Confronting pitfalls of AI-augmented molecular dynamics using statistical physics

Shashank Pant[a], Zachary Smith[b,d†], Yihang Wang[b,d†], Emad Tajkhorshid[a] and Pratyush Tiwary[c,d*]

[a]*NIH 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, IL 61801, USA*

[b] *Biophysics Program,* [c] *Department of Chemistry and Biochemistry,* [d]*Institute for Physical Science and Technology, University of Maryland, College Park, 20742, USA*

[†]*These authors contributed equally to this work*

E-mail: 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 seen use for 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 for many problems in chemistry and biophysics will be one that maximizes the time scale separation between slow and fast processes. To learn this timescale separation even from limited data, we use a maximum path entropy or caliber-based framework. We show the applicability of this automatic protocol for 3 classic benchmark problems, namely capturing the conformational dynamics of a model peptide (alanine dipeptide), ligand unbinding dynamics from a protein, and extensive sampling of the folding/unfolding energy landscape of the C-terminal domain of protein G (GB1-C16). We believe our work will lead to increased and robust use of trustworthy AI in molecular simulations of complex systems.

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 time-scale 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 like 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 like metadynamics,[22–24] umbrella-sampling (US),[25] adaptive biasing force (ABF)[26–28] and many others.[26,29–31] Although CV-based methods can be

computationally more efficient than tempering-based approaches, given a poor choice of low-dimensional modes (a non-trivial task to intuit for complex systems), CV biasing can fail miserably.[32] 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.[33,34]

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.[35–39] A number of such AI-based approaches have been proposed recently[35–37,40] 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[41–44] 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.[45] 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 time-scale 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. Here we use the "Spectral Gap Optimization of Order Parameters (SGOOP)" framework,[46] which uses notions of maximum caliber or path entropy[47] to cheaply construct 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 Auotencoded Variational Bayes for enhanced sampling (RAVE)",[38] even though we do expect it to be more generally applicable. 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 the presented algorithm marks a major step forward in the use of fully automated AI-enhanced MD for the study of complex bio-molecular processes.

Theory

AI can mislead

In this work our starting point is the recent AI-based method RAVE,[35,38,48] but the problems identified and solutions proposed are quite generic to the use of AI in molecular sciences. 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 s. These could be generic variables such as dihedrals or protein-ligand distances,[49] 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.[50–54] 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 future. RAVE uses the PIB as a computationally tractable approximation for the RC which is traditionally considered as the definition of a slow mode.[55] 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.[38,56,57]

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 $s = (s_1, s_2, ..., s_d)$, c_i denotes different weights[49] and d denotes the dimension of the RC

space. The PIB objective function which is then minimized in every training round can be written as a difference of two mutual informations: [58]

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

where $I(.)$ denotes the mutual information between two random variables.[58] 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 parameters 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 $I(\chi, \mathbf{s}_{\Delta t})$. 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 $I(\chi, \mathbf{s})$. 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[38,59] wherein optimizing the decoder is guaranteed to lead to a convex optimization problem, we are not concerned with over-fitting in the decoder. It is fitting the encoder, which directly leads to an interpretable RC, that is of concern to us here. This can be best illustrated through a simple numerical example involving protein conformational dynamics, which we describe in detail in **Methods** and in Fig. 1. We performed 6 different, independently initialized trials of PIB learning using the same input trajectory for a model peptide (alanine dipeptide), each running for same number of epoches. The RC was expressed as a linear combination of the sines and cosines of various Ramachandran

dihedral angles. As can be seen in Fig. 1B, we obtain different RCs with different trials even though they are all stopped at the same low value of the loss function (within 4 decimal digits of precision). Given the use of an interpretable linear encoder, one can see a sense of symmetry even in the at first glance different looking RCs in Fig. 1B. However as we show 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,[35,39,60] and/or where one does not really know *a priori* when to stop the training.

A

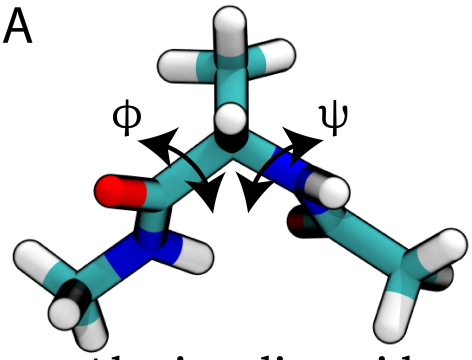

Alanine dipeptide

B

Trials	cosϕ	sinϕ	cosψ	sinψ	loss function
1	0.11	-0.59	-0.51	-0.62	0.2803
2	-0.07	0.63	0.48	0.60	0.2803
3	-0.05	0.63	0.47	0.61	0.2803
4	-0.12	0.58	0.52	0.62	0.2803
5	0.11	-0.59	-0.51	-0.61	0.2803
6	-0.11	0.58	0.52	0.62	0.2803

Figure 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 towards 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. .

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. Firstly, 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. Secondly, 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, namely the choice of the predictive time-delay Δt in Eq. 1. As shown in Ref.,[61] theoretically speaking the method is robust to the choice of this parameter as long as it non-zero yet small enough. In practice, it can be hard to judge whether it is indeed small enough or not.

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 for example through concepts such as Born-Oppenheimer approximation[62] and Michaelis-Menten principle.[63] 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.[64] 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,[46] which uses a maximum path entropy or caliber model[47,65] to construct a minimal model of the dynamics along a given low-dimensional projection. 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.[38,61] 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.[66] While improving the quality of the dynamical observables can lead to increasingly accurate eigenvalues,[67] here we use a computationally inexpensive dynamical observable denoted $\langle N \rangle$ and defined as the average number of nearest neighbor transitions per unit time along any RC. 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:[47,66]

$$K_{mn} = \frac{< N >}{\sum \sqrt{p_n p_m}} \sqrt{\frac{p_n}{p_m}} \tag{2}$$

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. The RC so identified 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.

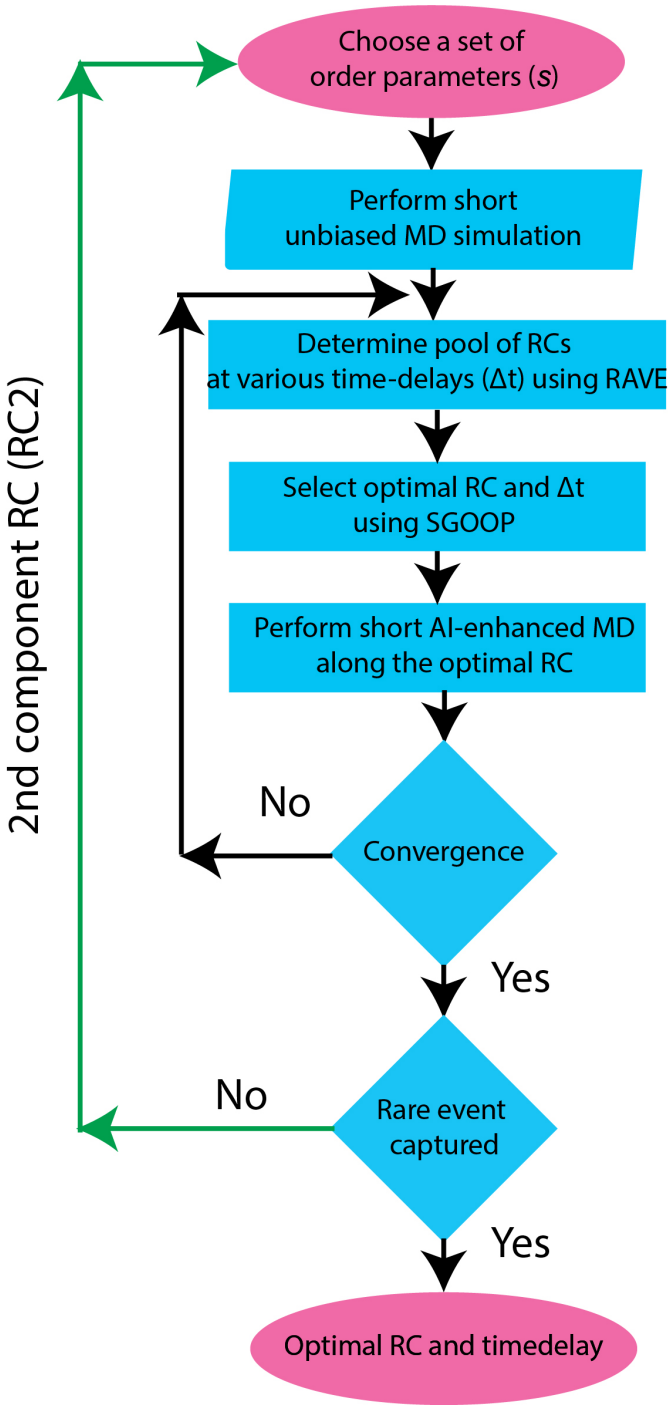

Figure 2: Flowchart illustrates our novel and computationally efficient protocol to screen AI solutions. Starting from short unbiased MD simulations, our protocol semi-automatically screen the spurious solutions obtained in AI-based method and learns the optimal RC. In this work we demonstrate the applicability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

Results

In the previous section 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. We do however expect our framework to be applicable to any molecular problem. Specifically, we consider (A) conformational dynamics of a model peptide in vaccum, (B) dissociation of a millimolar-affinity ligand from 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 of 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 which 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.[68] 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.[69] These static biases were directly obtained by inverting the probability distribution learnt during RAVE.[38,48] In system (C), we employ time-dependent biasing through well-tempered metadynamics[24,57] to capture folding of the GB1-C16 peptide.

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. 1B). Enabled by the small size of the

system, we performed 3 independent simulations, each 2 μs long. The corresponding trajectories along with the dihedral angles ϕ and ψ are provided as Supplementary Information (SI) in Fig. S1. In line with standard practice,[70,71] 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 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 to 40 ps. Furthermore, 10 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 value after the training as well as the RC so-obtained were recorded.

As hinted at in **Introduction**, we obtain very different RCs for the different Δt values and for different independent trials. Furthermore, different trials that were stopped at similar loss functions gave different RCs and spectral gaps (Fig. 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-C) we demonstrate the noisy correlation that we find between the loss function and the spectral gap for all 3 input trajectories. In SI (Fig. S2B), we provide an illustrative figure for one particular trajectory showing how the same loss function results in RCs with different free energy profiles, and that the one with the highest spectral gap stands out with 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. 3D). Irrespective of the choice of input trajectory, we find that the optimal RC shows higher weights for ϕ (as compared to ψ) (Table 1), in line with previous

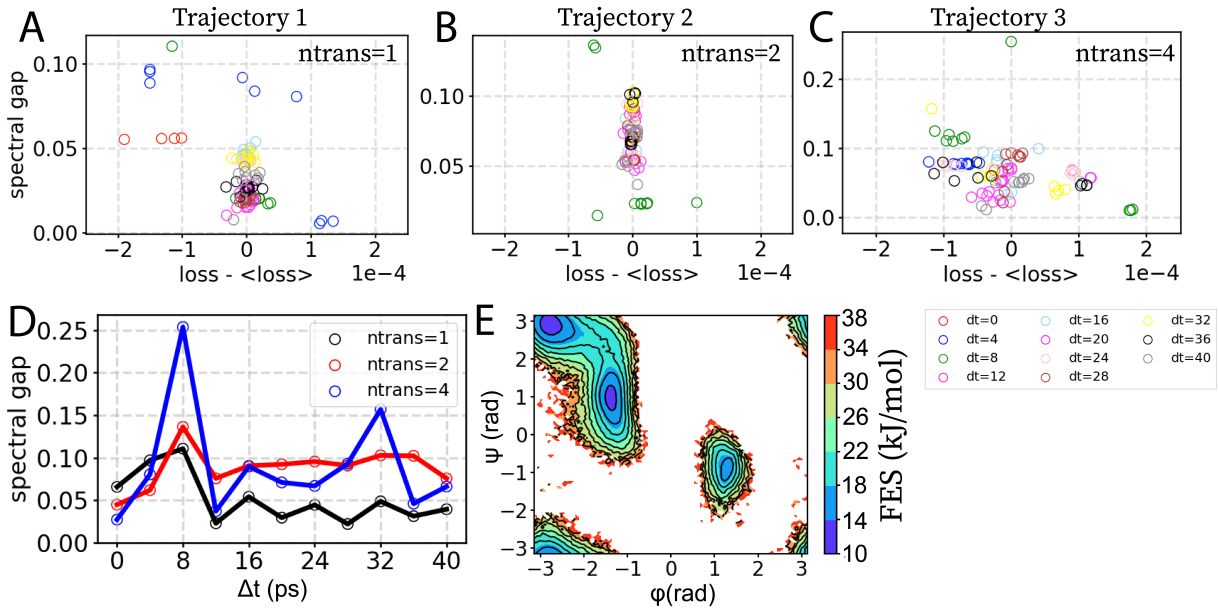

Figure 3: **Capturing the spurious AI solutions in alanine dipeptide.** Spectral gap and loss function were calculated for each of the three unbiased trajectories at multiple time-delays Δt between 0 and 40 ps, indicated using circles of different colors in the bottom right side of the figure. (A-C) show noisy correlation between the loss and the spectral gap for 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 by subtracting out the average of all losses. (D) Maximum spectral gap (out of 10 different trials of RAVE) vs. Δt was plotted for 3 different unbiased trajectories. (E) Free energy surface (FES) along the two dihedrals Φ and Ψ, obtained from 20 ns-long simulation in the presence of static bias. Energy contours every 4 kJ/mol.

studies which highlighted ϕ to be a more important degree of freedom than ψ.[38,72] Using this RC and its probability distribution as a fixed bias,[38] we then explored the conformational space of the peptide. The two-dimensional free-energy landscape along the two dihedrals ϕ and ψ from this biased run after only 20 ns of simulation (Fig. 3E) is in excellent agreement with the previous studies.[24,38]

Table 1: **Optimal weights of OPs obtained through a combination of RAVE and SGOOP**

Trajectory	Transitions	cosϕ	sinϕ	cosψ	sinψ
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

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. 4A). Force-field parametrization and other MD details are provided in SI. Here, our objective was to use RAVE to learn the most optimal RC on-the-fly as well as the absolute binding free energy of this protein-ligand complex. This is a difficult and important problem for which many useful methods have been already employed with varying levels of success.[73] 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).[74] The MD simulations were stopped when the ligand unbound, specifically when it was 2 Å away from the binding pocket (Fig. 4B). All trajectories were expressed in terms of 8 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.[49,75] We combined the results of the four independent MD trajectories to perform RAVE with 11 different choices of predictive time-delay, ranging from 0 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 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 **Theory** and done for alanine dipeptide. We again find noisy correlation between the loss function and the spectral gap (Fig. 4C), for the case of $\Delta t = 40$ ps (additional plots are given in Fig. S3). The same value of loss function gives rise to very different values of the spectral gap and of the RC (Fig. S3). Furthermore, the spectral gap not only captures the most optimal RC, it is also able to select the most optimal time-delay (Fig 4D). By using this RC and its probability distribution as a fixed bias,[38] we then performed 800 ns of biased simulations starting from the bound pose, but allowing the ligand to re-associate (Fig. S4A,B). Through this we then calculate the absolute binding affinity of the protein-ligand complex to be 6.6 kJ/mol (Fig 4E), in good agreement with values reported through metadynamics[49] and extended unbiased simulations on ANTON supercomputer.[76] It is worth pointing out that the ANTON simulations took 39 μs 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. 4F). 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 studies[74] have highlighted the importance of I56 as it forms hydrogen bonding interactions with the carbonyl group of the bound ligand, thus highlighting the significance of this residue which our algorithm also captured as the most significant OP. Followed by this highest weight component, the second and third highest com-

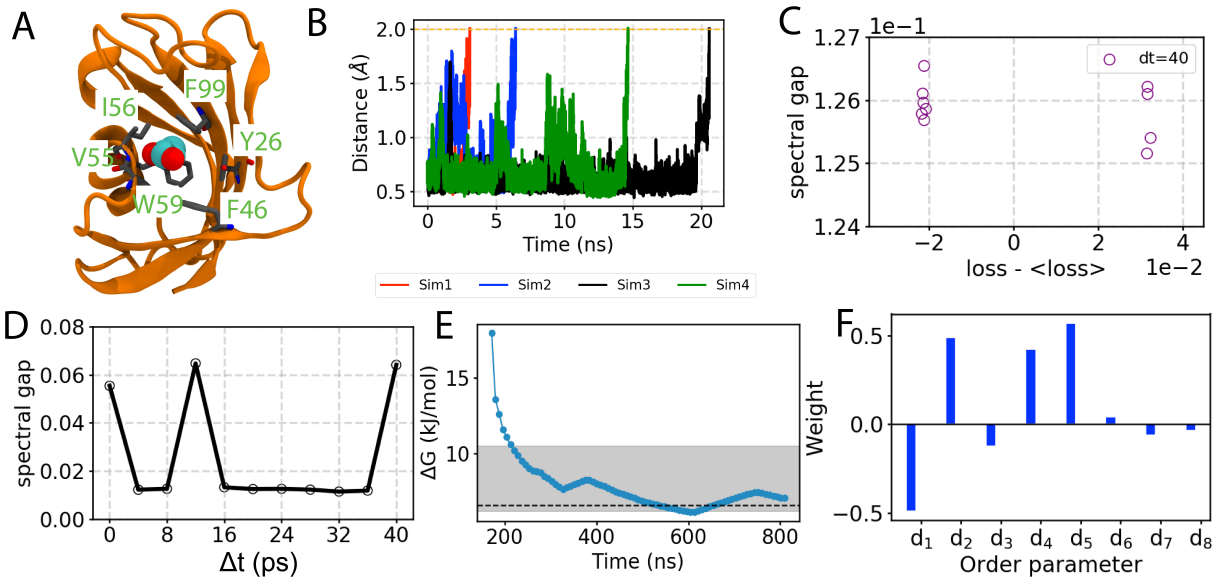

Figure 4: **Unbinding of 4-hydroxy-2-butanone (BUT) from FKBP.** (A) Molecular image of the bound FKBP/BUT protein–ligand complex, with binding pocket residues highlighted. The distances to the residues were used as different OPs detailed in Table S1. (B) Time evolution of the distance between the center 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 10 different trials were calculated after combining all four independent trajectories at multiple time-delay Δt between 0 and 40 ps, indicated using circles. Rest of the time-delays are shown in Fig. SS1. For visual clarity, at each iteration, we have plotted a mean-free version of the loss function by subtracting out the average of all losses. (D) Plot of maximum spectral gap (out of 10 different trials of RAVE) vs time-delay (Δt). (E) Absolute binding free energy G in kJ/mol of FKBP/BUT system as a function of simulation time with static external bias. The dotted black line shows the reference value reported through metadynamics[49] and the shaded region shows the free energy estimate from long unbiased MD simulations performed on Anton including the ± 0.5 kJ error reported.[76] (F) A visual depiction of the OP weights.

ponents are for d_1 and d_2, denoting respectively distances from the residues V55 and W59. 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.[77–81] Force-field parametrization and other MD details are provided in the SI. The free-energy landscape of this peptide has been extensively explored by replica-exchange MD simulations and clustering based methods.[77,81,82] 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 input in advanced slow mode analysis methods such as TICA[83–85] recovers a more superior two-dimensional description.[77] That work however used more than 12 μs of enhanced sampling, specifically replica exchange MD trajectories, for this purpose. Here instead we use just 1.6 μs of unbiased trajectories as our starting point. From this point onwards, using the same OPs as in Ref.,[77] 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[24,57] simulations as the choice of 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 SI (Fig. 5, Fig. S5). All the unbiased trajectories were mixed and fed into RAVE

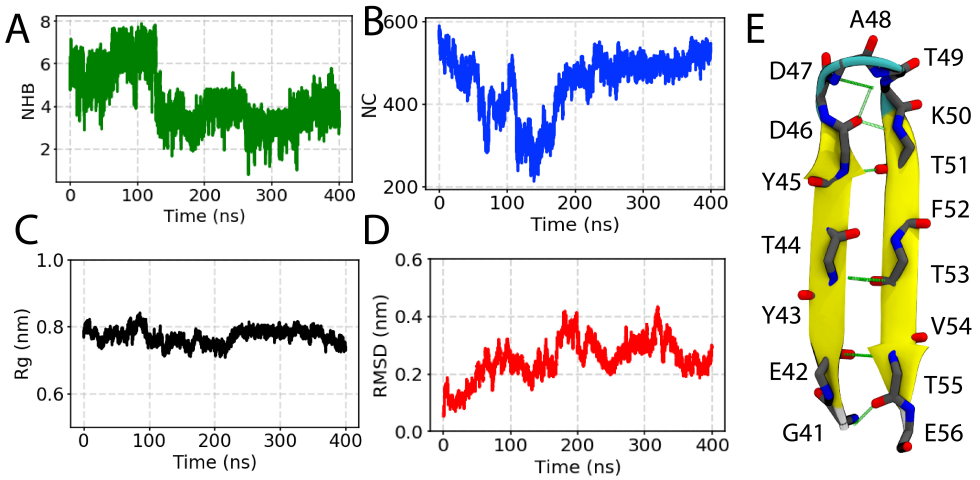

Figure 5: **Dynamics of GB1-C16 captured from unbiased MD.** One of the four representative trajectory of the peptide in explicit solvent is projected along different order parameters (A) number of hydrogen bonds (NHB), (B) native contacts (NC), (C) radius of gyration (Rg), and (D) root-mean square displacement (RMSD). (E) Molecular image of the GB1-C16. Native backbone hydrogen bonds are highlighted with green lines.

for learning the RC. We performed 10 different trials of RAVE for different time-delays Δt, ranging from 0 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 **Theory** and as done for the other two systems. Similar to the previous systems, we find noisy correlation between the loss function and the spectral gap (Fig. S7). The most optimal RC was selected for biased simulations using well-tempered metadynamics (Fig. S7A,B). Based on the maximum spectral gap, we chose Δt=8 ps for the next round of 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 sampling with this one-dimensional RC, which we call χ_1. With a 1-d 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.[86] 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 then performed 50 ns long 2D metadynamics simulations (Fig. S7 E,F), which were used to train χ_2. The most optimal 2-dimensional RC obtained after three iterations of training χ_2 is detailed in (Fig. 6A). The backbone heavy atom RMSD contribute the most toward the construction of slowest dimension χ_1, whereas Rg contributed more towards the second slowest dimension.

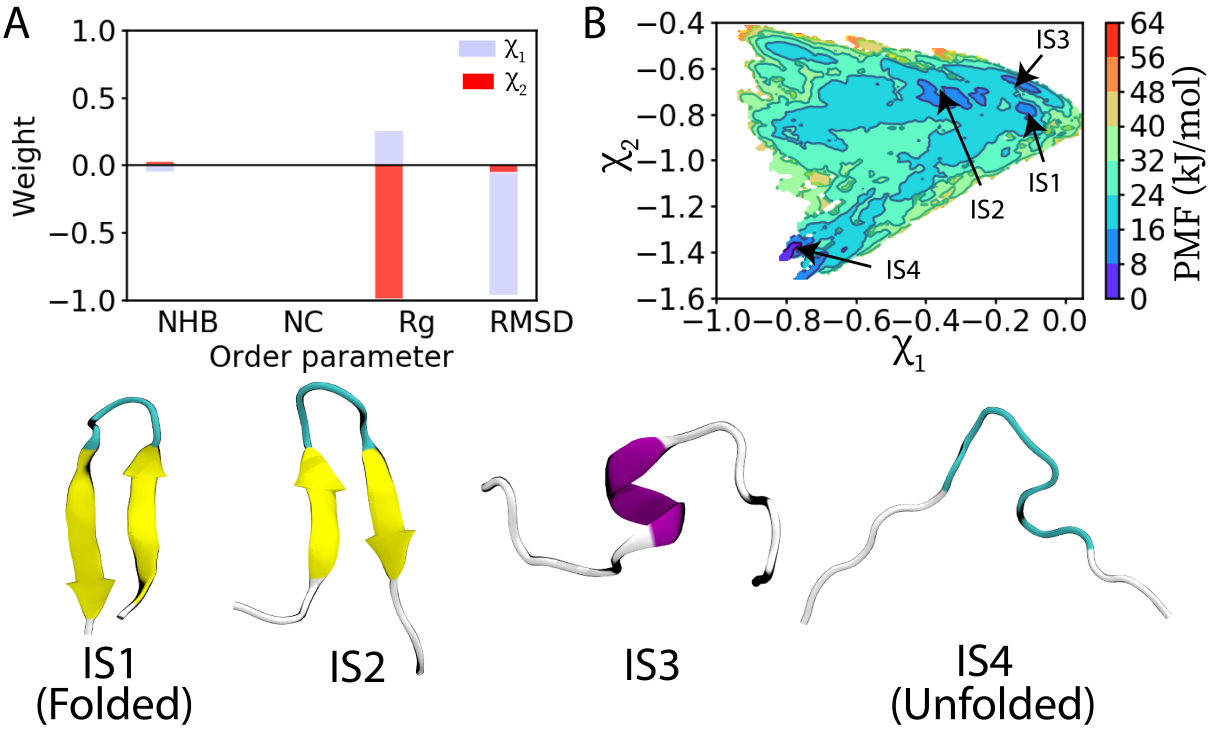

Figure 6: **Free-energy landscape and OPs contribution.** (A) Contribution of the different OPs to the 2-dimensional RC χ. The two components χ_1 and χ_2 are shown in blue and red bars, respectively. (B) A highly rugged 2-dimensional free-energy landscape of GB1-C16 folding/unfolding. We were able to capture multiple states, corresponding to the folded (IS1), unfolded (IS4), and intermediate states (IS2 and IS3). Interestingly, 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.

The two-dimensional RC is then used in longer well-tempered metadynamics simulations to facilitate movement between different metastable states (SI Video 1) and to obtain converged free energy surfaces. We performed 1.2 μs-long metadynamics simulations at 300 K, starting from the crystal structure (Fig. S6). The 2-dimensional metadynamics simulations were performed with initial hill height 0.5 kJ, bias factor = 10, Gaussian widths 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 SI). 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. Fig. 6B 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 using conventional OPs.[87] The two-dimensional free-energy landscape when projected along the pair of OPs, yields results consistent with the previous studies and suggests the presence of two metastable states (Fig. S8).

Conclusion

To conclude, we have introduced a new approach to sieve out the spurious solutions from AI-augmented enhanced sampling simulations.[35,36] AI-based approaches have had 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 mis-

leading 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.[88] Our algorithm is based on the simple but powerful notion that a more reliable AI solution will be 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.[46,47] We would like to emphasize that our approach, and spectral gap based optimization in general[46] might have as of yet unexplored connections with the Variational Approach for Markov processes (VAMP) .[89] The framework developed here is quite generic and should be applicable to many recent methods (Ref.[35] 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.[38] 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 4 order parameters that are generic for folding/unfolding processes and demonstrated how to semi-automatically learn a 2-dimensional RC, which we then used in well-tempered metadynamics protocol to obtain folding/unfolding trajectories. This directly allows us to gain atomic level insights 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 arguably for molecular simulations in general.

Acknowledgements

We thank the Donors of the American Chemical Society Petroleum Research Fund (No. PRF60512-DNI6) (PT), National Institutes of Health under award number P41-GM104601 (ET), the NCI-UMD Partnership for Integrative Cancer Research (YW), the COMBINE fellowship (DGE-1632976) (ZS), and Beckman Institute Graduate Fellowship (SP) for financial support. This work used the Extreme Science and Engineering Discovery Environment (XSEDE) Bridges through allocation CHE180027P, which is supported by National Science Foundation grant number ACI-1548562 (PT). We also thank UMD's Deepthought2 and MARCC's Bluecrab HPC clusters for computing resources. We would like to thank Navjeet Ahalawat and Jagannath Mondal for help with setting up the GB1 peptide system for MD simulation. We would also like to thank Andrew Ferguson for useful discussion and Joáo Marcelo Lamim Ribeiro for the careful reading of the manuscript.

References

(1) Bernardi, R. C.; Melo, M. C. R.; Schulten, K. Enhanced Sampling Techniques in Molecular Dynamics Simulations of Biological Systems. Biochim. Biophys. Acta **2015**, 1850, 872–877.

(2) Karplus, M.; Petsko, G. A. Molecular dynamics simulations in biology. Nature **1990**, 347, 631–639.

(3) Rohrdanz, M. A.; Zheng, W.; Clementi, C. Discovering mountain passes via torchlight: Methods for the definition of reaction coordinates and pathways in complex macromolecular reactions. Annu. Rev. Phys. Chem. **2013**, 64, 295–316.

(4) Abrams, C.; Bussi, G. Enhanced sampling in molecular dynamics using metadynamics, replica-exchange, and temperature-acceleration. Entropy **2014**, 16, 163–199.

(5) Hashemian, B.; Millán, D.; Arroyo, M. Modeling and enhanced sampling of molecular systems with smooth and nonlinear data-driven collective variables. J. Chem. Phys. **2013**, 139, 12B601_1.

(6) Onuchic, J. N.; Wolynes, P. G. Theory of protein folding. Curr. Opin. Struct. Biol. **2004**, 14, 70–75.

(7) Dill, K. A.; Ozkan, S. B.; Shell, M. S.; Weikl, T. R. The protein folding problem. Annu. Rev. Biophys. **2008**, 37, 289–316.

(8) Tiwary, P.; Limongelli, V.; Salvalaglio, M.; Parrinello, M. Kinetics of protein–ligand unbinding: Predicting pathways, rates, and rate-limiting steps. Proc. Natl. Acad. Sci. USA **2015**, 112, E386–E391.

(9) Tiwary, P.; Mondal, J.; Berne, B. J. How and when does an anticancer drug leave its binding site? Sci. Adv. **2017**, 3, e1700014.

(10) Moradi, M.; Tajkhorshid, E. Mechanistic picture for conformational transition of a membrane transporter at atomic resolution. Proc. Natl. Acad. Sci. USA **2013**, 110, 18916–18921.

(11) Moradi, M.; Tajkhorshid, E. Computational recipe for efficient description of large-scale conformational changes in biomolecular systems. J. Chem. Theory Comput. **2014**, 10, 2866–2880.

(12) Moradi, M.; Enkavi, G.; Tajkhorshid, E. Atomic-level characterization of transport cycle thermodynamics in the glycerol-3-phosphate:phosphate transporter. Nat. Commun. **2015**, 6, 8393.

(13) Pant, S.; Tajkhorshid, E. Microscopic Characterization of GRP1 PH Domain Interaction with Anionic Membranes. J. Comput. Chem. **2019**, 41, 489–499.

(14) Brooks, S. P.; Morgan, B. J. Optimization using simulated annealing. J. of the Royal Statistical Society: Series D (The Statistician) **1995**, 44, 241–257.

(15) Hansmann, U. H. Parallel tempering algorithm for conformational studies of biological molecules. Chemical Physics Letters **1997**, 281, 140–150.

(16) Sugita, Y.; Okamoto, Y. Replica-exchange molecular dynamics method for protein folding. Chem. Phys. Lett. **1999**, 314, 141–151.

(17) Sugita, Y.; Kitao, A.; Okamoto, Y. Multidimensional replica-exchange method for free-energy calculations. J. Chem. Phys. **2000**, 113, 6042–6051.

(18) Mitsutake, A.; Sugita, Y.; Okamoto, Y. Generalized-Ensemble Algorithms for Molecular Simulations of Biopolymers. Biopolymers **2001**, 60, 96–123.

(19) Fukunishi, H.; Watanabe, O.; Takada, S. On the Hamiltonian replica exchange method for efficient sampling of biomolecular systems: Application to protein structure prediction. J. Chem. Phys. **2002**, 116, 9058–9067.

(20) Maragliano, L.; Vanden-Eijnden, E. A temperature accelerated method for sampling free energy and determining reaction pathways in rare events simulations. Chem. Phys. Lett. **2006**, 426, 168–175.

(21) Miao, Y.; Feher, V. A.; McCammon, J. A. Gaussian accelerated molecular dynamics: Unconstrained enhanced sampling and free energy calculation. J. Chem. Theory Comput. **2015**, 11, 3584–3595.

(22) Laio, A.; Parrinello, M. Escaping free-energy minima. Proc. Natl. Acad. Sci. USA **2002**, 99, 12562–12566.

(23) Laio, A.; Gervasio, F. L. Metadynamics: a method to simulate rare events and reconstruct the free energy in biophysics, chemistry and material science. Rep. Progr. Phys. **2008**, 71, 126601.

(24) Barducci, A.; Bussi, G.; Parrinello, M. Well-tempered metadynamics: a smoothly converging and tunable free-energy method. Phys. Rev. Lett. **2008**, 100, 020603.

(25) Torrie, G. M.; Valleau, J. P. Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. J. Chem. Phys. **1977**, 23, 187–199.

(26) Darve, E.; Rodríguez-Gómez, D.; Pohorille, A. Adaptive biasing force method for scalar and vector free energy calculations. J. Chem. Phys. **2008**, 128, 144120.

(27) Fu, H.; Shao, X.; Chipot, C.; Cai, W. Extended adaptive biasing force algorithm. An on-the-fly implementation for accurate free-energy calculations. J. Chem. Theory Comput. **2016**, 12, 3506–3513.

(28) Comer, J.; Gumbart, J. C.; Hénin, J.; Lelièvre, T.; Pohorille, A.; Chipot, C. The adaptive biasing force method: Everything you always wanted to know but were afraid to ask. J. Phys. Chem. B **2015**, 119, 1129–1151.

(29) Abrams, J. B.; Tuckerman, M. E. Efficient and direct generation of multidimensional free energy surfaces via adiabatic dynamics without coordinate transformations. J. Phys. Chem. B **2008**, 112, 15742–15757.

(30) Kirkwood, J. G. Statistical mechanics of fluid mixtures. J. Chem. Phys. **1935**, 3, 300–313.

(31) den Otter, W. K.; Briels, W. J. The calculation of free-energy differences by constrained molecular-dynamics simulations. J. Chem. Phys. **1998**, 109, 4139–4146.

(32) Hazel, A.; Chipot, C.; Gumbart, J. C. Thermodynamics of deca-alanine folding in water. J. Chem. Theory Comput. **2014**, 10, 2836–2844.

(33) Liu, P.; Kim, B.; Friesner, R. A.; Berne, B. Replica exchange with solute tempering: A method for sampling biological systems in explicit water. Proc. Natl. Acad. Sci. USA **2005**, 102, 13749–13754.

(34) Huang, X.; Hagen, M.; Kim, B.; Friesner, R. A.; Zhou, R.; Berne, B. J. Replica exchange with solute tempering: efficiency in large scale systems. J. Phys. Chem. B **2007**, 111, 5405–5410.

(35) Wang, Y.; Ribeiro, J. M. L.; Tiwary, P. Machine learning approaches for analyzing and enhancing molecular dynamics simulations. Curr. Opin. Struct. Biol. **2020**, 61, 139–145.

(36) Noé, F.; De Fabritiis, G.; Clementi, C. Machine learning for protein folding and dynamics. Curr. Opin. Struct. Biol. **2020**, 60, 77–84.

(37) Noé, F.; Olsson, S.; Köhler, J.; Wu, H. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. Science **2019**, 365, eaaw1147.

(38) Wang, Y.; Ribeiro, J. M. L.; Tiwary, P. Past–future information bottleneck for sampling molecular reaction coordinate simultaneously with thermodynamics and kinetics. Nat. Commun. **2019**, 10, 1–8.

(39) Chen, W.; Tan, A. R.; Ferguson, A. L. Collective variable discovery and enhanced sampling using autoencoders: Innovations in network architecture and error function design. J. Comput. Chem. **2018**, 149, 072312.

(40) Tuckerman, M. E. Machine learning transforms how microstates are sampled. Science **2019**, 365, 982–983.

(41) Rotskoff, G.; Vanden-Eijnden, E. Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. Advances in neural information processing systems. 2018; pp 7146–7155.

(42) Cybenko, G. Approximation by superpositions of a sigmoidal function. Mathematics of Control, Signals and Systems **1992**, 5, 455–455.

(43) Barron, A. R. Universal approximation bounds for superpositions of a sigmoidal function. IEEE Transactions on Information theory **1993**, 39, 930–945.

(44) Bach, F. Breaking the curse of dimensionality with convex neural networks. The Journal of Machine Learning Research **2017**, 18, 629–681.

(45) Evtimov, I.; Eykholt, K.; Fernandes, E.; Kohno, T.; Li, B.; Prakash, A.; Rahmati, A.; Song, D. Robust physical-world attacks on deep learning models. arXiv preprint arXiv:1707.08945 **2017**,

(46) Tiwary, P.; Berne, B. Spectral gap optimization of order parameters for sampling complex molecular systems. Proc. Natl. Acad. Sci. USA **2016**, 113, 2839–2844.

(47) Ghosh, K.; Dixit, P. D.; Agozzino, L.; Dill, K. A. The Maximum Caliber Variational Principle for Nonequilibria. Annu. Rev. Phys. Chem. **2020**, 71, 213–238.

(48) Ribeiro, J. M. L.; Bravo, P.; Wang, Y.; Tiwary, P. Reweighted autoencoded variational Bayes for enhanced sampling (RAVE). J. Chem. Phys. **2018**, 149, 072301.

(49) Ravindra, P.; Smith, Z.; Tiwary, P. Automatic mutual information noise omission (AMINO): generating order parameters for molecular systems. Molecular Systems Design & Engineering **2020**,

(50) Tishby, N.; Pereira, F. C.; Bialek, W. The information bottleneck method. arXiv preprint physics/0004057 **2000**,

(51) Palmer, S. E.; Marre, O.; Berry, M. J.; Bialek, W. Predictive information in a sensory population. Proc. Natl. Acad. Sci. USA **2015**, 112, 6908–6913.

(52) Berman, G. J.; Bialek, W.; Shaevitz, J. W. Predictability and hierarchy in Drosophila behavior. Proc. Natl. Acad. Sci. USA **2016**, 113, 11943–11948.

(53) Alemi, A. A.; Fischer, I.; Dillon, J. V.; Murphy, K. Deep variational information bottleneck. arXiv preprint arXiv:1612.00410 **2016**,

(54) Still, S. Information bottleneck approach to predictive inference. Entropy **2014**, 16, 968–989.

(55) Krivov, S. V. On reaction coordinate optimality. Journal of chemical theory and computation **2013**, 9, 135–146.

(56) Smith, Z.; Ravindra, P.; Wang, Y.; Cooley, R.; Tiwary, P. Discovering loop conformational flexibility in T4 lysozyme mutants through Artificial Intelligence aided Molecular Dynamics. bioRxiv **2020**,

(57) Valsson, O.; Tiwary, P.; Parrinello, M. Enhancing important fluctuations: Rare events and metadynamics from a conceptual viewpoint. **2016**, 67, 159–184.

(58) Cover, T. M.; Thomas, J. A. Elements of information theory; John Wiley & Sons, 2012.

(59) Goodfellow, I.; Bengio, Y.; Courville, A. Deep Learning; MIT Press, 2016; `http://www.deeplearningbook.org`.

(60) Noé, F.; Nuske, F. A variational approach to modeling slow processes in stochastic dynamical systems. Multiscale Modeling & Simulation **2013**, 11, 635–655.

(61) Wang, Y.; Tiwary, P. Understanding the role of predictive time delay and biased propagator in RAVE. J. Chem. Phys. **2020**, 152, 144102.

(62) Levine, I. N.; Busch, D. H.; Shull, H. Quantum chemistry; Pearson Prentice Hall Upper Saddle River, NJ, 2009; Vol. 6.

(63) Nelson, D. L.; Lehninger, A. L.; Cox, M. M. Lehninger principles of biochemistry; Macmillan, 2008.

(64) Truhlar, D. G.; Garrett, B. C. Variational transition state theory. Annu. Rev. Phys. Chem. **1984**, 35, 159–189.

(65) Dixit, P. D.; Dill, K. A. Caliber corrected Markov modeling (C2M2): Correcting equilibrium Markov models. J. Chem. Theory Comput. **2018**, 14, 1111–1119.

(66) Smith, Z.; Pramanik, D.; Tsai, S.-T.; Tiwary, P. Multi-dimensional spectral gap optimization of order parameters (SGOOP) through conditional probability factorization. J. Chem. Phys. **2018**, 149, 234105.

(67) Meral, D.; Provasi, D.; Filizola, M. An efficient strategy to estimate thermodynamics and kinetics of G protein-coupled receptor activation using metadynamics and maximum caliber. J. Chem. Phys. **2018**, 149, 224101.

(68) Tiwary, P.; van de Walle, A. Multiscale Materials Modeling for Nanomechanics; Springer, 2016; pp 195–221.

(69) Debnath, J.; Parrinello, M. Gaussian Mixture Based Enhanced Sampling For Statics And Dynamics. J. Phys. Chem. Lett. **2020**,

(70) Mu, Y.; Nguyen, P. H.; Stock, G. Energy landscape of a small peptide revealed by dihedral angle principal component analysis. Proteins: Struct., Func., Bioinf. **2005**, 58, 45–52.

(71) Altis, A.; Nguyen, P. H.; Hegger, R.; Stock, G. Dihedral angle principal component analysis of molecular dynamics simulations. J. Chem. Phys. **2007**, 126, 244111.

(72) Salvalaglio, M.; Tiwary, P.; Parrinello, M. Assessing the reliability of the dynamics reconstructed from metadynamics. J. Chem. Theory Comput. **2014**, 10, 1420–1425.

(73) Gumbart, J. C.; Roux, B.; Chipot, C. Standard binding free energies from computer simulations: What is the best strategy? Journal of chemical theory and computation **2013**, 9, 794–802.

(74) Burkhard, P.; Taylor, P.; Walkinshaw, M. D. X-ray structures of small ligand-FKBP complexes provide an estimate for hydrophobic interaction energies. J. Mol. Biol. **2000**, 295, 953–962.

(75) Pramanik, D.; Smith, Z.; Kells, A.; Tiwary, P. Can One Trust Kinetic and Thermodynamic Observables from Biased Metadynamics Simulations?: Detailed Quantitative Benchmarks on Millimolar Drug Fragment Dissociation. J. Phys. Chem. B **2019**, 123, 3672–3678.

(76) Pan, A. C.; Xu, H.; Palpant, T.; Shaw, D. E. Quantitative characterization of the binding and unbinding of millimolar drug fragments with molecular dynamics simulations. Journal of chemical theory and computation **2017**, 13, 3372–3377.

(77) Ahalawat, N.; Mondal, J. Assessment and optimization of collective variables for protein conformational landscape: GB1 β-hairpin as a case study. J. Chem. Phys. **2018**, 149, 094101.

(78) Munoz, V.; Thompson, P. A.; Hofrichter, J.; Eaton, W. A. Folding dynamics and mechanism of β-hairpin formation. Nature **1997**, 390, 196–199.

(79) Fesinmeyer, R. M.; Hudson, F. M.; Andersen, N. H. Enhanced hairpin stability through loop design: the case of the protein G B1 domain hairpin. J. Am. Chem. Soc. **2004**, 126, 7238–7243.

(80) Hazel, A. J.; Walters, E. T.; Rowley, C. N.; Gumbart, J. C. Folding free energy landscapes of β-sheets with non-polarizable and polarizable CHARMM force fields. J. Chem. Phys. **2018**, 149, 072317.

(81) Best, R. B.; Mittal, J. Free-energy landscape of the GB1 hairpin in all-atom explicit solvent simulations with different force fields: Similarities and differences. Proteins: Struct., Func., Bioinf. **2011**, 79, 1318–1328.

(82) Ardevol, A.; Tribello, G. A.; Ceriotti, M.; Parrinello, M. Probing the unfolded configurations of a β-hairpin using sketch-map. J. Chem. Theory Comput. **2015**, 11, 1086–1093.

(83) Pérez-Hernández, G.; Paul, F.; Giorgino, T.; De Fabritiis, G.; Noé, F. Identification of slow molecular order parameters for Markov model construction. J. Chem. Phys. **2013**, 139, 07B604_1.

(84) Nuske, F.; Keller, B. G.; Pérez-Hernández, G.; Mey, A. S.; Noé, F. Variational approach to molecular kinetics. J. Chem. Theory Comput. **2014**, 10, 1739–1752.

(85) McGibbon, R. T.; Pande, V. S. Variational cross-validation of slow dynamical modes in molecular kinetics. J. Chem. Phys. **2015**, 142, 03B621_1.

(86) Lamim Ribeiro, J. M.; Tiwary, P. Toward achieving efficient and accurate Ligand-Protein unbinding with deep learning and molecular dynamics through RAVE. J. Chem. Theory Comput. **2018**, 15, 708–719.

(87) Bonomi, M.; Branduardi, D.; Gervasio, F. L.; Parrinello, M. The unfolded ensemble and folding mechanism of the C-terminal GB1 β-hairpin. J. Am. Chem. Soc. **2008**, 130, 13938–13944.

(88) Pressé, S.; Ghosh, K.; Lee, J.; Dill, K. A. Principles of maximum entropy and maximum caliber in statistical physics. Rev. Mod. Phys. **2013**, 85, 1115.

(89) Chen, W.; Sidky, H.; Ferguson, A. L. Nonlinear discovery of slow molecular modes using state-free reversible VAMPnets. J. Chem. Phys. **2019**, 150, 214114.

(90) Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. SoftwareX **2015**, 1, 19–25.

(91) Tribello, G. A.; Bonomi, M.; Branduardi, D.; Camilloni, C.; Bussi, G. PLUMED 2: New feathers for an old bird. Comput. Phys. Commun. **2014**, 185, 604–613.

(92) Bonomi, M.; Bussi, G.; Camilloni, C.; Tribello, G. A.; Banáš, P.; Barducci, A.; Bernetti, M.; Bolhuis, P. G.; Bottaro, S.; Branduardi, D., et al. Promoting transparency and reproducibility in enhanced molecular simulations. Nat. Methods

(93) Bussi, G.; Donadio, D.; Parrinello, M. Canonical sampling through velocity rescaling. J. Chem. Phys. **2007**, 126, 014101.

(94) Hornak, V.; Abel, R.; Okur, A.; Strockbine, B.; Roitberg, A.; Simmerling, C. Comparison of multiple Amber force fields and development of improved protein backbone parameters. Proteins: Struct., Func., Bioinf. **2006**, 65, 712–725.

(95) Best, R. B.; Hummer, G. Optimized molecular dynamics force fields applied to the helix- coil transition of polypeptides. J. Phys. Chem. B **2009**, 113, 9004–9015.

(96) Lindorff-Larsen, K.; Piana, S.; Palmo, K.; Maragakis, P.; Klepeis, J. L.; Dror, R. O.; Shaw, D. E. Improved side-chain torsion potentials for the Amber ff99SB protein force field. Proteins: Struct., Func., Bioinf. **2010**, 78, 1950–1958.

(97) Wang, J.; Wolf, R. M.; Caldwell, J. W.; Kollman, P. A.; Case, D. A. Development and testing of a general amber force field. J. Comput. Chem. **2004**, 25, 1157–1174.

(98) Martyna, G. J.; Tobias, D. J.; Klein, M. L. Constant pressure molecular dynamics algorithms. J. Chem. Phys. **1994**, 101, 4177–4189.

(99) Tiwary, P.; Parrinello, M. A time-independent free energy estimator for metadynamics. J. Phys. Chem. B **2015**, 119, 736–742.

Supporting Information

Methods

Molecular dynamics simulations All the systems were simulated with GROMACS version 5.0[90] patched with PLUMED version 2.4.2.[91,92] We constrained the bonds involving hydrogen atoms using LINCS algorithm and employed an integraion timestep of 2 fs. More details on individual systems and simulation protocols are provided below.

(A) **Simulation setup for alanine dipeptide in vacuum.** We followed Ref.[38] to set up our simulations for alanine dipeptide in vacuum. AMBER03 force field was used for the dipeptide. Three independent 2 μs-long MD simulations each with randomized initial velocities were performed to improve the sampling. The temperature was kept constant at 300 K using the velocity rescaling thermostat.[93]

(B) **Simulation setup for the FKBP-BUT complex.** As a second system, we studied the dissociation of 4-hydroxy-2-butanone (BUT) from the protein FKBP. AMBER99SB-ILDN force field[94–96] was used for the protein along with TIP3P water, and the ligand was parametrized with the generalized amber force field (GAFF).[97] The simulations were initiated from the BUT-bound crystal structure of FKBP (PDB ID: 1D7J). Initial equilibration was performed in NPT ensemble for 1 ns followed by four independent production runs with Nosé-Hoover thermostat[98] maintaining the temperature at 300 K.

(C) **Simulation setup for GB1 peptide.** As a third system, we studied the folding dynamics of GB1 peptide in water. The system was setup and equilibrated following the protocol described in Ref .[77] After this we performed four independent production runs, each lasting for 400 ns.

Metadynamics The central idea behind metadynamics is that by adding a history-dependent biasing potential it encourages the given system to escape the free energy minima, and explores various stable and metastable regions of the energy landscape.[24] Following the reweighting procedure, one can then construct the underlying energy landscape along any RC.[99] To capture the unfolding energy landscape of GB1 peptide, we employed well-tempered metadynamics using PLUMED implementation.[24,91,92] The simulations were performed along the optimal RC obtained through the combination of RAVE and SGOOP. The metadynamics simulations were performed with the biasfactor of 10, initial hill height of 0.5 kJ/mol, and bias deposition rate of 4 ps. The sigma values of the added Gaussian bias was set to the standard deviation of the biasing RC obtained from the unbiased simulation 400 ns.

Once a 50 ns metadynamics run was completed, the biased trajectory was reweighted and used as an input for the second round of RAVE.[38] The system in this 50 ns biased runs explored more conformational space than in

initial 400 ns of unbiased simulations. After this, we performed another round of biased simulations, where we employed 2D metadynamics along the optimal RCs obtained through the combination of RAVE and SGOOP. To prevent the peptide from getting trapped in entropic non-physical conformations due to the accumulated bias, we also restrain the movement of the peptide along the RMSD order parameter by addition of another bias. This bias had the form of Eq. 3. The final round of 2-dimensional metadynamics simulations were performed for 1200 ns.

$$V_{RMSD} = 2000(RMSD - 0.7)^{0.4}(kJ/mol) \tag{3}$$

Definitions of Order Parameters. For each system we first constructed a dictionary of order parameters (OPs), summarized in Table 2.

Table S1: **Details of the Pre-Selected Dictionary of Order Parameters (OPs)**

system	**order parameters**	**definition**
Alanine dipeptide	dihedral angle	cosϕ
	dihedral angle	sinϕ
	dihedral angle	cosψ
	dihedral angle	sinψ
FKBP-BUT	d_1	Y59-BUT
	d_2	V55-BUT
	d_3	F46-BUT
	d_4	Y26-BUT
	d_5	I56-BUT
	d_6	F99-BUT
	d_7	Y59-end A
	d_8	Y59-end B
GB1 peptide	NHB	backbone-backbone
	NC	all native contacts
	RG	backbone heavy atoms
	RMSD	backbone heavy atoms

We have calculated NC and NHB using a switching function

$$S = \sum_i \sum_j s_{ij} \tag{4}$$

$$s_{ij} = \begin{cases} \frac{1-(\frac{r_{ij}-d_o}{r_o})^6}{1-(\frac{r_{ij}-d_o}{r_o})^{12}} & if r_{ij} > r_o \\ 1 & otherwise \end{cases} \tag{5}$$

where d_o= 0.05 nm and r_o = 0.65 nm were used for NC calculations and d_o= 0.0 nm and r_o = 0.25 nm were used for NHB calculations. In the case of NC contacts, r_{ij} is a distance between i and j heavy atoms in the native state. While, in NHB, i and j are the pairs of O and HN (amino group) atoms of the backbone. RMSD of all the backbone heavy atoms with respect to crystal structure and RG of all the backbone heavy atoms were also calculated using PLUMED.

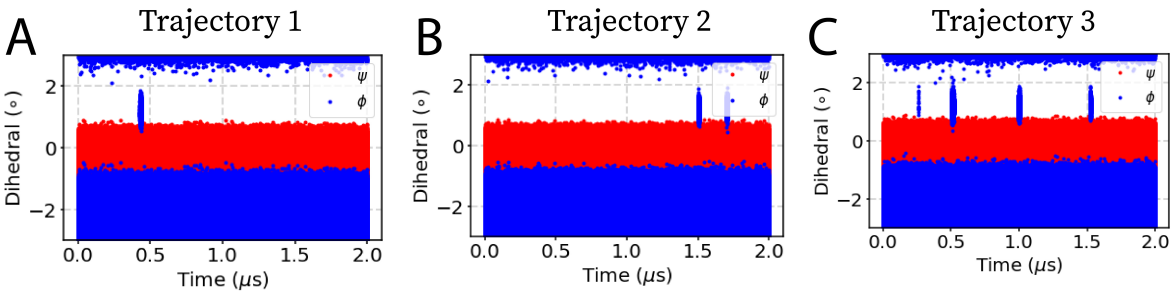

Figure S1: **Unbiased MD simulations of model peptide in vacuum.** (A-B) Three independent MD simulations of alanine dipeptide projected along dihedral (ϕ, ψ) angle space. Differing number of transitions along the ϕ order parameter were captured in the different independent simulations.

A

Trials	$\cos\phi$	$\sin\phi$	$\cos\psi$	$\sin\psi$	loss function	spectral gap
1	-0.695	-0.701	0.067	-0.138	0.3902	0.1121
2	0.691	0.704	-0.076	0.144	0.3902	0.1094
3	0.627	-0.654	0.029	0.423	0.3902	0.2536

B

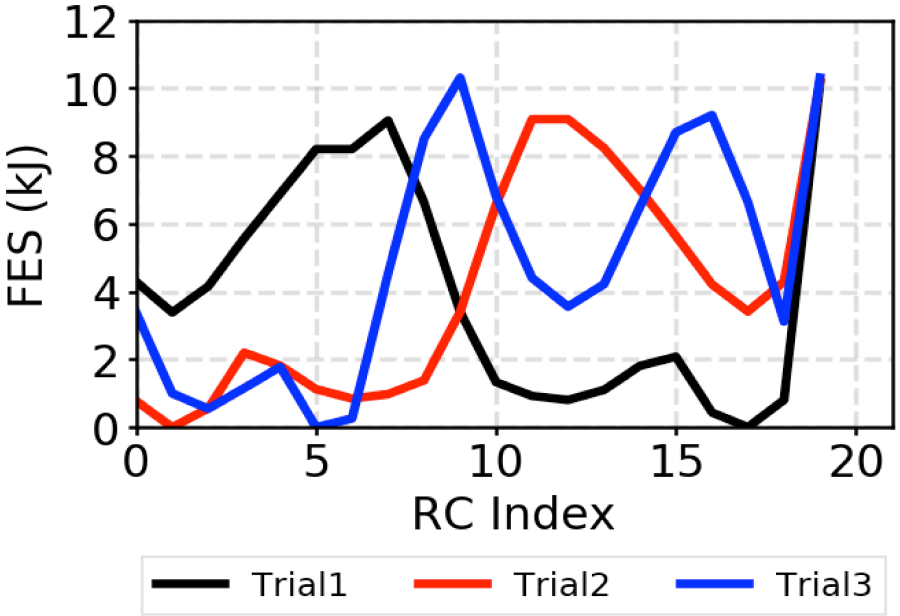

Figure S2: **Same loss function can give differing RCs** (A) Table highlights the sensitivity of spectral gap towards changes in the weights of the order parameters. While objective function remain insensitive to the changes in the weights of the order parameter. (B) Free energy landscape of one particular trajectory showing same loss function results in RCs with very different free energy profiles. The RC with maximum spectral gap highlights distinct metastable state.

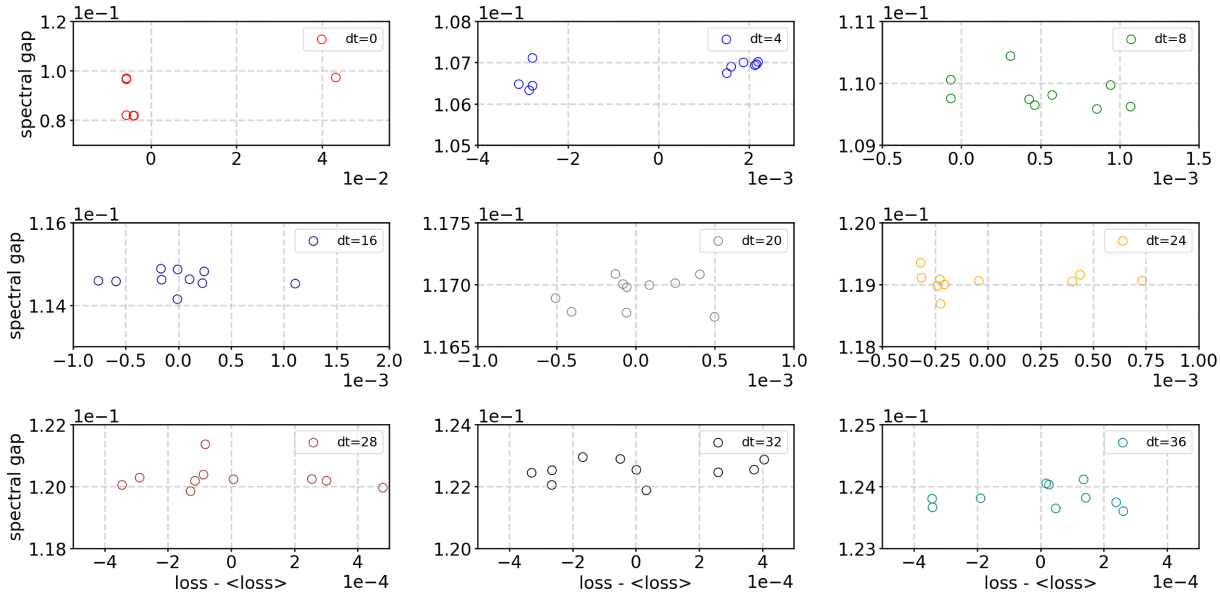

Figure S3: **Loss function versus spectral gap for ligand unbinding from FKBP along the optimal reaction coordinate.** Spectral gap and loss was calculated for the unbiased trajectory at multiple time-delays Δt between 0 and 40 ps, indicated using circles of different colors. At each time-delay we observed a noisy correlation between the loss and the spectral gap. For visual clarity, we have plotted a mean-free version of the loss function by subtracting out the average of all losses.

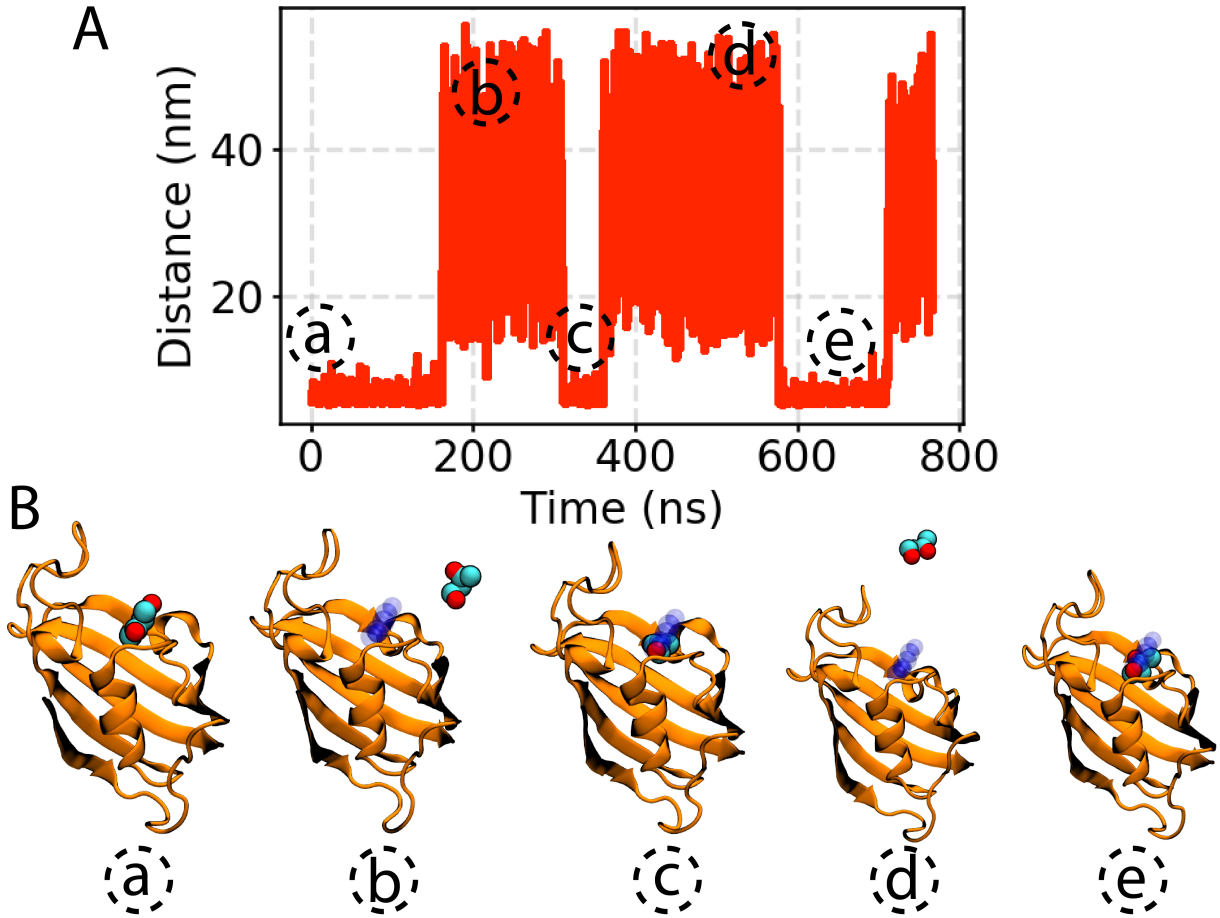

Figure S4: **Biased simulation of ligand unbinding from FKBP along the optimal reaction coordinate.** Biased simulation of ligand unbinding was performed under static bias. (A) Distance between the center of mass of the ligand from its binding pocket highlights multiple unbinding/binding events. (B) Representative snapshots of the trajectory are shown. Ligand is shown in vdW representation and the protein is shown in new cartoon (orange). Starting orientation of the ligand is shown in transparent blue vdW representation.

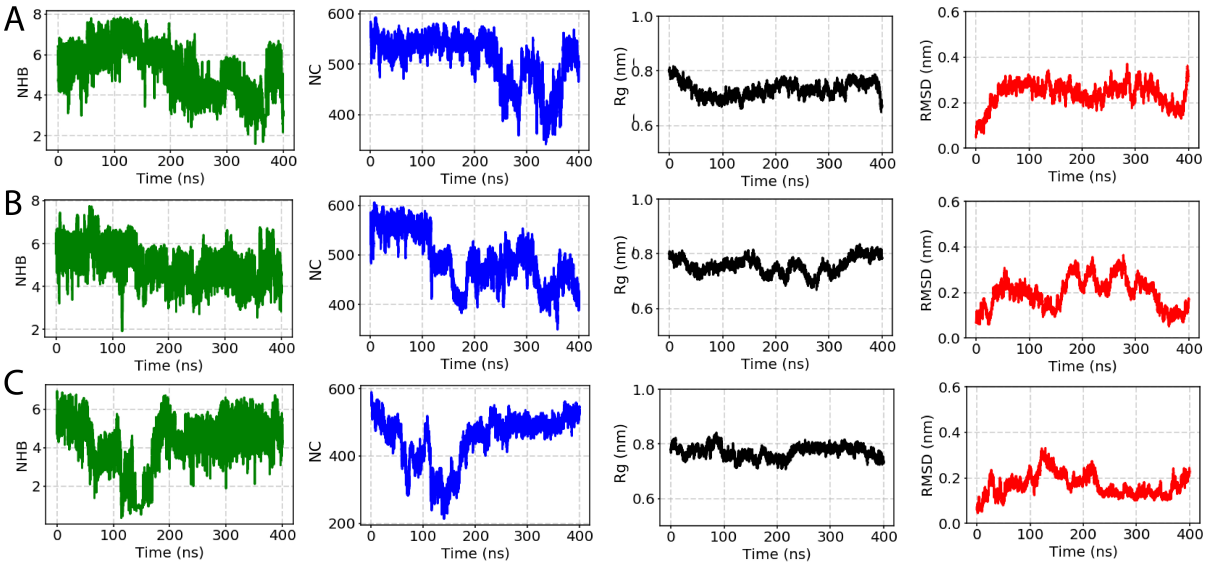

Figure S5: **Unbiased MD simulations of GB1 peptide in explicit solvent.** (A-C) Three independent MD simulations of GB1 peptide projected along different order parameters. Time evolution of OPs: root-mean square displacement (RMSD), radius of gyration (Rg), native contacts (NC), and number of backbone hydrogen bonds (NHB) suggests that the GB1 peptide was stable in all the simulated replicas.

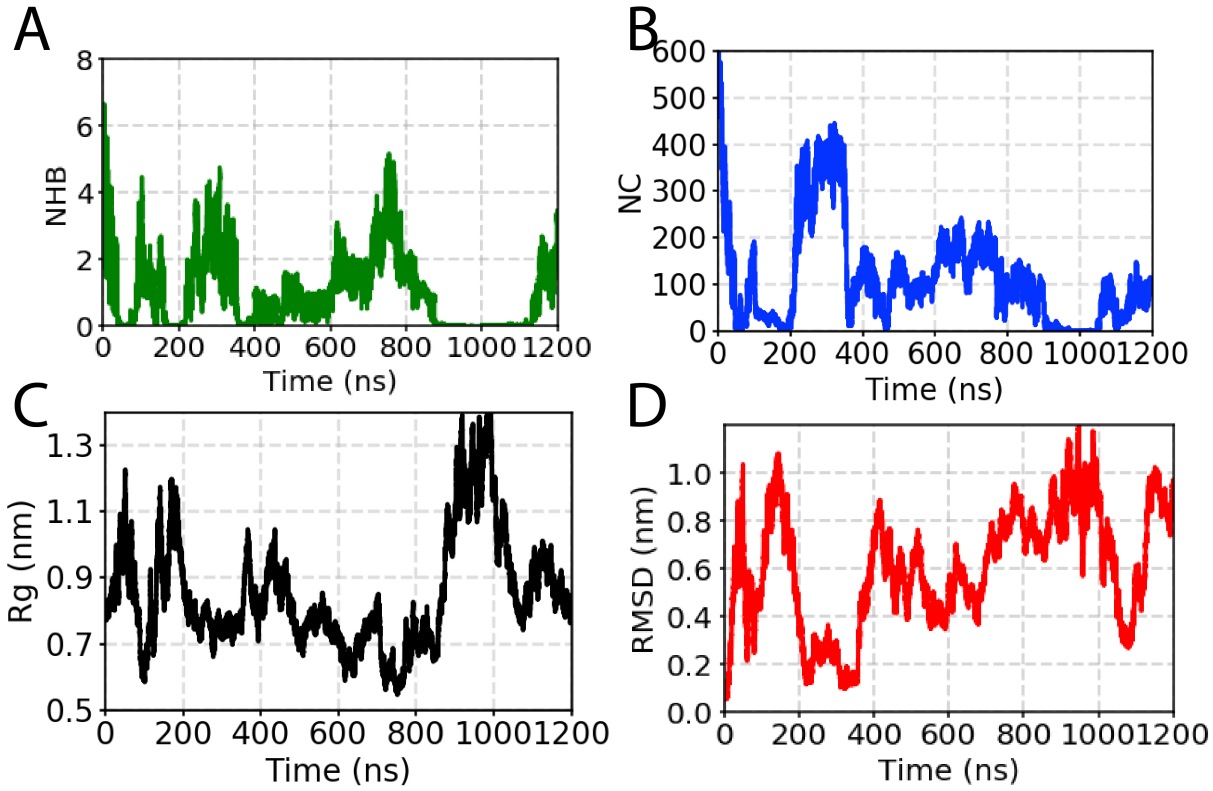

Figure S6: **Biased simulation along the optimal 2-d reaction coordinate.** The biased simulation was projected along the dictionary of order parameters used in this study. We captured multiple back and forth movement of the peptide between folded and unfolded states.

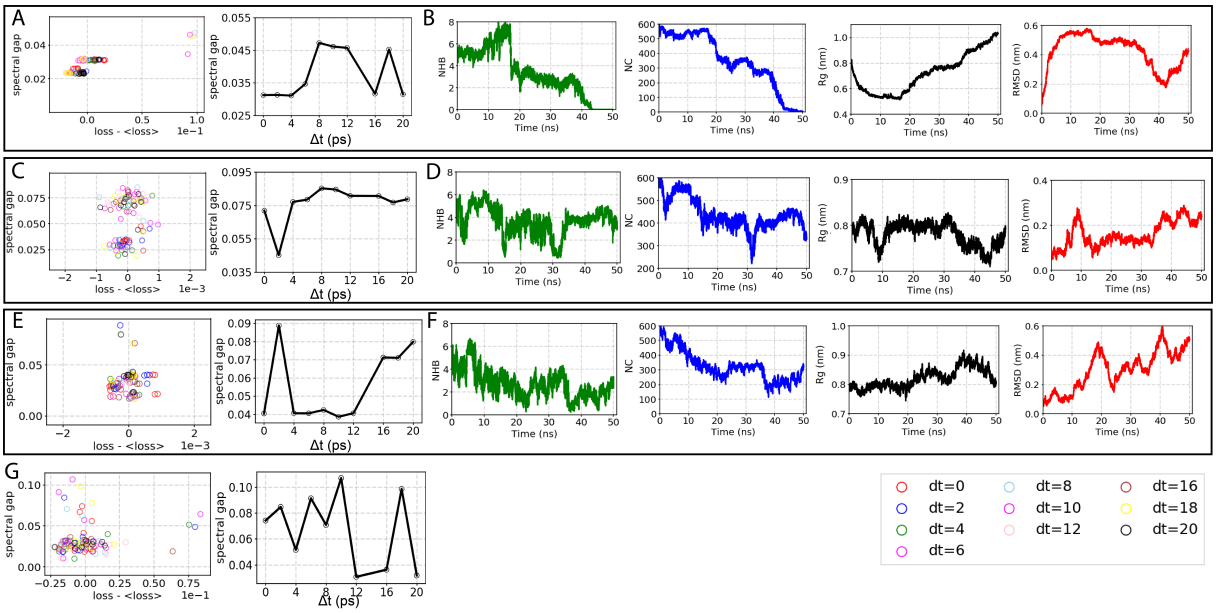

Figure S7: **Capturing optimal reaction coordinate for GB1 peptide.** Using an iterative ML-MD approach which incorporates RAVE and SGOOP, we have constructed an optimal RC to capture folding/unfolding of GB1 peptide. Spectral gap and loss was calculated in each round at multiple time-delays Δt between 0 and 20 ps, indicated using circles of different colors. (A,C,E,G) Show noisy correlation between the loss and the spectral gap at each iteration. Different circles denote different independent trials, with color denoting Δt. For visual clarity, at each iteration, we have plotted a mean-free version of the loss function by subtracting out the average of all losses. Corresponding maximum spectral gap (out of10 different trials of RAVE) vs. Δt was plotted at each iteration. (B,D,F) 50 ns-long biased simulation along the constructed RC (at ieach iteration) was performed and the trajectory was projected along the order different parameters employed in this study.

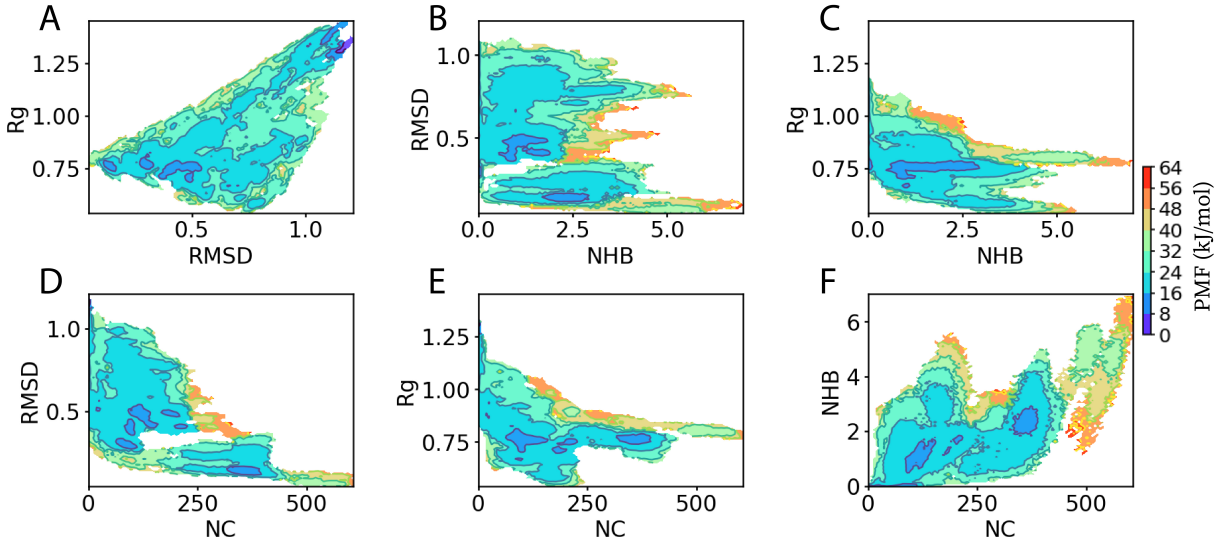

Figure S8: 2-dimensional free-energy profiles along the different OPs: root-mean square displacement (RMSD), radius of gyration (Rg), native contacts (NC), and number of backbone hydrogen bonds (NHB) used in this study.