

Confronting pitfalls of AI-augmented molecular dynamics using statistical physics

Cite as: J. Chem. Phys. 153, 234118 (2020); doi: 10.1063/5.0030931

Submitted: 25 September 2020 • Accepted: 29 November 2020 •

Published Online: 21 December 2020



Shashank Pant,¹  Zachary Smith,^{2,3}  Yihang Wang,^{2,3}  Emad Tajkhorshid,¹  and Pratyush Tiwary^{3,4,a)} 

AFFILIATIONS

¹NIH Center for Macromolecular Modeling and Bioinformatics, Beckman Institute for Advanced Science and Technology, Department of Biochemistry, Center for Biophysics and Quantitative Biology, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA

²Biophysics Program, University of Maryland, College Park, Maryland 20742, USA

³Institute for Physical Science and Technology, University of Maryland, College Park, Maryland 20742, USA

⁴Department of Chemistry and Biochemistry, University of Maryland, College Park, Maryland 20742, USA

^{a)}Author to whom correspondence should be addressed: ptiwary@umd.edu

ABSTRACT

Artificial intelligence (AI)-based approaches have had indubitable impact across the sciences through the ability to extract relevant information from raw data. Recently, AI has also found use in enhancing the efficiency of molecular simulations, wherein AI derived slow modes are used to accelerate the simulation in targeted ways. However, while typical fields where AI is used are characterized by a plethora of data, molecular simulations, per construction, suffer from limited sampling and thus limited data. As such, the use of AI in molecular simulations can suffer from a dangerous situation where the AI-optimization could get stuck in spurious regimes, leading to incorrect characterization of the reaction coordinate (RC) for the problem at hand. When such an incorrect RC is then used to perform additional simulations, one could start to deviate progressively from the ground truth. To deal with this problem of spurious AI-solutions, here, we report a novel and automated algorithm using ideas from statistical mechanics. It is based on the notion that a more reliable AI-solution will be one that maximizes the timescale separation between slow and fast processes. To learn this timescale separation even from limited data, we use a maximum caliber-based framework. We show the applicability of this automatic protocol for three classic benchmark problems, namely, the conformational dynamics of a model peptide, ligand-unbinding from a protein, and folding/unfolding energy landscape of the C-terminal domain of protein G. We believe that our work will lead to increased and robust use of trustworthy AI in molecular simulations of complex systems.

Published under license by AIP Publishing. <https://doi.org/10.1063/5.0030931>

I. INTRODUCTION

With the development of more accurate force fields and powerful computers, molecular dynamics (MD) has become a ubiquitous tool to study complex structural, thermodynamic, and kinetic processes of real-world systems across disciplines. However, the predictive capacity of the methodology is limited by the large timescale gap between the conformational dynamics of the complex processes of interest and the short periods accessible to it.^{1,2} This disparity is mostly attributed to the rough energy landscape typically characterized by numerous energy minima with hard to cross barriers between them,^{1,3,4} which trap the system in metastable states, leading to an incomplete sampling of the configuration space.

Comprehensive sampling of the configuration space not only provides high temporal and spatial resolutions of the complex process but also allows us to compute converged thermodynamic properties, sample physiologically relevant molecular conformations, and explore complex motions critical to biological and chemical processes such as protein folding, ligand binding, energy transfer, and countless others.^{5–13}

To overcome the limitations of timescales and accurately characterize such complex landscapes, a plethora of enhanced sampling techniques have been developed. We can broadly divide these methods into (1) tempering based, and (2) collective variable (CV) or reaction coordinate (RC) based,⁴ either of which can then also be coupled with multiple replica based exchange schemes.

In tempering based methods, the underlying landscape is sampled by either modifying the temperature and/or Hamiltonian of the system through approaches, such as temperature replica exchange, simulated annealing, and accelerated MD.^{14–21} On the other hand, CV based methods involve enhancing fluctuations along pre-defined low-dimensional modes through biased sampling approaches such as metadynamics,^{22–24} umbrella-sampling (US),²⁵ adaptive biasing force (ABF),^{26–30} and many others.^{27,31–33} Although CV-based methods can be computationally more efficient than tempering-based approaches, and given a poor choice of low-dimensional modes (a non-trivial task to intuit for complex systems), CV biasing can fail miserably.³⁴ Indeed, one could also argue that one way to make tempering methods more efficient is to select a specific part of the system, akin to a CV, which is then subjected to the tempering protocol.^{35,36}

Artificial intelligence (AI) potentially provides a systematic means to differentiate signal from noise in generic data and thus discover relevant CVs to accelerate the simulations.^{37–41} A number of such AI-based approaches have been proposed recently^{37–39,42,43} and remain the subject of extensive research. A common underlying theme in these methods is to exploit AI tools to gradually uncover the underlying effective geometry, parametrize it on-the-fly, and exploit it to bias the design of experiments with the MD simulator by emphasizing informative configuration space areas that have not been explored before. This iterative MD-AI procedure is repeated until desired sampling has been achieved. Conceptually, these approaches effectively restrain the $3N$ -dimensional space to a very small number of dimensions (typically 1 or 2), which encode all the relevant slow dynamics in the system, effectively discarding the remaining fast dynamics. Every round of AI estimates the slow modes given sampling so far, and this information is used to launch new biased rounds of simulations. Biasing along the slow modes leads to increased exploration, which can then be used in another round of AI to estimate the relevant slow modes even more accurately. The use of standard reweighting procedures can then recover unbiased thermodynamic and kinetic information from the AI-augmented MD trajectories so obtained.

However, there is a fundamental problem in such an approach. Most AI tools are designed for data-rich systems. It has been argued^{44–47} that given good quality training data and with a neural network with infinitely many parameters, the objective function for associated stochastic gradient optimization schemes is convex. However, in enhanced MD, we are, per construction, in a poorly sampled data-sparse regime, and moreover, it is impractical to use a dense network with too many parameters. The AI optimization function is therefore no longer guaranteed to be convex and can give spurious or multiple solutions for the same data set—in the same spirit as a self-driving car miscategorizing a “STOP” sign as an indication to speed up or some other action.⁴⁸ This would happen because gradient minimization got stuck in some spurious local minima or even a saddle point on the learning landscape. The slow modes thus derived would be spurious and using them as a biasing CV or RC would lead to incorrect and inefficient sampling. This could naturally lead one to derive misleading conclusions.

While the concerns stated above and the approach in this work to address them should be applicable to more general instances of AI application in molecular simulations, here, we focus on the

problem of enhanced sampling through MD-AI iterations. We report a new and computationally efficient algorithm designed to screen the spurious solutions obtained in AI-based methods. Our central hypothesis is that spurious AI solutions can be identified by tell-tale signatures in the associated dynamics, specifically through poor timescale separation between slow and fast processes. Thus, different slow mode solutions obtained from different instances of AI applied to the same data set can be ranked on the basis of how much slower the slow mode is relative to the fast modes. This difference between slow and fast mode dynamics is known as spectral gap. We would like to emphasize that the concept of largest spectral gap correlating with CV optimality is a well-founded and theoretically justified concept at the heart of many previous studies.^{49–53} However, it has not yet been applied in a computationally tractable manner to representations arising from AI frameworks used on biased datasets, as done in this work. Here, this is made feasible through the use of the “Spectral Gap Optimization of Order Parameters (SGOOP)” framework.⁵⁴ This builds a maximum caliber or path entropy⁵⁵ based model of the unbiased dynamics along different AI based representations even when the underlying observables arise from biased simulations, which then yields spectral gaps along different slow modes obtained from AI trials. We demonstrate this path entropy based screening procedure in the context of our recent iterative AI-MD scheme “Reweighted Autoencoded Variational Bayes for Enhanced sampling (RAVE).”⁴⁰ Here, we show how this automated protocol can be applied to the study of a variety of molecular problems of increasing complexity. These include conformational dynamics in a model peptide, ligand unbinding from a protein, and extensive sampling of the folding/unfolding of the C-terminal domain of protein G (GB1-C16). We believe that the presented algorithm marks a major step forward in the use of fully automated AI-enhanced MD for the study of complex biomolecular processes.

II. THEORY

A. AI can mislead

In this work, our starting point is the recent AI-based method RAVE.^{37,40,56} RAVE is an iterative MD-AI approach wherein rounds of MD for sampling are alternated with rounds of AI for learning slow modes. Specifically, RAVE begins with an initial unbiased MD trajectory comprising values of some order parameters $\mathbf{s} = (s_1, s_2, \dots, s_d)$. These could be generic variables such as dihedrals or protein–ligand distances,⁵⁷ as well as other CVs deemed to best describe the behavior of the system of interest. This trajectory is then treated with the past–future information bottleneck (PIB) framework.^{58–62} Per construction, the PIB is a low-dimensional representation with the best trade-off between minimal complexity and maximal predictive capability of the trajectory’s evolution slightly ahead in the future. RAVE uses the PIB as a computationally tractable approximation for the RC, which is traditionally considered as the definition of a slow mode.⁶³ PIB is then used in an importance sampling framework to perform the next round of biased MD. Assuming that the biased PIB is close enough to the true slow mode or modes of the system, one expects the exploration of the configuration space in this new biased round of MD to be greater than in the previous round. The biased MD itself can be performed using one of the many available biased sampling schemes.^{40,64,65}

In order to learn the PIB, RAVE uses an encoder–decoder framework. The PIB or RC χ is expressed as a linear combination of order parameters $\chi = \sum_i c_i s_i$, where the order parameters are $\mathbf{s} = (\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_d)$, c_i denotes different weights,⁵⁷ and d denotes the dimension of the order parameter space. The PIB objective function that is then minimized in every training round can be written as a difference of two mutual information,⁶⁶

$$\mathcal{L} \equiv I(\mathbf{s}, \chi) - \gamma I(\mathbf{s}_{\Delta t}, \chi), \quad (1)$$

where $I(\cdot)$ denotes the mutual information between two random variables.⁶⁶ The term $I(\mathbf{s}_{\Delta t}, \chi)$ describes the predictive power of the model, which is quantified by the amount of information shared by the information bottleneck χ and the future state of the system $\mathbf{s}_{\Delta t}$ when the information bottleneck is decoded back to the order parameter space. To optimize the objective function, the information bottleneck χ should be as informative as possible about the future state of the system, quantified through increasing $I(\mathbf{s}_{\Delta t}, \chi)$. At the same time, we seek to minimize the complexity of the low dimensional representation. Therefore, when the encoder maps the present state of the system $\mathbf{s}$ to information bottleneck χ , we aim to minimize the amount of information shared between them by decreasing $I(\mathbf{s}, \chi)$. The parameter γ is introduced to tune the trade-off between predictive power and the complexity of the model.

In Eq. (1), the encoder is a linear combination of the input coordinates, thereby keeping it interpretable and relatively robust to overfitting. The decoder is a deep artificial neural network (ANN). Due to the principle of variational inference^{40,67} wherein optimizing the decoder is guaranteed to lead to a convex optimization problem, we are not concerned with over-fitting in the decoder. Fitting the encoder, which directly leads to an interpretable RC, is our concern here. This can be best illustrated through a simple numerical example involving protein conformational dynamics, which we described in detail in the [supplementary material](#) and in [Fig. 1](#). We performed six different independently initialized trials of PIB learning using the same input trajectory for a model peptide (alanine dipeptide), each running for the same number of epochs. The RC was expressed as a linear combination of the sines and cosines of various Ramachandran dihedral angles. As shown in [Fig. 1\(b\)](#), we obtain different RCs with different trials even though they are all stopped at the same low value of the loss function (within four decimal digits of precision).




FIG. 1. Spurious AI solutions for RCs describing conformational dynamics of alanine dipeptide. (a) Molecular representation of alanine dipeptide showing relevant Ramachandran dihedral angles, ϕ and ψ . (b) Table highlights the insensitivity of the objective function toward the changes in the weights of the order parameters. Six independently initiated trials of RAVE, on the same input trajectory, resulted in different RCs. The RCs are expressed as a linear combination of sines and cosines of ϕ and ψ with coefficients/weights listed in the table.

Given the use of an interpretable linear encoder, one can see a sense of symmetry, at first glance, even in the differently looking RCs in [Fig. 1\(b\)](#). However, as will be shown later, the situation exemplified here exacerbates quickly with more complicated systems, and we expect this degeneracy to get only worse in more ambitious AI-based applications where even the encoder is non-linear^{37,41,51} and/or where one does not really know *a priori* when to stop the training.

The above numerical example demonstrates the problem at heart of what we wish to tackle in this manuscript: How does one screen through spurious solutions resulting from attempts to optimize an objective function in AI applications to molecular simulations and more broadly in chemistry and other physical sciences? The problem is especially difficult in two scenarios: first, when one does not know the ground truth against which different AI solutions could be ranked, as is expected in any application where one seeks to gain new insight. Second, as is the case in AI-augmented MD, this problem will have critical, unquantifiable ramifications in iterative learning scenarios when any such AI-derived insight is used to make new decisions and drive new rounds of biased simulations. For instance, in RAVE, we have yet another parameter that is not obvious how to select the choice of the predictive time-delay Δt in Eq. (1). As shown in Ref. 68, theoretically speaking the method is robust to the choice of this parameter as long as it is non-zero yet small enough. In practice, it can be hard to judge whether it is indeed small enough or not.

B. Path entropy model of dynamics can be used to screen AI solutions

In order to rank a set of AI-generated putative RCs, we appeal to the fundamental notion of timescale separation, which is ubiquitous across physics and chemistry through concepts such as Born–Oppenheimer approximation⁶⁹ and Michaelis–Menten principle.⁷⁰ We posit that given a basket of RC solutions generated through AI, we can rank them as being closer to the true but unknown RC if they have a higher timescale separation between slow and fast modes. Thus, a spurious AI solution should have a tell-tale signature in its dynamics, with poor separation between slow and fast modes. Indeed, one of the many definitions of an RC in chemistry is one that maximizes such a timescale separation.⁷¹ To estimate this timescale separation efficiently and rank a large number of putative AI based solutions for the true RC or PIB, here, we use the SGOOP framework,⁵⁴ which uses a maximum path entropy or caliber model^{55,72} to construct a minimal model of the dynamics along a given low-dimensional projection ([Fig. 2](#)). To construct such a model, SGOOP requires two key inputs. First, it needs the stationary probability density along any putative RC, which we directly obtain after each round of RAVE.^{40,68} Second, it needs estimates of unbiased dynamical observables, which we obtained from short MD simulations. With these two key sets of inputs, SGOOP can construct a matrix of transition rates along any putative RC. Diagonalizing this matrix gives the eigenvalues for the dynamical propagator. The spectral gap from these eigenvalues is then our estimate of the timescale separation.⁷³ While improving the quality of the dynamical observables can lead to increasingly accurate eigenvalues,⁷⁴ here, we use a computationally inexpensive dynamical observable denoted as $\langle N \rangle$ and defined as the average number of nearest neighbor transitions




FIG. 2. Flowchart illustrates our novel and computationally efficient protocol to screen AI solutions. Starting from short unbiased MD simulations, our protocol automatically screen the spurious solutions obtained in AI-based method and learns the optimal RC. In this work, we demonstrate the applicability of our protocol in the context of RAVE, and the screening of the spurious solutions is achieved by a path entropy based procedure.

per unit time along any RC. The SGOOP protocol requires a standard grid parameter (also used for histogramming), which, in all the studied systems, was set to 20. We use p_n to denote the stationary probability density along any suitably discretized putative RC at grid index n . With these inputs, the SGOOP transition probability matrix K for moving between two grid points m and n is given by^{55,73}

$$K_{mn} = \frac{\langle N \rangle}{\sum \sqrt{p_n p_m}} \sqrt{\frac{p_n}{p_m}}. \quad (2)$$

Our net product is an iterative framework that leverages the predictive power of RAVE and the fundamental notion of timescale separation of SGOOP to generate an optimal RC. The use of AI in RAVE allows one to generate several possible candidate RCs, and

by constructing a minimal path entropy based dynamical model, we efficiently screen out spurious solutions generated from AI. We would like to note that maximum path entropy does not require any additional simulations beyond those already available from RAVE; rather, it is a post processing protocol that can be employed after each set of RAVE runs to sieve-out spurious solutions. The RC identified so is then used as a biasing variable in enhanced sampling, and the biased trajectory itself is fed back to the AI module to further optimize the RC. The iteration between this framework and sampling continues until multiple transitions between different intermediate states are achieved. We also apply this framework to cleanly select the best choice of predictive time-delay (Δt) in Eq. (1)—the optimal predictive time delay in our model for PIB is the one that achieves the highest timescale separation.

III. RESULTS

In Sec. II B, we described a path entropy and timescale separation based paradigm to capture spurious solutions in AI-enhanced MD. In this section, we illustrate the effectiveness of our framework through three generically relevant biophysical examples of increasing complexity. Specifically, we consider (A) conformational dynamics of a model peptide in vacuum, (B) dissociation of a millimolar-affinity ligand from FK506-binding protein (FKBP) protein, and (C) folding of the GB1-C16 peptide. All simulations are done at an all-atom resolution, including explicit water in (B) and (C). In all three systems, starting with an initial unbiased MD trajectory comprising generic order parameters s , we perform iterative rounds of RAVE followed by biased enhanced sampling using SGOOP to screen RC candidates generated in RAVE and to select the optimal time delay Δt in Eq. (1). Apart from the starting choice of order parameters that are kept quite generic (Table S1), all steps are carried out with minimal use of human intuition. To display the versatility of our framework, we combined it with two different enhanced sampling algorithms.⁷⁵ In systems (A) and (B), we employ static biases to further enhance the conformational sampling of the model peptide and ligand dissociation along the reaction path.⁷⁶ These static biases were directly obtained by inverting the probability distribution learnt during RAVE.^{40,56} In system (C), we employ time-dependent biasing through well-tempered metadynamics^{24,65} to capture folding of the GB1-C16 peptide. All the simulations were performed with GROMACS version 5.0⁷⁷ patched with PLUMED version 2.4.2.^{78,79}

A. Conformational dynamics of alanine dipeptide

The first system we consider here is the well-studied case of alanine dipeptide in vacuum. It can exist in multiple conformations separated by barriers and commonly characterized by differing values of its backbone dihedral angles ϕ and ψ [Fig. 1(b)]. Enabled by the small size of the system, we performed three independent simulations, each 2 μ s long. The corresponding trajectories along with the dihedral angles ϕ and ψ are provided in Fig. S1 of the [supplementary material](#). In line with standard practice,^{80,81} the sines and cosines of these two dihedral angles provide natural input order parameters (OPs) $s = (\cos \phi, \sin \phi, \cos \psi, \sin \psi)$ for RAVE, which then learns the optimal RC χ as a linear combination of these four. In the three independent trajectories, even with such long simulation times, we

capture only 1, 2, and 4 transitions between the axial and the equatorial conformations of the dipeptide. Using such input trajectories with a different number of transitions helps us ascertain robustness of the protocol developed here. Each trajectory was used to perform RAVE with 11 different choices of the predictive time delay Δt in Eq. (1) ranging from 0 ps to 40 ps. Furthermore, ten different trials were performed for each Δt value corresponding to different input trajectories. This amounts to a total of 330 RAVE calculations, with 110 for each input trajectory. Each trial was stopped after the same training time, and the loss function value after the training and the RC so-obtained were recorded.

As hinted in Sec. I, we obtain very different RCs for the different Δt values and for different independent trials. Furthermore, different trials that were stopped at a similar loss function value gave different RCs and spectral gaps (Figs. 3 and S2). However, our protocol of using spectral gaps to rank these different solutions works well in screening out the RC. In Figs. 3(a)–3(c), we demonstrate the noisy correlation that we find between the loss function value and the spectral gap for all three input trajectories. In the [supplementary material](#) [Fig. S2(b)], we provide an illustrative figure for one particular trajectory showing how the same loss function value results in RCs with different free energy profiles and that the one with the highest spectral gap stands out with the most clearly demarcated metastable states. Similarly, the spectral gap captures the most optimal RC not just from the set of multiple trials at each time delay, but it can also be used to select the optimal

time delay itself [Fig. 3(d), red]. In the subsequent calculations, an optimal time delay of 8 ps, corresponding to the maximum spectral gap, was employed. Irrespective of the choice of input trajectory, we find that the optimal RC shows higher weights for ϕ (as compared to ψ) (Table I), in line with previous studies that highlighted ϕ to be a more important degree of freedom than ψ .^{40,82} Using the RC corresponding to the *ntrans* = 4 and its probability distribution as a fixed bias,⁴⁰ we then explored the conformational space of the peptide. The two-dimensional free-energy landscape along the dihedrals ϕ and ψ was able to capture axial and equatorial conformations of the peptide in only 20 ns of biased simulation [Fig. 3(e)]. This is in excellent agreement with previously published studies for this system.^{24,40} However, biased simulations with the RAVE-alone RC result in poorer sampling of the configuration space relative to biased simulations using the RC further screened with SGOOP, as shown by the lesser number of transitions between the energy basins in Fig. S3.

B. Unbinding of millimolar-affinity ligand from FKBP

In the second example, we applied our framework to a well-studied problem of dissociation of 4-hydroxy-2-butanone (BUT), a millimolar affinity ligand, from the FKBP protein [Fig. 4(a)]. Force-field parametrization^{83–85} and other MD details are provided in the [supplementary material](#). Here, our objective was to use RAVE to learn the most optimal RC on-the-fly as well as the absolute




FIG. 3. Capturing the spurious AI solutions in alanine dipeptide. Spectral gap and loss function values were calculated for each of the three unbiased trajectories at multiple time delays Δt between 0 ps and 40 ps, indicated using circles of different colors in the bottom right-hand side. (a)–(c) show noisy correlation between the loss and the spectral gap for the number of transitions *ntrans* equaling 1, 2, and 4 respectively. Different circles denote different independent trials, with color denoting Δt . For visual clarity, for every *ntrans*, we have plotted a mean-free version of the loss function value by subtracting out the average of all losses. (d) Maximum spectral gap (out of ten different trials of RAVE) vs Δt was plotted for three different unbiased trajectories. Optimal time delay of 8 ps was employed in subsequent calculations. (e) Free energy surface (FES) along the two dihedrals ϕ and ψ obtained from the 20 ns-long simulation in the presence of static bias. Energy contours are shown at every 4 kJ/mol.

TABLE I. Optimal weights of OPs obtained through a combination of RAVE and SGOOP.

Trajectory	Transitions	$\cos \phi$	$\sin \phi$	$\cos \psi$	$\sin \psi$
1	1	-0.71	-0.68	0.09	-0.15
2	2	0.35	0.81	0.32	0.34
3	4	0.63	-0.65	0.03	0.42

binding free energy of this protein–ligand complex. This is a difficult and important problem for which many useful methods have already been employed with varying levels of success.⁸⁶ At this stage at least, our intention is not to compete with these other existing methods, but instead validate that our framework works for a well-studied benchmark problem. We begin by performing four independent MD simulations of FKBP in its ligand-bound form (PDB:1D7J).⁸⁷ The MD simulations were stopped when the ligand unbound, specifically when it was 2 Å away from the binding pocket [Fig. 4(b)]. All trajectories were expressed in terms of eight OPs representing various distances between the center of mass (COM) of the ligand and the COM of the residues in the binding pocket (Table S1), which comprise a natural choice for the process of ligand unbinding from the protein

and have been employed in previous studies.^{57,88} We combined the results of the four independent MD trajectories to perform RAVE with 11 different choices of predictive time delay ranging from 0 ps to 40 ps. At each Δt , 10 different trials were performed, resulting in a total of 110 RAVE calculations. Each trial was stopped after the same training time, and the loss function value after the training as well as the RC so-obtained were recorded. Different RCs were screened using a path entropy based model, as discussed in Sec. II and done for alanine dipeptide. We again find noisy correlation between the loss function values and the spectral gap [Fig. 4(c)] for the case of $\Delta t = 40$ ps (additional plots are given in Fig. S4). The same value of the loss function gives rise to very different values of the spectral gap and the RC [Fig. 4(c)]. Furthermore, the spectral gap not only captures the most optimal RC but also is able to select the most optimal time delay [Fig. 4(d)]. By using this RC and its probability distribution as a fixed bias,⁴⁰ we then performed 800 ns of biased simulations starting from the bound pose, but allowing the ligand to re-associate [Figs. S5(a) and S5(b)]. Through this, we then calculated the absolute binding affinity of the protein–ligand complex to be 6.6 kJ/mol [Fig. 4(e)], in good agreement with the values reported through metadynamics.⁵⁷ Interestingly, the binding affinity of the protein–ligand complex was also in good agreement with the values reported through extended unbiased simulations by Pan *et al.*,⁸⁹ although the ligand was parameterized with the generalized amber




FIG. 4. Unbinding of 4-hydroxy-2-butanone (BUT) from FKBP. (a) Molecular image of the bound FKBP/BUT protein–ligand complex, with binding pocket residues highlighted. The distances to the residues were used as different OPs detailed in Table S1. (b) Time evolution of the distance between the center of mass (COM) of the bound ligand and COM of residue W59. It is to be noted that in order to avoid entropy dominant process, only the ligand-bound trajectories were considered in our protocol. (c) Spectral gap and loss (at time delay of 40 ps) for ten different trials were calculated after combining all four independent trajectories at multiple time delay Δt between 0 ps and 40 ps, indicated using circles. Rest of the time delays are shown in Fig. S1. For visual clarity, at each iteration, we have plotted a mean-free version of the loss function value by subtracting out the average of all losses. (d) Plot of the maximum spectral gap (out of ten different trials of RAVE) vs time delay (Δt). (e) Absolute binding free energy G in kJ/mol of the FKBP/BUT system as a function of simulation time with static external bias. The black dotted and orange dotted line show the reference value reported through metadynamics⁵⁷ and long unbiased MD simulations performed on ANTON, respectively.⁸⁹ The shaded region shows the free energy estimate from long unbiased MD simulations performed on Anton including the ± 2.092 kJ error reported.⁸⁹ (f) A visual depiction of the OP weights.

force field (GAFF).⁸⁸ It is worth pointing out that the ANTON simulations took 39 μ s, while we obtained converged estimates in around 800 ns, reflecting roughly a factor of 48 speed-up with minimal use of prior human intuition.

The use of a linear encoder in RAVE allows us to directly interpret the weights of the different OPs in the RC [Fig. 4(f)]. The highest weight corresponds to the OP d_5 , which is the ligand separation from residue I56. This residue forms direct interactions with the bound ligand in the crystal structure. Interestingly, previous studies⁸⁷ have highlighted the importance of I56 as it forms hydrogen bonding interactions with the carbonyl group of the bound ligand; our algorithm also captured it as the most significant OP. Followed by this highest weight component, the second and third highest components are for d_1 and d_2 , denoting distances from the residues V55 and W59, respectively. These are roughly equal in magnitude, reflecting that the ligand moves closer to V55 and W59 as it moved away from I56.

C. Folding/unfolding dynamics of GB1 peptide

Finally, we tested our method on the folding/unfolding dynamics of GB1-C16, which is known to adopt a β -hairpin structure.^{90–94} Force-field parametrization and other MD details are provided in the [supplementary material](#). The free-energy landscape of this peptide has been extensively explored by replica-exchange MD simulations and clustering based methods.^{90,94,95} These studies reported the presence of multiple intermediate conformations by projecting the simulation data along multiple OPs, such as radius of gyration (Rg), root-mean-squared deviation (RMSD), fraction of native contacts (NC), and native state hydrogen bonds (NHB). These OPs on their own were not able to distinguish between intermediate conformations with proper energy barriers. However, using a combination of

these OPs as the input in advanced slow mode analysis methods such as TICA^{52,96,97} recovers a more superior two-dimensional description.⁹⁰ That work, however, used more than 12 μ s of enhanced sampling, specifically replica exchange MD trajectories, for this purpose. Here, instead, we use just 1.6 μ s of unbiased trajectories as our starting point. From this point onwards, using the same OPs as in Ref. 90, our work provides a semi-automated solution in deriving an optimal two-dimensional RC for GB1-C16, which is capable of resolving the intermediate conformations. Here, in contrast to the previous two examples, we use well-tempered metadynamics^{24,65} simulations as the choice of the enhanced sampling engine coupled with RAVE.

We start by performing four independent 400 ns of unbiased MD simulations of the peptide in explicit solvent. All the simulated systems were observed to be fairly stable when projected along a library of OPs comprising RMSD, NC, Rg, and NHB, with their detailed construction described in the [supplementary material](#) (see Figs. 5 and S6). All the unbiased trajectories were mixed and fed into RAVE for learning the RC. We performed 10 different trials of RAVE for different time delays Δt ranging from 0 ps to 20 ps, which amounts to a total of 110 RAVE calculations. Different putative RCs learnt from RAVE were screened using the path entropy based model, as discussed in the Sec. II and as done for the other two systems. Similar to the previous systems, we find a noisy correlation between the loss function value and the spectral gap (Fig. S7). The most optimal RC was selected for biased simulations using well-tempered metadynamics [Figs. S7(a) and S7(b)]. Based on the maximum spectral gap, we chose $\Delta t = 8$ ps for the next round of the 50 ns-long metadynamics simulation. We then alternatively iterate between the rounds of learning improved RC, using our framework, and running metadynamics using the optimal RC in every iteration. After two iterations, we did not find any further improvement in




FIG. 5. Dynamics of GB1-C16 captured from unbiased MD. One of the four representative trajectory of the peptide in explicit solvent is projected along different order parameters: (a) number of hydrogen bonds (NHB), (b) native contacts (NC), (c) radius of gyration (Rg), and (d) root-mean square deviation (RMSD). (e) Molecular image of the GB1-C16. Native backbone hydrogen bonds are highlighted with green lines.




FIG. 6. Free-energy landscape and OPs contribution. (a) Contribution of the different OPs to the two-dimensional RC χ . The two components χ_1 and χ_2 are shown in blue and red bars, respectively. (b) A highly rugged two-dimensional free-energy landscape of GB1-C16 folding/unfolding. We were able to capture multiple states corresponding to the folded (IS1), unfolded (IS4), and intermediate states (IS2 and IS3). Interestingly, it is only by projecting the free energy as a function of the two RCs that we were able to capture a partially helical state (IS3), which, otherwise, was not easy to distinguish solely using traditional OPs. Representative snapshots of the captured structures are shown in the bottom panel, and their locations on the energy landscape shown in (b).

sampling with this one-dimensional RC, which we call χ_1 . With a 1d metadynamics, we were unable to attain back and forth transitions between different metastable states, suggesting the presence of missing/orthogonal degrees of freedom not encapsulated by χ_1 . In order to learn these other degrees of freedom through the second component of the RC, which we call χ_2 , we used the protocol from Ref. 98. For practical purposes, this corresponds to ignoring the already learnt χ_1 and treating the biased trajectory without any consideration of the bias along χ_1 . We would like to note that in the previous study,⁷³ we have extended the scope of SGOOP by employing the notion of conditional probability factorization where known features are effectively washed out to learn additional features of the underlying energy landscape. This is what we have used for RAVE as well in the current work. In principle, RAVE could be directly used to output a two-dimensional or even higher-dimensional RC, but this protocol ensures that we gradually ensure the RC dimensionality only when a lower dimension is found insufficient for sampling. We then performed 50 ns long 2D metadynamics simulations [Figs. S7(e) and S7(f)], which were used to train χ_2 . The most optimal two-dimensional RC obtained after three iterations of training χ_2 is detailed in [Fig. 6(a)]. The backbone heavy atom RMSD contributes the most toward the construction of the slowest dimension χ_1 , whereas Rg contributed more toward the second slowest dimension.

The two-dimensional RC is then used in longer well-tempered metadynamics simulations to facilitate movement between different metastable states (see Video 1 of the [supplementary material](#)) and to obtain converged free energy surfaces. We performed 1.2 μ s-long metadynamics simulations at 300 K, starting from the crystal structure (Fig. S8). The two-dimensional metadynamics simulations were performed with an initial hill height of 0.5 kJ, bias factor = 10, Gaussian widths of 0.03 for both χ_1 and χ_2 , and bias added every 4 ps. Additional restraint potential was applied along the RMSD order parameter preventing very high values from being attained (see details in the [supplementary material](#)). In principle, this step is not necessary as the simulation would eventually return back to low RMSD states, but in practice, due to the entropic nature of the

high RMSD states, such a restraint significantly helps with computational efficiency. Figure 6(b) shows the 2D free energy landscape as a function of the two RC components at 300 K. We find that the system shows multiple energy basins corresponding to the different stable and metastable intermediates. Interestingly, we captured a helical conformation of this peptide, which was not easy to distinguish by using a combination of conventional OPs such as RMSD, Rg, and contact map.⁹⁹ For example, previous metadynamics-based studies employed Rg and native hydrogen bonds to accelerate the folding process, but they were not able to clearly demarcate distinct conformational states with energy barriers.^{100,101} Interestingly, the two-dimensional free-energy landscapes when projected along the pair of OPs yield results consistent with the previous studies and suggest the presence of two metastable states (Fig. S9).

IV. CONCLUSION

To conclude, we have introduced a new approach to sieve out the spurious solutions from AI-augmented enhanced sampling simulations.^{37,38} AI-based approaches have had an indisputable impact across sciences, including their use in enhancing the efficiency of molecular simulations. However, when these AI-based approaches are applied to a data sparse regime, it can lead to spurious or multiple solutions. This would happen because gradient minimization can get stuck in some spurious local minima or even saddle points on the learning landscape, leading to misleading use of AI.

To deal with this issue of trustworthiness of AI in molecular simulations, we report a novel automated algorithm aided by ideas from statistical physics.¹⁰² Our algorithm is based on the simple but powerful notion that a more reliable AI solution will be the one that maximizes the timescale separation between slow and fast processes. This fundamental notion of timescale separation was implemented on the basis of maximum caliber- or path-entropy-based method, SGOOP.^{54,55} We would like to emphasize that our approach and spectral gap based optimization, in general,⁵⁴ might have as of yet unexplored connections with the Variational Approach for Markov Processes (VAMP).¹⁰³ The framework

developed here should be applicable to many recent methods (Ref. 37 and references therein), which involve iterating between MD and AI for sampling and learning, respectively. Here, we demonstrate its usefulness through our recent integrated AI-MD algorithm RAVE.⁴⁰ We illustrate the applicability of our algorithm through three illustrative examples, including the complex problem of capturing the energetic landscape of GB1 peptide folding in all-atom simulations. In this last case, we started from a library of four order parameters that are generic for folding/unfolding processes and demonstrated how to semi-automatically learn a two-dimensional RC, which we then used in a well-tempered metadynamics protocol to obtain folding/unfolding trajectories. This directly allows us to gain atomic level insight into different metastable states relevant to the folding/unfolding process. We thus believe that our method marks a useful and much needed step forward in increasing the utility of machine learning and AI-based methods in the context of enhanced sampling, and one can expect that such an approach could be applicable to molecular simulations in general, although this is purely speculative at this point and remains to be verified.

SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for simulation details, neural network architecture, unbiased/biased MD trajectories, and other numerical details.^{104–107}

AUTHORS' CONTRIBUTIONS

Z.S. and Y.W. contributed equally to this work.

ACKNOWLEDGMENTS

We thank the Donors of the American Chemical Society Petroleum Research Fund (Grant No. PRF60512-DNI6; P.T.), the National Institutes of Health under Award No. P41-GM104601 (E.T.), the NCI-UMD Partnership for Integrative Cancer Research (Y.W.), the COMBINE fellowship (Grant No. DGE-1632976; Z.S.), and the Beckman Institute Graduate Fellowship (S.P.) for financial support. This work used the Extreme Science and Engineering Discovery Environment (XSEDE) Bridges through allocation Grant No. CHE180027P, which is supported by the National Science Foundation Grant No. ACI-1548562 (P.T.). We also thank UMD's Deepthought2 and MARCC's Bluecrab HPC clusters for computing resources. We thank Navjeet Ahlawat and Jagannath Mondal for help with setting up the GB1 peptide system for MD simulation. We also thank Andrew Ferguson for useful discussion and João Marcelo Lamim Ribeiro for the careful reading of the manuscript.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request (RAVE code is available at <https://github.com/tiwarylab/RAVE> and SGOOP code can be accessed at <https://github.com/tiwarylab/SGOOP>).

REFERENCES

- ¹R. C. Bernardi, M. C. R. Melo, and K. Schulten, "Enhanced sampling techniques in molecular dynamics simulations of biological systems," *Biochim. Biophys. Acta* **1850**, 872–877 (2015).
- ²M. Karplus and G. A. Petsko, "Molecular dynamics simulations in biology," *Nature* **347**, 631–639 (1990).
- ³M. A. Rohrdanz, W. Zheng, and C. Clementi, "Discovering mountain passes via torchlight: Methods for the definition of reaction coordinates and pathways in complex macromolecular reactions," *Annu. Rev. Phys. Chem.* **64**, 295–316 (2013).
- ⁴C. Abrams and G. Bussi, "Enhanced sampling in molecular dynamics using metadynamics, replica-exchange, and temperature-acceleration," *Entropy* **16**, 163–199 (2014).
- ⁵B. Hashemian, D. Millán, and M. Arroyo, "Modeling and enhanced sampling of molecular systems with smooth and nonlinear data-driven collective variables," *J. Chem. Phys.* **139**, 214101 (2013).
- ⁶J. N. Onuchic and P. G. Wolynes, "Theory of protein folding," *Curr. Opin. Struct. Biol.* **14**, 70–75 (2004).
- ⁷K. A. Dill, S. B. Ozkan, M. S. Shell, and T. R. Weikl, "The protein folding problem," *Annu. Rev. Biophys.* **37**, 289–316 (2008).
- ⁸P. Tiwary, V. Limongelli, M. Salvalaglio, and M. Parrinello, "Kinetics of protein–ligand unbinding: Predicting pathways, rates, and rate-limiting steps," *Proc. Natl. Acad. Sci. U. S. A.* **112**, E386–E391 (2015).
- ⁹P. Tiwary, J. Mondal, and B. J. Berne, "How and when does an anticancer drug leave its binding site?," *Sci. Adv.* **3**, e1700014 (2017).
- ¹⁰M. Moradi and E. Tajkhorshid, "Mechanistic picture for conformational transition of a membrane transporter at atomic resolution," *Proc. Natl. Acad. Sci. U. S. A.* **110**, 18916–18921 (2013).
- ¹¹M. Moradi and E. Tajkhorshid, "Computational recipe for efficient description of large-scale conformational changes in biomolecular systems," *J. Chem. Theory Comput.* **10**, 2866–2880 (2014).
- ¹²M. Moradi, G. Enkavi, and E. Tajkhorshid, "Atomic-level characterization of transport cycle thermodynamics in the glycerol-3-phosphate:phosphate transporter," *Nat. Commun.* **6**, 8393 (2015).
- ¹³S. Pant and E. Tajkhorshid, "Microscopic characterization of GRP1 PH domain interaction with anionic membranes," *J. Comput. Chem.* **41**, 489–499 (2019).
- ¹⁴S. P. Brooks and B. J. Morgan, "Optimization using simulated annealing," *J. R. Stat. Soc.: D* **44**, 241–257 (1995).
- ¹⁵U. H. Hansmann, "Parallel tempering algorithm for conformational studies of biological molecules," *Chem. Phys. Lett.* **281**, 140–150 (1997).
- ¹⁶Y. Sugita and Y. Okamoto, "Replica-exchange molecular dynamics method for protein folding," *Chem. Phys. Lett.* **314**, 141–151 (1999).
- ¹⁷Y. Sugita, A. Kitao, and Y. Okamoto, "Multidimensional replica-exchange method for free-energy calculations," *J. Chem. Phys.* **113**, 6042–6051 (2000).
- ¹⁸A. Mitsutake, Y. Sugita, and Y. Okamoto, "Generalized-ensemble algorithms for molecular simulations of biopolymers," *Biopolymers* **60**, 96–123 (2001).
- ¹⁹H. Fukunishi, O. Watanabe, and S. Takada, "On the Hamiltonian replica exchange method for efficient sampling of biomolecular systems: Application to protein structure prediction," *J. Chem. Phys.* **116**, 9058–9067 (2002).
- ²⁰L. Maragliano and E. Vanden-Eijnden, "A temperature accelerated method for sampling free energy and determining reaction pathways in rare events simulations," *Chem. Phys. Lett.* **426**, 168–175 (2006).
- ²¹Y. Miao, V. A. Feher, and J. A. McCammon, "Gaussian accelerated molecular dynamics: Unconstrained enhanced sampling and free energy calculation," *J. Chem. Theory Comput.* **11**, 3584–3595 (2015).
- ²²A. Laio and M. Parrinello, "Escaping free-energy minima," *Proc. Natl. Acad. Sci. U. S. A.* **99**, 12562–12566 (2002).
- ²³A. Laio and F. L. Gervasio, "Metadynamics: A method to simulate rare events and reconstruct the free energy in biophysics, chemistry and material science," *Rep. Progr. Phys.* **71**, 126601 (2008).
- ²⁴A. Barducci, G. Bussi, and M. Parrinello, "Well-tempered metadynamics: A smoothly converging and tunable free-energy method," *Phys. Rev. Lett.* **100**, 020603 (2008).
- ²⁵G. M. Torrie and J. P. Valleau, "Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling," *J. Chem. Phys.* **23**, 187–199 (1977).

- ²⁶T. Lelièvre, M. Rousset, and G. Stoltz, "Computation of free energy profiles with parallel adaptive dynamics," *J. Chem. Phys.* **126**, 134111 (2007).
- ²⁷E. Darve, D. Rodríguez-Gómez, and A. Pohorille, "Adaptive biasing force method for scalar and vector free energy calculations," *J. Chem. Phys.* **128**, 144120 (2008).
- ²⁸J. Comer, J. C. Gumbart, J. Hénin, T. Lelièvre, A. Pohorille, and C. Chipot, "The adaptive biasing force method: Everything you always wanted to know but were afraid to ask," *J. Phys. Chem. B* **119**, 1129–1151 (2015).
- ²⁹H. Fu, X. Shao, C. Chipot, and W. Cai, "Extended adaptive biasing force algorithm. An on-the-fly implementation for accurate free-energy calculations," *J. Chem. Theory Comput.* **12**, 3506–3513 (2016).
- ³⁰A. Lesage, T. Lelièvre, G. Stoltz, and J. Hénin, "Smoothed biasing forces yield unbiased free energies with the extended-system adaptive biasing force method," *J. Phys. Chem. B* **121**, 3676–3685 (2017).
- ³¹J. B. Abrams and M. E. Tuckerman, "Efficient and direct generation of multidimensional free energy surfaces via adiabatic dynamics without coordinate transformations," *J. Phys. Chem. B* **112**, 15742–15757 (2008).
- ³²J. G. Kirkwood, "Statistical mechanics of fluid mixtures," *J. Chem. Phys.* **3**, 300–313 (1935).
- ³³W. K. den Otter and W. J. Briels, "The calculation of free-energy differences by constrained molecular-dynamics simulations," *J. Chem. Phys.* **109**, 4139–4146 (1998).
- ³⁴A. Hazel, C. Chipot, and J. C. Gumbart, "Thermodynamics of deca-alanine folding in water," *J. Chem. Theory Comput.* **10**, 2836–2844 (2014).
- ³⁵P. Liu, B. Kim, R. A. Friesner, and B. Berne, "Replica exchange with solute tempering: A method for sampling biological systems in explicit water," *Proc. Natl. Acad. Sci. U. S. A.* **102**, 13749–13754 (2005).
- ³⁶X. Huang, M. Hagen, B. Kim, R. A. Friesner, R. Zhou, and B. J. Berne, "Replica exchange with solute tempering: Efficiency in large scale systems," *J. Phys. Chem. B* **111**, 5405–5410 (2007).
- ³⁷Y. Wang, J. M. L. Ribeiro, and P. Tiwary, "Machine learning approaches for analyzing and enhancing molecular dynamics simulations," *Curr. Opin. Struct. Biol.* **61**, 139–145 (2020).
- ³⁸F. Noé, G. De Fabritiis, and C. Clementi, "Machine learning for protein folding and dynamics," *Curr. Opin. Struct. Biol.* **60**, 77–84 (2020).
- ³⁹F. Noé, S. Olsson, J. Köhler, and H. Wu, "Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning," *Science* **365**, eaaw1147 (2019).
- ⁴⁰Y. Wang, J. M. L. Ribeiro, and P. Tiwary, "Past–future information bottleneck for sampling molecular reaction coordinate simultaneously with thermodynamics and kinetics," *Nat. Commun.* **10**, 3573 (2019).
- ⁴¹W. Chen, A. R. Tan, and A. L. Ferguson, "Collective variable discovery and enhanced sampling using autoencoders: Innovations in network architecture and error function design," *J. Comput. Chem.* **149**, 072312 (2018).
- ⁴²M. E. Tuckerman, "Machine learning transforms how microstates are sampled," *Science* **365**, 982–983 (2019).
- ⁴³S.-L. J. Lahey and C. N. Rowley, "Simulating protein–ligand binding with neural network potentials," *Chem. Sci.* **11**, 2362–2368 (2020).
- ⁴⁴G. Rotskoff and E. Vanden-Eijnden, "Parameters as interacting particles: Long time convergence and asymptotic error scaling of neural networks," in *Advances in Neural Information Processing Systems* (Curran Associates, 2018), pp. 7146–7155.
- ⁴⁵G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Math. Control, Signals, Syst.* **5**, 455 (1992).
- ⁴⁶A. R. Barron, "Universal approximation bounds for superpositions of a sigmoidal function," *IEEE Trans. Inf. Theory* **39**, 930–945 (1993).
- ⁴⁷F. Bach, "Breaking the curse of dimensionality with convex neural networks," *J. Mach. Learn. Res.* **18**, 629–681 (2017).
- ⁴⁸I. Evtimov, K. Eykholt, E. Fernandes, T. Kohno, B. Li, A. Prakash, A. Rahmati, and D. Song, "Robust physical-world attacks on deep learning models," *arXiv:1707.08945* (2017).
- ⁴⁹F. Noé, I. Horenko, C. Schütte, and J. C. Smith, "Hierarchical analysis of conformational dynamics in biomolecules: Transition networks of metastable states," *J. Chem. Phys.* **126**, 155102 (2007).
- ⁵⁰M. A. Rohrdanz, W. Zheng, M. Maggioni, and C. Clementi, "Determination of reaction coordinates via locally scaled diffusion map," *J. Chem. Phys.* **134**, 124116 (2011).
- ⁵¹F. Noé and F. Nüske, "A variational approach to modeling slow processes in stochastic dynamical systems," *Multiscale Model. Simul.* **11**, 635–655 (2013).
- ⁵²G. Pérez-Hernández, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, "Identification of slow molecular order parameters for Markov model construction," *J. Chem. Phys.* **139**, 015102 (2013).
- ⁵³Q. Li, F. Dietrich, E. M. Boltt, and I. G. Kevrekidis, "Extended dynamic mode decomposition with dictionary learning: A data-driven adaptive spectral decomposition of the Koopman operator," *Chaos* **27**, 103111 (2017).
- ⁵⁴P. Tiwary and B. Berne, "Spectral gap optimization of order parameters for sampling complex molecular systems," *Proc. Natl. Acad. Sci. U. S. A.* **113**, 2839–2844 (2016).
- ⁵⁵K. Ghosh, P. D. Dixit, L. Agazzino, and K. A. Dill, "The maximum caliber variational principle for nonequilibrium," *Annu. Rev. Phys. Chem.* **71**, 213–238 (2020).
- ⁵⁶J. M. L. Ribeiro, P. Bravo, Y. Wang, and P. Tiwary, "Reweighted autoencoded variational Bayes for enhanced sampling (RAVE)," *J. Chem. Phys.* **149**, 072301 (2018).
- ⁵⁷P. Ravindra, Z. Smith, and P. Tiwary, "Automatic mutual information noise omission (AMINO): Generating order parameters for molecular systems," *Mol. Syst. Des. Eng.* **5**, 339 (2020).
- ⁵⁸N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," *arXiv:physics/0004057* (2000).
- ⁵⁹S. E. Palmer, O. Marre, M. J. Berry, and W. Bialek, "Predictive information in a sensory population," *Proc. Natl. Acad. Sci. U. S. A.* **112**, 6908–6913 (2015).
- ⁶⁰G. J. Berman, W. Bialek, and J. W. Shaevitz, "Predictability and hierarchy in *Drosophila* behavior," *Proc. Natl. Acad. Sci. U. S. A.* **113**, 11943–11948 (2016).
- ⁶¹A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy, "Deep variational information bottleneck," *arXiv:1612.00410* (2016).
- ⁶²S. Still, "Information bottleneck approach to predictive inference," *Entropy* **16**, 968–989 (2014).
- ⁶³S. V. Krivov, "On reaction coordinate optimality," *J. Chem. Theory Comput.* **9**, 135–146 (2013).
- ⁶⁴Z. Smith, P. Ravindra, Y. Wang, R. Cooley, and P. Tiwary, "Discovering loop conformational flexibility in T4 lysozyme mutants through Artificial Intelligence aided Molecular Dynamics," *J. Phys. Chem. B* **124**, 8221–8229 (2020).
- ⁶⁵O. Valsson, P. Tiwary, and M. Parrinello, "Enhancing important fluctuations: Rare events and metadynamics from a conceptual viewpoint," *Annu. Rev. Phys. Chem.* **67**, 159–184 (2016).
- ⁶⁶T. M. Cover and J. A. Thomas, *Elements of Information Theory* (John Wiley & Sons, 2012).
- ⁶⁷I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016), <http://www.deeplearningbook.org>.
- ⁶⁸Y. Wang and P. Tiwary, "Understanding the role of predictive time delay and biased propagator in RAVE," *J. Chem. Phys.* **152**, 144102 (2020).
- ⁶⁹I. N. Levine, D. H. Busch, and H. Shull, *Quantum Chemistry* (Pearson Prentice Hall Upper, Saddle River, NJ, 2009), Vol. 6.
- ⁷⁰D. L. Nelson, A. L. Lehninger, and M. M. Cox, *Lehninger Principles of Biochemistry* (Macmillan, 2008).
- ⁷¹D. G. Truhlar and B. C. Garrett, "Variational transition state theory," *Annu. Rev. Phys. Chem.* **35**, 159–189 (1984).
- ⁷²P. D. Dixit and K. A. Dill, "Caliber corrected Markov modeling (C_2M_2): Correcting equilibrium Markov models," *J. Chem. Theory Comput.* **14**, 1111–1119 (2018).
- ⁷³Z. Smith, D. Pramanik, S.-T. Tsai, and P. Tiwary, "Multi-dimensional spectral gap optimization of order parameters (SGOOP) through conditional probability factorization," *J. Chem. Phys.* **149**, 234105 (2018).
- ⁷⁴D. Meral, D. Provasi, and M. Filizola, "An efficient strategy to estimate thermodynamics and kinetics of G protein-coupled receptor activation using metadynamics and maximum caliber," *J. Chem. Phys.* **149**, 224101 (2018).
- ⁷⁵P. Tiwary and A. van de Walle, *Multiscale Materials Modeling for Nanomechanics* (Springer, 2016), pp. 195–221.

- ⁷⁶J. Debnath and M. Parrinello, "Gaussian mixture based enhanced sampling for statics and dynamics," *J. Phys. Chem. Lett.* **11**, 5076 (2020).
- ⁷⁷M. J. Abraham, T. Murtola, R. Schulz, S. Páll, J. C. Smith, B. Hess, and E. Lindahl, "GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers," *SoftwareX* **1-2**, 19–25 (2015).
- ⁷⁸G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi, "PLUMED 2: New feathers for an old bird," *Comput. Phys. Commun.* **185**, 604–613 (2014).
- ⁷⁹M. Bonomi, G. Bussi, C. Camilloni, G. A. Tribello, P. Banáš, A. Barducci, M. Bernetti, P. G. Bolhuis, S. Bottaro, D. Branduardi *et al.*, "Promoting transparency and reproducibility in enhanced molecular simulations," *Nat. Methods* **16**, 670–673 (2019).
- ⁸⁰Y. Mu, P. H. Nguyen, and G. Stock, "Energy landscape of a small peptide revealed by dihedral angle principal component analysis," *Proteins* **58**, 45–52 (2005).
- ⁸¹A. Altis, P. H. Nguyen, R. Hegger, and G. Stock, "Dihedral angle principal component analysis of molecular dynamics simulations," *J. Chem. Phys.* **126**, 244111 (2007).
- ⁸²M. Salvalaglio, P. Tiwary, and M. Parrinello, "Assessing the reliability of the dynamics reconstructed from metadynamics," *J. Chem. Theory Comput.* **10**, 1420–1425 (2014).
- ⁸³V. Hornak, R. Abel, A. Okur, B. Strockbine, A. Roitberg, and C. Simmerling, "Comparison of multiple Amber force fields and development of improved protein backbone parameters," *Proteins* **65**, 712–725 (2006).
- ⁸⁴R. B. Best and G. Hummer, "Optimized molecular dynamics force fields applied to the helix-coil transition of polypeptides," *J. Phys. Chem. B* **113**, 9004–9015 (2009).
- ⁸⁵K. Lindorff-Larsen, S. Piana, K. Palmo, P. Maragakis, J. L. Klepeis, R. O. Dror, and D. E. Shaw, "Improved side-chain torsion potentials for the Amber ff99SB protein force field," *Proteins* **78**, 1950–1958 (2010).
- ⁸⁶J. C. Gumbart, B. Roux, and C. Chipot, "Standard binding free energies from computer simulations: What is the best strategy?," *J. Chem. Theory Comput.* **9**, 794–802 (2013).
- ⁸⁷P. Burkhard, P. Taylor, and M. D. Walkinshaw, "X-ray structures of small ligand-FKBP complexes provide an estimate for hydrophobic interaction energies," *J. Mol. Biol.* **295**, 953–962 (2000).
- ⁸⁸D. Pramanik, Z. Smith, A. Kells, and P. Tiwary, "Can one trust kinetic and thermodynamic observables from biased metadynamics simulations?: Detailed quantitative benchmarks on millimolar drug fragment dissociation," *J. Phys. Chem. B* **123**, 3672–3678 (2019).
- ⁸⁹A. C. Pan, H. Xu, T. Palpant, and D. E. Shaw, "Quantitative characterization of the binding and unbinding of millimolar drug fragments with molecular dynamics simulations," *J. Chem. Theory Comput.* **13**, 3372–3377 (2017).
- ⁹⁰N. Ahalawat and J. Mondal, "Assessment and optimization of collective variables for protein conformational landscape: GB1 β -hairpin as a case study," *J. Chem. Phys.* **149**, 094101 (2018).
- ⁹¹V. Muñoz, P. A. Thompson, J. Hofrichter, and W. A. Eaton, "Folding dynamics and mechanism of β -hairpin formation," *Nature* **390**, 196–199 (1997).
- ⁹²R. M. Fesinmeyer, F. M. Hudson, and N. H. Andersen, "Enhanced hairpin stability through loop design: The case of the protein G B1 domain hairpin," *J. Am. Chem. Soc.* **126**, 7238–7243 (2004).
- ⁹³A. J. Hazel, E. T. Walters, C. N. Rowley, and J. C. Gumbart, "Folding free energy landscapes of β -sheets with non-polarizable and polarizable CHARMM force fields," *J. Chem. Phys.* **149**, 072317 (2018).
- ⁹⁴R. B. Best and J. Mittal, "Free-energy landscape of the GB1 hairpin in all-atom explicit solvent simulations with different force fields: Similarities and differences," *Proteins* **79**, 1318–1328 (2011).
- ⁹⁵A. Ardevol, G. A. Tribello, M. Ceriotti, and M. Parrinello, "Probing the unfolded configurations of a β -hairpin using sketch-map," *J. Chem. Theory Comput.* **11**, 1086–1093 (2015).
- ⁹⁶F. Nüske, B. G. Keller, G. Pérez-Hernández, A. S. Mey, and F. Noé, "Variational approach to molecular kinetics," *J. Chem. Theory Comput.* **10**, 1739–1752 (2014).
- ⁹⁷R. T. McGibbon and V. S. Pande, "Variational cross-validation of slow dynamical modes in molecular kinetics," *J. Chem. Phys.* **142**, 124105 (2015).
- ⁹⁸J. M. Lamim Ribeiro and P. Tiwary, "Toward achieving efficient and accurate Ligand-Protein unbinding with deep learning and molecular dynamics through RAVE," *J. Chem. Theory Comput.* **15**, 708–719 (2018).
- ⁹⁹M. Bonomi, D. Branduardi, F. L. Gervasio, and M. Parrinello, "The unfolded ensemble and folding mechanism of the C-terminal GB1 β -hairpin," *J. Am. Chem. Soc.* **130**, 13938–13944 (2008).
- ¹⁰⁰G. Bussi, F. L. Gervasio, A. Laio, and M. Parrinello, "Free-energy landscape for β hairpin folding from combined parallel tempering and metadynamics," *J. Am. Chem. Soc.* **128**, 13435–13441 (2006).
- ¹⁰¹G. Saladino, S. Pieraccini, S. Rendine, T. Recca, P. Francescato, G. Speranza, and M. Sironi, "Metadynamics study of a β -hairpin stability in mixed solvents," *J. Am. Chem. Soc.* **133**, 2897–2903 (2011).
- ¹⁰²S. Pressé, K. Ghosh, J. Lee, and K. A. Dill, "Principles of maximum entropy and maximum caliber in statistical physics," *Rev. Mod. Phys.* **85**, 1115 (2013).
- ¹⁰³W. Chen, H. Sidky, and A. L. Ferguson, "Nonlinear discovery of slow molecular modes using state-free reversible VAMPnets," *J. Chem. Phys.* **150**, 214114 (2019).
- ¹⁰⁴G. Bussi, D. Donadio, and M. Parrinello, "Canonical sampling through velocity rescaling," *J. Chem. Phys.* **126**, 014101 (2007).
- ¹⁰⁵J. Wang, R. M. Wolf, J. W. Caldwell, P. A. Kollman, and D. A. Case, "Development and testing of a general amber force field," *J. Comput. Chem.* **25**, 1157–1174 (2004).
- ¹⁰⁶G. J. Martyna, D. J. Tobias, and M. L. Klein, "Constant pressure molecular dynamics algorithms," *J. Chem. Phys.* **101**, 4177–4189 (1994).
- ¹⁰⁷P. Tiwary and M. Parrinello, "A time-independent free energy estimator for metadynamics," *J. Phys. Chem. B* **119**, 736–742 (2015).