-step (i.e. before a simulation is generated), the probability of being in the starting state is 1, and everywhere else, 0:

$$P(v_{ij}^k = 0 = 1) = I$$

where I is the identity matrix. Since $P(v_{ij}^k = 1)$ is a list of probability vectors, we can propagate the probabilities a time step (the lag-time, τ) using the transition probability matrix, T .

$$P(v_{ij}^k = 1) = \begin{cases} I & \text{if } k = 0 \\ P(v_{ij}^{k-1} = 1)T & \text{if } k > 0 \end{cases} \quad [3]$$

This expression is useful for determining $P(D_{ij}^{\mathbf{K}, \mathbf{M}} = 1)$, since the probability of ever visiting a state is the complement of not visiting it at each time step. For example, the probability of discovering state n_j after one step is the complement of not discovering it before and after one step:

$$P\left(D_{ij}^{\mathbf{K} = \{1\}, M = 1} = 1\right) = 1 - \left(1 - P(v_{ij}^0 = 1)_{ij}\right) * \left(1 - P(v_{ij}^1 = 1)_{ij}\right) = 1 - (1 - I_{ij}) * (1 - T_{ij}) = \begin{cases} 1 & \text{if } i = j \\ T_{ij} & \text{if } i \neq j \end{cases}$$

This reasoning holds true for a single step in a simulation, although does not for more than one step. What is required is an expression for the probability of being in a state at time step, k , conditional on not having discovered state n_j for all of the previous steps. We represent this expression as, $P(v_{ij}^k = 1 \mid \{v_{ij}^{k'} = 0 \forall k' < k\})$, which can be evaluated with the following:

$$P\left(v_{ij}^k = 1 \mid \{v_{ij}^{k'} = 0 \forall k' < k\}\right) = \begin{cases} I & \text{if } k = 0 \\ P(v_{ij}^{k-1} = 1 \mid \{v_{ij}^{k'} = 0 \forall k' \leq (k-1)\})T & \text{if } k > 0 \end{cases} \quad [4]$$

For each step in the recursive calculation, the j^{th} column of $P(v_{ij}^{k-1} = 1)$ is set to 0, and each row is then normalized to unity. This is described in more detail in the supporting information.

Using this definition, we can extend our expression of the discover probabilities to include an arbitrary number of steps, K . In a single simulation, we can see that the probability of discovering a state is:

$$P\left(D_{ij}^{\mathbf{K} = \{K\}, M = 1} = 1\right) = 1 - \Pi_{k=0}^K \left(1 - P(v_{ij}^k = 1 \mid \{v_{ij}^{k'} = 0 \forall k' < k\})_{ij}\right) \quad [5]$$

Since the probability of discovering a state within a simulation is independent of the probability in another simulation, the discover probabilities for multiple simulations is the complement of not discovering a state in any of the individual simulations. For example, in the case of two simulations with lengths K_0 and K_1 ,

$$P\left(D_{ij}^{\mathbf{K}=\{K_0, K_1\}, M=2}\right) = 1 - \left(1 - P\left(D_{ij}^{\mathbf{K}=\{K_0\}, M=1}\right)\right) * \left(1 - P\left(D_{ij}^{\mathbf{K}=\{K_1\}, M=1}\right)\right)$$

This can be generalized to an arbitrary number of simulations with arbitrary lengths:

$$P(D_{ij}^{\mathbf{K}, \mathbf{M}} = 1) = 1 - \prod_{m=1}^M \left[\prod_{k=0}^{K_m} \left(1 - P(v_{ij}^k = 1 | \{v_{ij}^{k'} = 0 \ \forall \ k' < k\}) \right) \right] \quad [6]$$

This gives us our final expression for state discovery as a function of our equilibrium-sampling parameters.

Results and Discussion

There are different advantages to running many short or a few long simulations

From equation 6, it is clear that the probability of discovering a state is influenced by four parameters: 1) the number of trajectories, 2) the lengths of the trajectories, 3) the starting state, and 4) the shape of the landscape being sampled. Strikingly, this implies that the probability of discovering a state can be drastically distinct between a single long simulation and many short simulations, though this is only true for finite sampling since

$P(D_{ij}^{\mathbf{K}, \mathbf{M}} = 1) \rightarrow 1$ as either, $\mathbf{M} \rightarrow \infty$, or $K_m \rightarrow \infty$. It may seem tempting to seek the global optimum sampling parameters, however, sampling is strongly dependent on the specifics of the landscape itself. Additionally, different goals may necessitate different sampling strategies, i.e. is the goal to discover as many states as possible, or to discover a pathway between a particular set of states? From this, our goal is to characterize different sampling strategies for a variety of different landscapes to gain insight into their appropriate uses.

As a first test, we constructed a simple landscape where four states are connected in a linear arrangement (Figure 1A). Here, each state can transition to either a neighbor or itself, with differing probabilities. We imagine that these states represent a conformational landscape where each successive state is progress along some order parameter. Starting from state 0, the first state in the chain, we calculate the probability of discovering the other states from either a single trajectory or many parallel trajectories with an equivalent aggregate amount of simulation. Figure 1C depicts the probability of discovering states 2 (blue curves) and 3 (red curves) from a single trajectory at various time-steps (solid lines), or some number of parallel trajectories with 4 time-steps each (dashed lines). We see that the long simulations have a higher probability of reaching states 3 and 4 than do parallel simulations. For this shaped landscape, the discrepancy between long and parallel simulations widens with the number of states. This makes intuitive sense from equation 6, because we see that the

probability of a simulation making 2 successive transitions is different than one of two simulations making 2 successive transitions.

We should note that the fully connected landscape in Figure 1B also displays this property, indicating that it is not an artifact of the way we have drawn the landscape. Here, the probability of transitioning between state 0 to 2 is very low, making the more probable route go through the transition state, 1. This leaves parallel simulations at a disadvantage of having to take the longer route to observe the transition, making this observation less probable. Interestingly, this also indicates that it is possible to consistently stumble upon an incorrect conclusion for the transition pathway; a trivial example being that many 1-step simulations started from state 0 would incorrectly predict the pathway as going directly from state 0 to 2. It is an important point that this result arises from the probability of discovering certain states, and their transitions, but not from the estimates of each states conditional transition probabilities, which should remain preserved across sampling methods. Therefore, in addition to understanding how sampling affects state discovery, we are interested in how the state discovery influences the predicted mechanism of conformational changes (e.g. the highest probability transition pathways between two sets of states). We investigate this idea in more detail in later sections.

So far, linear and fully connected landscapes might lead one to believe that long simulations are always advantageous in state discovery, but this is not the case when landscapes have entropic barriers. For many realistic systems, it is likely that a particular conformational state has many other states that it can transition to. To capture this transitional entropy, we generated the star-shaped landscape depicted in Figure 2A. This landscape has a central state and 5 arms, which is reminiscent of a “kinetic hub” where unfolded/high-energy states typically pass through the folded state to reach other unfolded/high-energy states.⁴⁶ Parallel simulations have a significantly higher probability to discover any of the states on this landscape, compared with equal aggregate time of the long simulation. We reason that the long simulations are penalized by having to backtrack to explore each of the arms, whereas the parallel simulations have a high probability of sampling multiple arms simultaneously. This effect becomes more drastic as the dimensionality of the state-space increases. Furthermore, this landscape provides a nice example that the optimal sampling scheme is strongly dependent on the shape of the landscape.

These simple landscapes provide valuable insight into how long or parallel simulations affect state discovery, setting a baseline for characterizing more complicated sampling schemes, such as adaptive sampling. Towards this goal, we generated a series of larger landscapes, which emulate common challenges in the sampling of proteins. To aid in human intuition, these landscapes are two-dimensional energy surfaces projected onto a grid, where each point on the grid represents a conformational state with a single potential energy. Each state can have up to four connected neighbors, with transitions governed by the Metropolis criterion. In the next few sections, we make use of kinetic Monte Carlo simulations on these landscapes using four different sampling methods: 1) a single long simulation (referred to as “long”), 2) many short simulations (referred to as “parallel”), 3) counts-based adaptive sampling (referred to as “counts”), and 4) our goal-oriented FAST algorithm (referred to as “FAST”). Although there are many adaptive sampling algorithms, we chose to use counts

because it has been shown to be the best at indiscriminately discovering new states.^{34,42} The specifics of sampling are described in greater detail in the methods section. Furthermore, we aim to characterize each method based on three different criteria: 1) ability to discover a target state, 2) ability to predict realistic transition pathways, and 3) ability to estimate accurate transition probabilities.

FAST is most likely to discover the target state

The first landscape that we generated was inspired by the challenge of using MD simulations to find the native state of a cooperatively folding protein. Two common tasks include: 1) to determine the native conformational state given an amino-acid sequence, also known as a structure prediction problem,^{47–49} and 2) explore the preferred pathway(s) from an unfolded state to the native state.^{50,51} We chose to start with one of the simplest possible models, a minimally frustrated folding-funnel (Figure 3).^{52,53} Here, there is a reasonably smooth energetic gradient from a high-energy starting-state to the low-energy target state. The solid colored lines represent the three highest-flux pathways from the start to the target.

To characterize state-discovery on this landscape, we performed 5,000 independent trials of each sampling method, with equivalent aggregate simulation times, as is described in the methods. We then calculate the probability of discovering a given state (which we refer to as the discover probabilities) for the four methods, by averaging the results of equation 1 for each trial. We note that we terminate the algorithm after reaching the target state, since we are mainly concerned with the initial pathway to the target; including excessive sampling after reaching the target convolutes the results with what happens afterwards. Additionally, trimming the data after discovering the end state does not affect the discover probabilities of the end state itself.

If the goal is to simply reach the end-state, FAST does so with the highest probability. The discover probabilities of the target state are $1.0 \pm 7 \times 10^{-4}$, 0.94, $0.62 \pm 7 \times 10^{-3}$, and 2.2×10^{-5} for FAST, long, counts, and parallel simulations respectively (this value for long and parallel simulations come from equation 6). It is not a surprise that FAST is the best at reaching the end state, since it is the only method tested that uses knowledge of the end state in its sampling and we have previously shown FAST's ability to reach a target state with orders of magnitude less simulation.⁴² Of greater interest here is the difference between the observation of states along the way to the target.

Towards this goal, we plot the discover probabilities for each method in Figure 4, which reveals distinct patterns for each sampling method. We find it extremely beneficial to view the discover probabilities for each state in this manner, since it provides intuition for the ways that each method explores the landscape before reaching the target. Analysis of the long simulations indicates that they have a propensity to sample around the native-well before reaching the target state. The 25 states closest to the target have over a 0.9 probability of being discovered first. Conversely, parallel simulations rarely venture near the target, but thoroughly explore the landscape around the starting state. Strikingly, this suggests that parallel simulations would require orders of magnitude more aggregate simulation time than the long simulation to reliably observe a transition to the target. In fact, this is what we observe for MD simulations of the λ -repressor in a later section.

Unlike the other sampling strategies, counts-based adaptive sampling has an elevated propensity to explore the high-energy edges of the funneled landscape. Compared to the long simulations, counts has almost twice the probability of discovering the states furthest from the start and the target, yet, nearly half the probability of discovering the target itself. This is because counts indiscriminately discovers new states, particularly in high-energy neighborhoods where low count states are prevalent. This property enables counts-based sampling to lead in state discovery, with an average of 183.3 ± 12.3 states discovered, in comparison to 168.5 ± 12.3 , 144.2 ± 24.0 , and 72.7 ± 10.1 for FAST, long, and parallel simulations respectively. Interestingly, counts-based sampling's propensity to climb energy barriers actually hinders its ability to follow a simple gradient to the global minimum. Therefore, counts-based simulations may actually be a poor choice for many applications, despite its ability to discover many states, because it will dedicate significant computational resources to sampling irrelevant (e.g. high-energy) states. On the other hand, FAST simulations are very directed.

On this funneled landscape, FAST not only has a higher probability of discovering the states along the highest-flux pathways to the global minimum, but also provides the best estimates of their transition probabilities. While there are many ways to estimate transition probabilities to construct an MSM from trajectories, we compare sampling results by row-normalizing transition counts. This is a straightforward method that works well with adaptive sampling data, as is described in a later section. Using a relative entropy metric to quantify the deviation of MSMs built with each method from the true landscape, as we have done previously,^{42,46} we find that FAST and long simulations have the lowest deviations for states in the top three highest-flux pathways. These relative entropies, ascending, are 0.58 ± 0.46 , 0.84 ± 0.80 , 1.96 ± 0.76 , and $2.46 \pm 5 \times 10^{-2}$ for FAST, long, counts, and parallel simulations, respectively. This result suggests that FAST matches long simulations' ability to reach distant conformations, parallel simulations' ability to thoroughly explore particular regions of conformational space, and adaptive sampling's flexibility.

Taken together, the funneled landscape provides a coarse view of each sampling method's behavior. With the perspective that aggregate simulation time is a finite resource, we can imagine the differences between sampling methods being the amount of this resource spent on each region of the conformational landscape. Parallel simulations spend the majority of this resource around the starting state. Counts-based simulations spread it across the landscape. Long simulations distribute it in proportion to neighboring states' energy. FAST spends computing resources on the states that optimize its objective. On the funneled landscape, there are minimal barriers to prevent counts from spreading, and the states that optimize FAST's objective are nearly a straight line from the start to the target. In the following section, we add a layer of complexity to see if adaptive sampling can truly adapt to roadblocks in energy landscapes.

Adaptive sampling navigates obstacles

To mimic the complexities of more realistic landscapes, we generated the rugged landscape in Figure 5A. This rugged landscape provides an interesting challenge to not just discover the target state, but also discover the preferred pathways. The three highest-flux pathways

between the start and the target state are shown in Figure 5A, which each require navigation around large energy barriers. As an added difficulty, there exist alternative routes around the barriers, with differing fluxes. Although sampling is stochastic, and any individual run has the potential to proceed along an arbitrary path, we expect the distribution of paths to resemble the actual highest-flux paths. Of special interest is how FAST navigates the landscape, since it strongly uses structural information in reseeding simulations. We wish to confirm that it does not cut across high-energy barriers in an effort to maximize its objective function.

Similar to the performance on the previous landscape, FAST outperforms the alternative approaches in discovering the target state. This is best seen from each method's discover probabilities on this landscape (Figure S3), where FAST clearly has the highest probability of discovering the target state. In addition, FAST is most likely to discover the states along the actual highest-flux pathways, which suggests that FAST also predicts the correct pathways. To better quantify this, we characterize the probability that a state is predicted to be on pathway from the start to the target. This is done by calculating the highest-flux pathway for each of our 5,000 trials and determining the number of times a state is observed. Averaging this leaves us with a state value of 1 if it is always observed when transitioning from the start to the target, and 0 if it is never observed.

Inspection of the pathway probabilities for FAST (Figure 5B) reveals its ability to navigate around obstacles. The predicted pathways from the start to the target do not pathologically cut across the energy barriers, but mimic the route taken by the three highest-flux pathways that were calculated from the underlying transition probabilities. Furthermore, the predicted pathways of FAST and counts resemble the predicted pathways obtained from the long simulations (Figure S4). This is consistent with the hypothesized benefits of adaptive and goal-oriented sampling: since each simulation is in local equilibrium, the probability of traversing any individual barrier remains unchanged, and thus, transitions will typically occur along realistic pathways.

Pathway tunneling: observing an unfavorable pathway due to sampling artifacts

The landscapes considered so far have been well suited for use with FAST, largely because the simple geometric function used in our FAST ranking (i.e. distance to the target state) is a reasonable surrogate for kinetic proximity to the target. However, there are many instances where finding a reasonable surrogate may be difficult. For example, there are many systems where transitioning between geometrically similar conformations may require partial unfolding of a protein.⁵⁴ In these cases, the optimal transition path would have, at times, unfavorable state rankings.

To explore the utility of FAST when the preferred pathway is suboptimally described by the geometric ranking function, we modeled a landscape with a large barrier separating the start and target states (Figure 6). Here, the three highest-flux pathways all circumnavigate this large barrier rather than taking the geometrically shortest path (across the barrier). Indeed, the long simulations also indicate that the preferred pathway does not cut across the barrier, but follows the longer, low-energy route (Figure 7A–B).

This landscape highlights a potential pathology of running many short parallel simulations, which consistently predict that the highest-flux pathway cuts across the high-energy barrier. From Figure 7C, we observe that the probability that one of the short simulations completes the long path is significantly less than the probability that it hops across the high energy barrier. This leads to the prediction of a very unrealistic highest-flux pathway, as shown in Figure 7D. We name this undesired phenomenon “pathway tunneling”, due to its loose similarity to the tunneling through high-energy barriers observed in quantum mechanics. If the length of all the parallel simulations is gradually increased, the probability of pathway tunneling falls monotonically and converges on the correct mechanism. In this example, pathway tunneling is a consequence of not discovering the set of states along the optimal path, although, it can also arise from poor estimates of transition probabilities that result in an overestimate of the probabilities of rare paths (due to insufficient sampling or inaccuracies in MSM construction). We explore the role of MSM estimators on adaptive sampling data in a later section.

From the discover probabilities in Figure 8A, we observe that FAST has a significantly higher probability of discovering the states along the preferred pathway compared to those of the tunneled pathway. It appears that even in the extreme case where the directed component is at times orthogonal to the preferred pathway, FAST’s statistical component mitigates pathway tunneling. This is evidenced from counts-based adaptive samplings’ ability to discover the correct pathway (Figure S6–7). However, despite this benefit, compared to the long simulations there is an increased probability of discovering the tunneled states. This isn’t an issue if the estimates of the transition probabilities are accurate enough to distinguish the likelihood of each path, although the pathway probabilities in Figure 8B show that FAST non-negligibly predicts the tunneled pathway as the preferred pathway. This result suggests that although FAST is able to circumnavigate orthogonal energy barriers, poor selection of a geometric criterion may lead to over estimating the probability of traversing unfavorable paths.

FAST-string quickly discriminates between alternative pathways

To minimize the probability that FAST falls victim to pathway tunneling, we introduce a new ranking scheme for FAST that refines the transition probabilities along the highest-flux pathways to quantify their relative weights. This method draws inspiration from the string method,^{55–57} which refines a proposed transition path by iteratively running short molecular dynamics simulations from regularly spaced conformations along the path and letting them relax towards the true lowest free energy path. Here, we begin FAST-string after first discovering a pathway, or set of pathways, to the target state using the original FAST rankings. Then, we change the ranking function to focus on refining the transition probabilities of the path(s) found. Specifically, we calculate the n -highest-flux pathways and rank states found in these paths by some statistical criterion. Thus, our state rankings become:

$$r(i) = \begin{cases} \bar{\psi}(i) & \text{if } i \in \{w_0, \dots, w_n\} \\ 0 & \text{otherwise} \end{cases} \quad [7]$$

where $r(i)$ is the ranking of state i , $\bar{\psi}(i)$ is the scaled statistical component of the original FAST ranking function, and $\{w_0, \dots, w_n\}$ represents the states found in the n -highest-flux paths. For our purposes, we use the counts of each state as our statistical component to favor less explored regions of the predicted pathways. We expect that sampling along these states will distinguish favorable paths from unfavorable, if multiple paths are discovered, and help relax the pathway to the preferred path if pathway tunneling has occurred.

With our FAST-string method, we are able to consistently determine the preferred transition path. Figure 9A shows that the tunneled pathway is no longer predicted as the transition pathway. We obtain this result with the same amount of aggregate simulation as the other methods; we run FAST until it discovers the end state, then switch to FAST-string for the remainder of the rounds. Instead of redundantly sampling around the target state once found, FAST-string productively refines estimates of the transition probabilities. From Figure 9B–D, we can see that FAST-string has the most accurate estimates of transition probabilities along the highest-flux pathways.

Normalizing raw counts provides a good balance of estimating rates and equilibrium populations with adaptive sampling data

In addition to comparing different sampling methods, it is important to ask what the best way of estimating the transition probabilities between states from a given data set is. In other words, what is the best way to use a count-matrix, which counts the observed transitions between every pair of states, to estimate the transition probabilities and equilibrium populations of each state?

The simplest way is to normalize each row in the count-matrix to get an unbiased estimate of each states conditional transition probabilities, where the first eigenvector provides the equilibrium populations.⁴⁵ However, this approach does not guarantee microscopic reversibility and can have serious pathologies if the transition probability matrix is not ergodic, especially if transitions are observed from state n_i to n_j but not in the opposite direction. To alleviate this issue, it is customary to assume that, prior to observing any data, each state has equal probability to transition to any other state. This can be represented by adding a pseudocount, $\tilde{C}$, to each possible transition,

$$T_{ij}^{normalize} = \frac{C_{ij} + \tilde{C}}{\sum_k (C_{ik} + \tilde{C})} \quad [8]$$

where,

$$\tilde{C} = \frac{1}{n} \quad [9]$$

and n is the number of states. An alternative estimator, called the transpose method, enforces detailed balance. At equilibrium, we know that each state transition should be equally populated by the reverse process (running an infinitely long simulation in reverse should not

alter the estimates for transition probabilities). Enforcing this is straightforward, by averaging with the transpose of the count-matrix:

$$C_{ij}^{transpose} = \frac{C_{ij} + C_{ji}}{2}$$

and

$$T_{ij}^{transpose} = \frac{C_{ij}^{transpose}}{\sum_k C_{ik}^{transpose}}$$

and the equilibrium populations are calculated as,

$$\pi_i = \frac{\sum_j C_{ij}^{transpose}}{\sum_{k,j} C_{kj}^{transpose}} \quad [10]$$

This has recently been extended for use with simulations at multiple temperatures.⁵⁸ More sophisticated methods have also been developed to enforce detailed balance, such as the use of maximum likelihood estimation (MLE),^{23,24} and the observable operator model (OOM).⁵⁹ In the MLE method, the likelihood of the transition probability matrix given an observed trajectory, $\mathbf{X}$, is determined to be,

$$P(T|\mathbf{X}) \propto \prod_{i,j} T_{ij}^{C_{ij}}$$

Consequently, the most likely transition probability matrix is solved as,

$$T_{ij}^{MLE} = \arg \max_{T_{ij}^*} P(T_{ij}^*|\mathbf{X})$$

A variant of the MLE method, which we will refer to as MLE-CP (MLE-with Constrained Populations), has also been developed to enforce a pre-determined equilibrium probability distribution.^{60,61} This is useful with experimental estimates of state populations. Lastly, the OOM was recently developed as a generalization to hidden Markov models,⁶² and restructured for use with MSMs.^{59,63}

Each of these methods has been studied theoretically, in the limit of infinite data, and on small systems where sampling is not an issue. However, we are interested in the likely scenario where sampling is far from exhaustive. To test MSM construction in this regime, we used the FAST simulation sets on the landscape in Figure 5A to generate MSMs using the five methods listed above: 1) normalize, 2) transpose, 3) MLE, 4) OOM, and 5) MLE with constrained populations. We then compared the MSMs predictions of thermodynamics (equilibrium populations) and kinetics (transition probabilities) to the true distributions

calculated from the underlying landscape. While the performance of each estimator may depend on the particular landscape being sampled, the case study we present here is representative of our results with other landscapes. Our metrics for performance consists of how well each estimator predicts the kinetics (transition rates) and thermodynamics (equilibrium populations) of the underlying energy landscape, as quantified with a relative entropy metric.

Upon inspection of the predicted equilibrium populations, we find that the MLE and OOM methods have a tendency to significantly overestimate the populations of an arbitrary set of states. Figure 10A is an example of this phenomenon for MLE, where four states are predicted to have a total probability of 0.58 even though the probability that they were sampled in the raw data is only 0.016. For reference, the true total probability of these states is 0.029 and the true probability of the most populated state in the underlying landscape is 0.032. In comparison, Figure S9 shows that normalize and transpose give more reasonable predictions. Characterizing this over the entire dataset, we observe that on average, the largest predicted state population for MLE and OOM is 10.5 ± 21.0 and 16.4 ± 53.5 times larger than its true population. Interestingly, the deviation in these predictions are sizable; the most egregious observances of an overinflated state population for MLE and OOM were predictions of a single state containing 0.38 and 0.55 of the total population for each method, respectively. Additionally, MLE and OOM do not regularly overpopulate the same state; the probability that the state with the largest predicted population is truly the most populated state is 0.16 and 0.14 for MLE and OOM, respectively, compared to 0.35 and 0.33 for normalize and transpose, respectively. On the other hand, normalize and transpose have a largest populated state that is only 2.2 ± 2.5 and 2.0 ± 2.7 times its true population. However, while OOM is subject to the same pathology as MLE, severely over estimating the populations for a set of states, it appears to describe most of the landscape quite well (Figure S9). Future developments of OOM could provide an accurate and robust estimator. To further quantify predictions for all states, we compute the relative entropy between each models' prediction of transition probabilities and equilibrium populations to the true distributions.

The MLE-CP is shown to generate an MSM with the most accurate estimates of kinetics and thermodynamics for FAST simulations on this particular landscape. Figure 10B shows the average deviations of transition probabilities and equilibrium populations from the true values for the underlying landscape for each MSM method. It is not surprising that constraining the populations to their true values performs well. Also, as has been previously reported, there are significant improvements to estimates of transition probabilities when the equilibrium populations are constrained.⁶¹ However, *a priori* knowledge of the equilibrium distribution is not typically available, so it is not currently possible to adopt this approach as standard practice.

The normalize and transpose methods produce the next most accurate estimates of transition probabilities and equilibrium populations. However, despite transposes' adequate performance on this landscape, it can be shown from equation 10 that the estimated equilibrium populations are directly related to the amount of sampling in each state. This is not thought to be ideal with adaptive sampling, since continually sampling from a state will

artificially inflate its estimated equilibrium population. Transpose does well on this particular landscape due to the relatively flat energy surface of the preferred path and would be less favorable with real landscapes. Therefore, we recommend the use of the normalize method with adaptive sampling data for its simplicity and accurate estimates of thermodynamics and kinetics.

Simulations of λ -repressor recapitulate the patterns observed for simple landscapes

Kinetic Monte Carlo simulations on physically inspired landscapes have provided valuable functional insight, but it is important to ensure that our conclusions hold true for the exploration of real protein landscapes. Protein conformational landscapes are hyperdimensional and likely have many barriers, both enthalpic and entropic. Thus, we turn to using all-atom MD simulations for three sampling methods: 1) long simulations, 2) massively parallel simulations, and 3) FAST-contacts (which ranks states by the fraction of native contacts that are present). Each method uses the same unfolded starting structure and simulation parameters, where extended details are described in the Methods. As for a model system, we chose to simulate a fast-folding variant of the λ -repressor.⁶⁴ Due to its speed of folding and size, the kinetics of this protein have been extensively studied, both experimentally and computationally, making it ideal for use when comparing sampling strategies.

Unlike the simple landscapes in previous sections, all-atom MD simulations are computationally expensive and sample along vast conformational landscapes. As a consequence, we cannot run thousands of iterations to robustly characterize the probability of discovering a particular state. Instead, we can compare the performance of each method by focusing on a more coarse-grained metric of interest, such as the computational time required to reach the folded state, as measured by the fraction of native contacts present.

Analysis of the three sampling methods reveals that adaptive sampling yields similar benefits to those found on our simple landscapes. Figure 11 shows the highest fraction of native contacts observed for each sampling method as a function of the aggregate simulation time. Remarkably, FAST-contacts folds the λ -repressor with $\sim 4 \mu\text{s}$ of aggregate simulation, which is faster than its experimental folding time. By comparison, it takes nearly $40 \mu\text{s}$ of long simulations to achieve a similar level of foldedness. Furthermore, the massively parallelized simulations, with over $700 \mu\text{s}$ of aggregate simulation time, and counts-based adaptive sampling do not discover the folded state. Due to the high dimensionality of the λ -repressor, compared with our generated landscapes, counts-based adaptive sampling appears to hinder discovery of the folded state; low count states are continually discovered and selected in orthogonal directions to the fraction of native contacts. These results are in strong agreement with the discovery predictions from the landscapes in Figures 3 and 5.

In addition to understanding the probability of observing a folded state, we are interested in the predicted folding pathways. However, the idea of characterizing a pathway for all-atom MD simulations is more complicated than on the theoretical landscapes; state-space is significantly larger, computational limitations prevent multiple trials to assess the stochasticity, and the optimal (human-intuitive) parameters to define a pathway are not straightforward. The long and parallel simulations require too much computational resources

to gather statistics on, although we were able to generate five independent trials of FAST in a reasonable timeframe. For the purposes of defining a pathway, others have successfully taken the approach of characterizing folding by the rate of formation of secondary structural elements.^{65–67} Thus, we also aim to characterize the rate of secondary structure formation by determining each residues' root-mean-square deviation (RMSD) of backbone dihedrals from the crystal structure, for states along the predicted highest-flux pathway. We plot these deviations for two representative runs of FAST-contacts and the long simulations in Figure 12.

The predicted pathways for each of these methods are reasonably consistent with one another. FAST-contacts predominantly predicts the folding of helices 1 and 4 before helices 2 and 3. This is consistent with our prediction using the single set of long simulations. Additionally, this is what has been seen with previous simulation reports,^{68,69} and hydrogen exchange experiments.⁷⁰ Interestingly, this is counter to the results from a Gō model, which has been previously used and describes helices 1–4 folding cooperatively.⁷¹ This difference suggests that FAST-contacts is not simply an expensive Gō model.

Conclusions

We have presented a systematic comparison of different sampling strategies on a variety of representative energy landscapes. We first developed an analytic expression for the probability of discovering states on a landscape that depends on the number, length, and starting state of simulations. From this we find that long simulations have a higher probability of discovering states on landscapes with reduced dimensionality, though parallel simulations have a higher probability of discovering states as the dimensionality increases. To build upon this, we used kinetic Monte Carlo simulations on more complex landscapes to compare four sampling strategies (long simulations, parallel simulations, counts adaptive sampling, and FAST), which each reveal a unique state discovery signature. Understanding the differences in how these sampling strategies discover states has provided insight into their advantages and disadvantages. Specifically, long simulations provide an unbiased estimates of transition paths, although requires significant computational resources compared to adaptive sampling or FAST and produces less accurate MSMs. Parallel simulations thoroughly explore around the starting state and provide excellent estimates of transition probabilities (for the states discovered) but are unlikely to explore distant regions of conformational space and may provide erroneous transition paths. Counts-based adaptive sampling discovers the most states along a variety of paths, although these states are likely to be unproductive for a given goal, especially on landscapes with large dimensionality.

Throughout our analysis, we have taken special interest in the performance of our recently developed goal-oriented sampling algorithm, FAST. On our simple landscapes, we find that FAST consistently has the highest probability of discovering a target state, predicts reasonable pathways, and provides the best estimates of transition probabilities for an entire MSM as well as of the true highest-flux pathway (Table S1). Furthermore, we demonstrate the utility of FAST using all-atom MD simulations of the λ -repressor. FAST produces an accurate folding pathway with an order of magnitude less aggregate simulation than long simulations, and orders of magnitude less than parallel simulations.

Methods

Generation and simulation of simple landscapes

The three physically inspired potential energy landscapes were generated by selectively adding Gaussian potentials to an otherwise flat surface. These potential energy landscapes were then converted to a transition probability matrix using the following relations:

$$\zeta_{ij} = \begin{cases} e^{\varepsilon_i - \varepsilon_j} & \text{if } \varepsilon_i < \varepsilon_j \\ 1 & \text{if } \varepsilon_i \geq \varepsilon_j \end{cases}$$

for all j that are neighbors of i , and where ε_i is the potential energy of state n_i in units of $k_B T$. This can then be row-normalized to obtain,

$$T_{ij} = \frac{\zeta_{ij}}{\sum_j \zeta_{ij}}$$

Kinetic Monte Carlo simulations were then performed with this transition probability matrix for four sampling schemes: 1) long simulations, 2) parallel simulations, 3) counts-based adaptive sampling, and 4) FAST simulations. For each of the sampling schemes, 5,000 independent sets of simulations were generated, each with a total of 1,000 time-steps. For the long simulations, this consisted of 5,000 single trajectories, of 1,000 steps. Each set for the parallel simulations consisted of 25 trajectories with 40 steps.

Counts-based adaptive sampling and FAST both followed the same basic protocol: 1) generate 5 trajectories of 20 steps each from the initial state, 2) build an MSM, 3) rank states, 4) generate 5 more trajectories of 20 steps each from the top 5 states with the highest ranking, 5) repeat steps 2-4 for a total of 10 rounds. The difference between counts-based adaptive sampling and FAST is in the manner of ranking states between each round. For counts adaptive sampling, states were ranked by their observed counts in the MSM, with lower counts being more favorable. For FAST, we used the following ranking,

$$r_\phi(i) = \bar{\phi}(i) + \alpha \bar{\psi}(i) + \beta \chi(i) \quad [11]$$

where $\bar{\phi}$ is the feature-scaled directed component (Euclidean distance to the target state), $\bar{\psi}$ is the feature scaled undirected component, χ is a similarity penalty, and α and β control the weights of $\bar{\psi}$ and χ , respectively, as has been published previously.²¹ Here, $\bar{\psi}(i)$ is taken to be the state counts and a value of 1 was used for both α and β . The directed component for each state on the landscapes was the grid distance to the target state. The similarity penalty for each state selected is defined with,

$$\chi(i) = \begin{cases} 0 & \text{if } N = 0 \\ \frac{1}{N} \sum_{j=1}^N \left(1 - e^{\frac{-d_{ij}^2}{2w^2}} \right) & \text{if } N > 0 \end{cases} \quad [12]$$

which is the average of the Gaussian weighted grid distance, d , from state n_i to the N states that have been selected for reseeding so far, where w is the Gaussian width (set to the clustering radius). Thus, selecting states proceeds as follows: 1) rank all states by the FAST ranking and select the top state, 2) add the similarity penalty and select the top-ranking state as the next state, 3) repeat step 2 until the desired number of states have been selected.

After generating the state trajectories on the landscapes from the sampling methods, state discover probabilities, pathway probabilities, and relative entropies were calculated. The discover probabilities were calculated by first using equation 1 to indicate if a state was discovered for each simulation set. These values for $D_{ij}^{\mathbf{K}, \mathbf{M}}$ were then averaged over the 5,000 trials to determine the probability of discovering a state in the simulation set, $P(D_{ij}^{\mathbf{K}, \mathbf{M}} = 1)$. Similar to the discover probabilities, the pathway probabilities were calculated by averaging the output of a selector function, over the simulation sets, that indicated if a state was present in the predicted highest-flux pathway. The highest-flux pathway for each simulation set was calculated using MSMBuilders.⁷² The relative entropies of each state were calculated as the Kullback-Leibler divergence between the estimated conditional transition probabilities from that state and those of the underlying energy distribution:

$$D_{KL}^i(P_i || Q_i) = - \sum_i P_i \log \left(\frac{Q_i}{P_i} \right)$$

where D_{KL}^i is the relative entropy for state i , P_i is the i -th row of the true transition probability matrix, and Q_i is the i -th row of the transition probability matrix reconstructed from synthetic trajectories. The relative entropy of the entire MSM is a population weighted average of these values, as is described previously.^{42,46} MSMs were constructed with either the MSMBuilders or PyEMMA software packages.⁷²⁻⁷⁴

Molecular Dynamics Simulations

Four sets of all-atom molecular dynamics simulations for the λ -repressor were generated: 1) 7,005 parallel simulations (103.4 ± 82.0 ns each), 2) 16 long simulations (2.5 μ s each), 3) FAST-contacts simulations (30 rounds of 10 simulations per round, with 30 ns per simulation), and 4) counts-based adaptive sampling (30 rounds of 10 simulations per round, with 30 ns per simulation). Each of these simulations were run with Gromacs 5.1.1⁷⁵ using the AMBER03 force field with explicit TIP3P solvent.^{76,77}

Each of these sets of simulations began from the same starting structure, which was prepared as follows. First, a linear structure of the λ D14A mutant⁶⁴ was generated using the VMD software package.⁷⁸ The linear structure was equilibrated for 1 ns at 420 K with OBC GBSA implicit solvent.⁷⁹ The final conformation was then placed in a dodecahedron box that extended 1.0 nm beyond the protein in any dimension, with a total of 46,450 atoms in the system. This system was then energy minimized with the steepest descent algorithm until the maximum force fell below 100 kJ/mol/nm using a step size of 0.01 nm and a cutoff distance of 1.2 nm for the neighbor list, Coulomb interactions, and van der Waals interactions.

For production runs, all bonds were constrained with the LINCS algorithm and virtual sites were used to allow a 4 fs time step. Cutoffs of 1.0 nm were used for the neighbor list, Coloumb interactions, and van der Waals interactions. The Verlet cutoff scheme was used for the neighbor list. The stochastic velocity rescaling (ν -rescale) thermostat was used to hold the temperature at 360 K and conformations were stored every 50 ps.⁸⁰

FAST Simulations

Five sets of FAST-contacts simulations were generated that each observed an independent folding trajectory for the λ -repressor. Each set of FAST-contacts consisted of 9 μ s of aggregate simulation time: 30 rounds, of 10 simulations per round, where each simulation was 30 ns. Between each round, discrete states were generated by clustering atomic coordinates of backbone atoms using a k -centers algorithm based on RMSD between conformations until every cluster center had a radius less than 3.0 Å. States were selected for reseeding based on the ranking function and selection criterion described with equations 11 and 12. The similarity penalty used was RMSD between cluster centers, where the Gaussian width, w , was set to the clustering radius of 3.0 Å. The directed component to the FAST ranking was the feature scaled values of the fraction of native contacts, described elsewhere.⁸¹

MSM Construction and Analysis

MSMs were built of each simulation set using MSMBUILDER.^{72,73} The construction of each MSM followed the same basic protocol: 1) cluster conformations into discrete states, 2) count transitions between these states at a specified lag-time, and 3) generate each states' conditional transition probabilities. For the first step, atomic coordinates of backbone heavy atoms (CO, C α , O, N) and C β atoms were clustered with a k -centers clustering algorithm until every cluster center had a radius of less than 3.0 Å. A lag-time of 5 ns was used for counting transitions between states. Each states' conditional transition probabilities were computed using the normalize method with a prior-counts, as described with equations 8 and 9. Structural analysis was aided with the use of MDTraj.⁸²

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This work was funded by NSF CAREER Award MCB-1552471 and NIH R01GM12400701. G.R.B. holds a Career Award at the Scientific Interface from the Burroughs Wellcome Fund and a Packard Fellowship for Science and Engineering from The David and Lucile Packard Foundation M.I.Z. holds a Monsanto Graduate Fellowship and a Center for Biological Systems Engineering Fellowship.

References

- (1). Shaw DE; Deneroff MM; Dror RO; Kuskin JS; Larson RH; Salmon JK; Young C; Batson B; Bowers KJ; Chao JC; Eastwood MP; Gagliardo J; Grossman JP; Ho CR; Ierardi DJ; Kolossváry I; Klepeis JL; Layman T; McLeavey C; Moraes MA; Mueller R; Priest EC; Shan Y; Spengler J; Theobald M; Towles B; Wang SC Anton, a Special-Purpose Machine for Molecular Dynamics Simulation. *Communications of the ACM* 2008, 51 (7), 91–97.
- (2). Shaw DE; Grossman JP; Bank JA; the BBPO; 2014 Anton 2: Raising the Bar for Performance and Programmability in a Special-Purpose Molecular Dynamics Supercomputer. dl.acm.org.
- (3). Takada S Go-Ing for the Prediction of Protein Folding Mechanisms. *Proc. Natl. Acad. Sci. U.S.A.* 1999, 96 (21), 11698–11700. [PubMed: 10518512]
- (4). Go N Theoretical Studies of Protein Folding. *Annu. Rev. Biophys. Bioeng* 1983, 12 (1), 183–210. [PubMed: 6347038]
- (5). Sugita Y; Okamoto Y Replica-Exchange Molecular Dynamics Method for Protein Folding. *Chemical Physics Letters* 1999, 314 (1-2), 141–151.
- (6). Fukunishi H; Watanabe O; Takada S On the Hamiltonian Replica Exchange Method for Efficient Sampling of Biomolecular Systems: Application to Protein Structure Prediction. *The Journal of Chemical Physics* 2002, 116 (20), 9058–9067.
- (7). Faraldo-Gómez JD; Roux B Characterization of Conformational Equilibria Through Hamiltonian and Temperature Replica-Exchange Simulations: Assessing Entropic and Environmental Effects. *J Comput Chem* 2007, 28 (10), 1634–1647. [PubMed: 17342721]
- (8). Izrailev S; Stepaniants S; Isralewitz B; Kosztin D; Lu H; Molnar F; Wriggers W; Schulten K Steered Molecular Dynamics In Computational Molecular Dynamics: Challenges, Methods, Ideas; *Lecture Notes in Computational Science and Engineering*; Springer Berlin Heidelberg: Berlin, Heidelberg, 1999; Vol. 4, pp 39–65.
- (9). Isralewitz B; Gao M; Schulten K Steered Molecular Dynamics and Mechanical Functions of Proteins. *Current Opinion in Structural Biology* 2001, 11 (2), 224–230. [PubMed: 11297932]
- (10). Voter AF Hyperdynamics: Accelerated Molecular Dynamics of Infrequent Events. *Physical Review Letters* 1997, 78 (20), 3908–3911.
- (11). Hamelberg D; Mongan J; McCammon JA Accelerated Molecular Dynamics: a Promising and Efficient Simulation Method for Biomolecules. *The Journal of Chemical Physics* 2004, 120 (24), 11919–11929. [PubMed: 15268227]
- (12). Laio A; Parrinello M Escaping Free-Energy Minima. *Proc. Natl. Acad. Sci. U.S.A* 2002, 99 (20), 12562–12566. [PubMed: 12271136]
- (13). Laio A; Gervasio FL Metadynamics: a Method to Simulate Rare Events and Reconstruct the Free Energy in Biophysics, Chemistry and Material Science. *Reports on Progress in Physics* 2008, 71 (12), 126601.
- (14). Perez A; MacCallum J; Dill KA Meld: Modeling Peptide-Protein Interactions. *Biophysical Journal* 2013, 104 (2), 399a.
- (15). Perez A; MacCallum JL; Dill KA Accelerating Molecular Simulations of Proteins Using Bayesian Inference on Weak Information. *Proc. Natl. Acad. Sci. U.S.A* 2015, 112 (38), 11846–11851. [PubMed: 26351667]
- (16). Zheng L; Chen M; Yang W Random Walk in Orthogonal Space to Achieve Efficient Free-Energy Simulation of Complex Systems. *Proc. Natl. Acad. Sci. U.S.A* 2008, 105 (51), 20227–20232. [PubMed: 19075242]
- (17). Shirts MR; Pande VS Mathematical Analysis of Coupled Parallel Simulations. *Physical Review Letters* 2001, 86 (22), 4983–4987. [PubMed: 11384401]

- (18). Shirts M COMPUTING: Screen Savers of the World Unite! *Science* 2000, 290 (5498), 1903–1904. [PubMed: 17742054]
- (19). Bowman GR; Pande VS; Noé F Introduction and Overview of This Book. In *An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation*; *Advances in Experimental Medicine and Biology*; Springer Netherlands: Dordrecht, 2014; Vol. 797, pp 1–6. [PubMed: 24297271]
- (20). Hart KM; Ho CMW; Dutta S; Gross ML; Bowman GR Modelling Proteins' Hidden Conformations to Predict Antibiotic Resistance. *Nature Communications* 2016, 7, 12965.
- (21). Zimmerman MI; Hart KM; Sibbald CA; Frederick TE; Jimah JR; Knoverek CR; Tolia NH; Bowman GR Prediction of New Stabilizing Mutations Based on Mechanistic Insights From Markov State Models. *ACS Cent. Sci* 2017, 3 (12), 1311–1321. [PubMed: 29296672]
- (22). Hart KM; Moeder KE; Ho CMW; Zimmerman MI; Frederick TE; Bowman GR Designing Small Molecules to Target Cryptic Pockets Yields Both Positive and Negative Allosteric Modulators. *PLOS ONE* 2017, 12 (6), e0178678. [PubMed: 28570708]
- (23). Bowman GR; Beauchamp KA; Boxer G; Pande VS Progress and Challenges in the Automated Construction of Markov State Models for Full Protein Systems. *The Journal of Chemical Physics* 2009, 131 (12), 124101. [PubMed: 19791846]
- (24). Metzner P; Noe F; Schütte C Estimating the Sampling Error: Distribution of Transition Matrices and Functions of Transition Matrices for Given Trajectory Data. *Physical Review E* 2009, 80 (2), 021106.
- (25). Pande VS; Beauchamp K; Bowman GR Everything You Wanted to Know About Markov State Models but Were Afraid to Ask. *Methods* 2010, 52 (1), 99–105. [PubMed: 20570730]
- (26). Chodera JD; Noe F Markov State Models of Biomolecular Conformational Dynamics. *Current Opinion in Structural Biology* 2014, 25, 135–144. [PubMed: 24836551]
- (27). Mukherjee S; Pantelopulos GA; Voelz VA Markov Models of the <I>Apo</I>-MDM2 Lid Region Reveal Diffuse Yet Two-State Binding Dynamics and Receptor Poses for Computational Docking. *Scientific Reports* 2016 6 2016, 6 (1), 31631.
- (28). Zhou G; Pantelopulos GA; Mukherjee S; Voelz VA Bridging Microscopic and Macroscopic Mechanisms of P53-MDM2 Binding with Kinetic Network Models. *Biophysical Journal* 2017, 113 (4), 785–793. [PubMed: 28834715]
- (29). Plattner N; Doerr S; De Fabritiis G; Noé F Complete Protein–Protein Association Kinetics in Atomic Detail Revealed by Molecular Dynamics Simulations and Markov Modelling. *Nature Chemistry* 2017 9:10 2017, 9 (10), 1005–1011.
- (30). Wang W; Cao S; Zhu L; Huang X Constructing Markov State Models to Elucidate the Functional Conformational Changes of Complex Biomolecules. *Wiley Interdisciplinary Reviews: Computational Molecular Science* 2018, 8 (1), e1343.
- (31). Husic BE; Pande VS Markov State Models: From an Art to a Science. *Journal of the American Chemical Society* 2018, 140 (7), 2386–2396. [PubMed: 29323881]
- (32). Hinrichs NS; Pande VS Calculation of the Distribution of Eigenvalues and Eigenvectors in Markovian State Models for Molecular Dynamics. *The Journal of Chemical Physics* 2007, 126 (24), 244101. [PubMed: 17614531]
- (33). Bowman GR; Ensign DL; Pande VS Enhanced Modeling via Network Theory: Adaptive Sampling of Markov State Models. *J. Chem. Theory Comput* 2010, 6 (3), 787–794. [PubMed: 23626502]
- (34). Weber JK; Pande VS Characterization and Rapid Sampling of Protein Folding Markov State Model Topologies. *J. Chem. Theory Comput* 2011, 7 (10), 3405–3411. [PubMed: 22140370]
- (35). Doerr S; De Fabritiis G On-the-Fly Learning and Sampling of Ligand Binding by High-Throughput Molecular Simulations. *J. Chem. Theory Comput* 2014, 10 (5), 2064–2069. [PubMed: 26580533]
- (36). Voelz VA; Elman B; Razavi AM; Zhou G Surprisal Metrics for Quantifying Perturbed Conformational Dynamics in Markov State Models. *J. Chem. Theory Comput* 2014, 10 (12), 5716–5728. [PubMed: 26583253]

- (37). Bacci M; Vitalis A; Caflisch A A Molecular Simulation Protocol to Avoid Sampling Redundancy and Discover New States. *Biochimica et Biophysica Acta (BBA) - General Subjects* 2015, 1850 (5), 889–902. [PubMed: 25193737]
- (38). Kukharensko O; Sawade K; Steuer J; Peter C Using Dimensionality Reduction to Systematically Expand Conformational Sampling of Intrinsically Disordered Peptides. *J. Chem. Theory Comput* 2016, 12 (10), 4726–4734. [PubMed: 27588692]
- (39). Doerr S; Harvey MJ; Noe F; De Fabritiis G HTMD: High-Throughput Molecular Dynamics for Molecular Discovery. *J. Chem. Theory Comput* 2016, 12 (4), 1845–1852. [PubMed: 26949976]
- (40). Sultan MM; Pande VS Decision Functions From Supervised Machine Learning Algorithms as Collective Variables for Accelerating Molecular Simulations. 2018, 28, 2018.
- (41). Noé F; Banisch R; Clementi C Commute Maps: Separating Slowly Mixing Molecular Configurations for Kinetic Modeling. *ACS Publications* 2016, 12 (11), 5620–5630.
- (42). Zimmerman MI; Bowman GR FAST Conformational Searches by Balancing Exploration/Exploitation Trade-Offs. *J. Chem. Theory Comput* 2015, 11 (12), 5747–5757. [PubMed: 26588361]
- (43). Zimmerman MI; Bowman GR How to Run FAST Simulations. *Methods in Enzymology* 2016, 578, 213–225. [PubMed: 27497168]
- (44). Bowman GR; Geissler PL Equilibrium Fluctuations of a Single Folded Protein Reveal a Multitude of Potential Cryptic Allosteric Sites. *Proc. Natl. Acad. Sci. U.S.A* 2012, 109 (29), 11681–11686. [PubMed: 22753506]
- (45). Grinstead CM; Snell JL Introduction to Probability; 2012.
- (46). Bowman GR; Pande VS Protein Folded States Are Kinetic Hubs. *Proc. Natl. Acad. Sci. U.S.A* 2010, 107 (24), 10890–10895. [PubMed: 20534497]
- (47). Wang Y; Lv J; Zhu L; Ma Y Crystal Structure Prediction via Particle-Swarm Optimization. *Phys. Rev. B* 2010, 82 (9), 094116.
- (48). Moult J A Decade of CASP: Progress, Bottlenecks and Prognosis in Protein Structure Prediction. *Current Opinion in Structural Biology* 2005, 15 (3), 285–289. [PubMed: 15939584]
- (49). Floudas CA; Fung HK; McAllister SR; Mönnigmann M; Rajgaria R Advances in Protein Structure Prediction and De Novo Protein Design: a Review. *Chemical Engineering Science* 2006, 61 (3), 966–988.
- (50). Voeltz VA; Bowman GR; Beauchamp K; Pande VS Molecular Simulation of Ab Initio Protein Folding for a Millisecond Folder NTL9(1-39). *Journal of the American Chemical Society* 2010, 132 (5), 1526–1528. [PubMed: 20070076]
- (51). Lindorff-Larsen K; Piana S; Dror RO; Shaw DE How Fast-Folding Proteins Fold. *Science* 2011, 334 (6055), 517–520. [PubMed: 22034434]
- (52). Leopold PE; Montal M; Onuchic JN Protein Folding Funnels: a Kinetic Approach to the Sequence-Structure Relationship. *Proc. Natl. Acad. Sci. U.S.A* 1992, 89 (18), 8721–8725 [PubMed: 1528885]
- (53). Dill KA; Chan HS From Levinthal to Pathways to Funnels. *Nature Structural & Molecular Biology* 1997 4:1 1997, 4 (1), 10–19.
- (54). Whitford PC; Onuchic JN What Protein Folding Teaches Us About Biological Function and Molecular Machines. *Current Opinion in Structural Biology* 2015, 30, 57–62. [PubMed: 25559307]
- (55). E W; Ren W.; Vanden-Eijnden E String Method for the Study of Rare Events. *Phys. Rev. B* 2002, 66 (5), 052301.
- (56). Maragliano L; Fischer A; Vanden-Eijnden E; Ciccotti G String Method in Collective Variables: Minimum Free Energy Paths and Isocommittor Surfaces. *The Journal of Chemical Physics* 2006, 125 (2), 024106.
- (57). Pan Albert C; Deniz Sezer A; Benoit Roux. Finding Transition Pathways Using the String Method with Swarms of Trajectories; American Chemical Society, 2008; Vol. 112, pp 3432–3440.
- (58). Leahy CT; Kells A; Hummer G; Buchete N-V; Rosta E Peptide Dimerization-Dissociation Rates From Replica Exchange Molecular Dynamics. *The Journal of Chemical Physics* 2017, 147 (15), 152725. [PubMed: 29055328]

- (59). Nüske F; Wu H; Prinz J-H; Wehmeyer C; Clementi C; Noé F Markov State Models From Short Non-Equilibrium Simulations—Analysis and Correction of Estimation Bias. *The Journal of Chemical Physics* 2017, 146 (9), 094104.
- (60). Trendelkamp-Schroer B; Noe F Efficient Bayesian Estimation of Markov Model Transition Matrices with Given Stationary Distribution. *The Journal of Chemical Physics* 2013, 138 (16), 164113. [PubMed: 23635117]
- (61). Trendelkamp-Schroer B; Noe F Efficient Estimation of Rare-Event Kinetics. *Phys. Rev. X* 2016, 6 (1), 011009.
- (62). Jaeger H Observable Operator Models for Discrete Stochastic Time Series. 10.1162/089976600300015411 2006, 12 (6), 1371–1398.
- (63). Wu H; Prinz J-H; Noé F Projected Metastable Markov Processes and Their Estimation with Observable Operator Models. *The Journal of Chemical Physics* 2015, 143 (14), 144101. [PubMed: 26472357]
- (64). Yang WY; Gruebele M Folding Λ -Repressor at Its Speed Limit. *Biophysical Journal* 2004, 87 (1), 596–608. [PubMed: 15240492]
- (65). Hu W; Walters BT; Kan Z-Y; Mayne L; Rosen LE; Marqusee S; Englander SW Stepwise Protein Folding at Near Amino Acid Resolution by Hydrogen Exchange and Mass Spectrometry. *Proc. Natl. Acad. Sci. U.S.A* 2013, 110 (19), 7684–7689. [PubMed: 23603271]
- (66). Piana S; Lindorff-Larsen K; Shaw DE How Robust Are Protein Folding Simulations with Respect to Force Field Parameterization? *Biophysical Journal* 2011, 100 (9), L47–L49. [PubMed: 21539772]
- (67). Henry ER; Best RB; Eaton WA Comparing a Simple Theoretical Model for Protein Folding with All-Atom Molecular Dynamics Simulations. *Proc. Natl. Acad. Sci. U.S.A* 2013, 110 (44), 17880–17885. [PubMed: 24128764]
- (68). Liu Y; Strumpfer J; Freddolino PL; Gruebele M; Schulten K Structural Characterization of Λ -Repressor Folding From All-Atom Molecular Dynamics Simulations. *ACS Publications* 2012, 3 (9), 1117–1123.
- (69). Bowman GR; Voelz VA; Pande VS Atomistic Folding Simulations of the Five-Helix Bundle Protein $\Lambda 6$ –85. *Journal of the American Chemical Society* 2010, 133 (4), 664–667.
- (70). Yu W; Baxa MC; Gagnon I; Freed KF; Sosnick TR Cooperative Folding Near the Downhill Limit Determined with Amino Acid Resolution by Hydrogen Exchange. *Proc. Natl. Acad. Sci. U.S.A* 2016, 113 (17), 4747–4752. [PubMed: 27078098]
- (71). Shoemaker BA; Wang J; Wolynes PG Exploring Structures in Protein Folding Funnels with Free Energy Functionals: the Transition State Ensemble. *Journal of Molecular Biology* 1999, 287 (3), 675–694. [PubMed: 10092467]
- (72). Harrigan MP; Sultan MM; Hernández CX; Husic BE; Eastman P; Schwantes CR; Beauchamp KA; McGibbon RT; Pande VS MSMBuilder: Statistical Models for Biomolecular Dynamics. *Biophysical Journal* 2017, 112 (1), 10–15. [PubMed: 28076801]
- (73). Beauchamp KA; Bowman GR; Lane TJ; Maibaum L; Haque IS; Pande VS MSMBuilder2: Modeling Conformational Dynamics on the Picosecond to Millisecond Scale. *J. Chem. Theory Comput* 2011, 7 (10), 3412–3419. [PubMed: 22125474]
- (74). Scherer MK; Trendelkamp-Schroer B; Paul F; Pérez-Hernández G; Hoffmann M; Plattner N; Wehmeyer C; Prinz J-H; Noe F PyEMMA 2: a Software Package for Estimation, Validation, and Analysis of Markov Models. *J. Chem. Theory Comput* 2015, 11 (11), 5525–5542. [PubMed: 26574340]
- (75). Abraham MJ; Murtola T; Schulz R; Páll S; Smith JC; Hess B; Lindahl E GROMACS: High Performance Molecular Simulations Through Multi-Level Parallelism From Laptops to Supercomputers. *SoftwareX* 2015, 1-2, 19–25.
- (76). Duan Y; Wu C; Chowdhury S; Lee MC; Xiong G; Zhang W; Yang R; Cieplak P; Luo R; Lee T; Caldwell J; Wang J; Kollman P A Point-Charge Force Field for Molecular Mechanics Simulations of Proteins Based on Condensed-Phase Quantum Mechanical Calculations. *J Comput Chem* 2003, 24 (16), 1999–2012. [PubMed: 14531054]

- (77). Jorgensen WL; Chandrasekhar J; Madura JD; Impey RW; Klein ML Comparison of Simple Potential Functions for Simulating Liquid Water. *The Journal of Chemical Physics* 1998, 79 (2), 926–935.
- (78). VMD: Visual Molecular Dynamics. *Journal of Molecular Graphics* 1996, 14 (1), 33–38. [PubMed: 8744570]
- (79). Onufriev A; Bashford D; Case DA Exploring Protein Native States and Large Scale Conformational Changes with a Modified - Generalized Born Model. *Proteins: Structure, Function, and Bioinformatics* 2004, 55 (2), 383–394.
- (80). Bussi G; Donadio D; Parrinello M Canonical Sampling Through Velocity Rescaling. *The Journal of Chemical Physics* 2007, 126 (1), 014101. [PubMed: 17212484]
- (81). Best RB; Hummer G; Eaton WA Native Contacts Determine Protein Folding Mechanisms in Atomistic Simulations. *Proc. Natl. Acad. Sci. U.S.A* 2013, 110 (44), 17874–17879. [PubMed: 24128758]
- (82). McGibbon RT; Beauchamp KA; Harrigan MP; Klein C; Swails JM; Hernandez CX; Schwantes CR; Wang L-P; Lane TJ; Pande VS MDTraj: a Modern Open Library for the Analysis of Molecular Dynamics Trajectories. *Biophysical Journal* 2015, 109 (8), 1528–1532. [PubMed: 26488642]

Figure 1:

The probability of discovering particular states on simplified landscapes as a function of the number and length of simulations from equation 6. (A) Four states, arranged linearly, have transitions to themselves and their direct neighbors to varying degrees. States 2 and 3 are colored blue and red for visual aid. (B) A fully connected 3 state system. The probability of transitioning from state 0 to 2 is very low. (C) The probability of discovering state 2 (blue) or state 3 (red) with either a single long simulation (solid line) or many simulations consisting of 4 steps (dashed line) for the landscape in panel A. (D) The probability of discovering state 2 (blue) with either a single long simulation (solid line) or many simulations consisting of 2 steps (dashed line) for the landscape in panel B.

Figure 2:

The probability of discovering particular states on a star-shaped landscape as a function of simulation length and number of simulations from equation 6. (A) Network representation of the star-shaped landscape. Due to symmetry in the transition probabilities, a simulation started from state 0 has equal probability of reaching any of the states labeled 1, as well as any of the states labeled 2. State 0 also has a self-transition probability of 0.17, but this edge is omitted for visual clarity. (B) The probability of discovering a particular state 1 (blue) or state 2 (red) with either a single long simulation (solid line) or parallel simulations consisting of 2 steps (dashed line).

Figure 3:

An energy landscape inspired by a simple folding funnel. Conformational states are located at the vertices of the grid, where the color at this point represents the energy of that state. States can have up to 4 neighbors to transition with. Solid lines (black, red, and green) indicate the three highest flux pathways from the start to the target state, where line thickness is proportional to the flux along the particular path.

Figure 4:

The discover probabilities (the probability that a simulation set observes a particular state) on funneled landscape in Figure 3. Shown are the probabilities for four sampling strategies, a single long simulation, many parallel simulations, counts-based adaptive sampling, and the goal-oriented FAST simulations.

Figure 5:

The performance of FAST on a rugged landscape. (A) An energy landscape inspired by a folding funnel with random obstacles. Conformational states are located at the intersection of the grid lines, where the color at this point represents the energy of that state. Solid lines (black, red, and green) indicate the three highest flux pathways from the start to the target state, where line thickness is proportional to the flux along the particular path. (B) The probability that a FAST simulation set will predict a state to be in the highest-flux path from the start to the target state.

Figure 6:

An energy landscape where the preferred pathway is not the shortest distance between the start and the target state. Conformational states are located at the intersection of the grid lines, where the color at this point represents the energy of that state. Solid lines (black, red, and green) indicate the three highest flux pathways from the start to the target state, where line thickness is proportional to the flux along the particular path.

Figure 7:

The discover probabilities and predicted pathways for long and parallel simulations on the landscape in Figure 6A. (A) The probability that a long simulation discovers a particular state. (B) The probability that a long simulation will predict a state to be in the highest-flux path from the start to the target state. (C) The probability that a set of parallel simulations discovers a particular state. (D) The probability that a set of parallel simulations will predict a state to be in the highest-flux path from the start to the target state.

Figure 8:
FAST simulations navigating a large energy barrier. (A) The probability that a FAST simulation set discovers a particular state. (B) The probability that a FAST simulation set will predict a state to be in the highest-flux path from the start to the target state.

Figure 9:

A comparison of predicted pathways and estimated transition probabilities between sampling methods on a landscape with a large barrier. (A) The probability that a FAST-string simulation set will predict a state to be in the highest-flux path from the start to the target state. (B-D) The Kullbeck-Liebler divergence of each states conditional transition probabilities to the true transition probabilities. Here, a lower value indicates a lower deviation from the true underlying landscape. Compared are FAST simulations, FAST simulations followed by FASTstring, and a long simulation. Each of these are produced from equivalent aggregate simulation.

Figure 10:

An analysis of MSM estimators' performance on the landscape depicted in Figure 5A. (A) The predicted state populations for a single FAST simulation using the MLE method. (B) A comparison of the MLE-CP, transpose, normalize, MLE, and OOM estimators. Solid points are the average relative entropy for transition probabilities and equilibrium populations. Red lines are the standard deviations of these values.

Figure 11:

The largest observed fraction of native contacts as a function of aggregate simulation time for three equilibrium-based sampling methods. Simulation sets were generated from the same initial structure, which had a fraction of native contacts of 0.17 formed. Structures indicate the largest fraction of native contacts observed in a single run of FAST (red) or parallel simulations (blue) in contrast with the crystal structure (gray) (PDBID: 1LMB).

Figure 12:

Analysis of predicted folding pathways for λ -repressor using the RMSD of each residues' backbone ϕ and ψ angles to the crystal structure (PDBID: 1LMB). Folding pathways are defined as an MSMs' highest-flux path from the starting state to the state with the largest fraction of native contacts. The time evolution of each residue's backbone RMSDs are shown along the x-axis for the predicted folding pathway from two separate runs of FAST-contacts and a single set of long simulations.