

# Achieving Reversible Ligand-Protein Unbinding with Deep Learning and Molecular Dynamics through RAVE

João Marcelo Lamim Ribeiro and Pratyush Tiwary\*

*Department of Chemistry and Biochemistry and Institute for Physical Science and Technology, University of Maryland, College Park 20742, USA.*

## Abstract

In this work we demonstrate how to leverage our recent iterative deep learning–all atom molecular dynamics (MD) technique “Reweighted autoencoded variational Bayes for enhanced sampling (RAVE)” (Ribeiro, Bravo, Wang, Tiwary, *J. Chem. Phys.* 149, 072301 (2018)) for sampling protein-ligand unbinding mechanisms and calculating absolute binding affinities when plagued with difficult to sample rare events. RAVE iterates between rounds of MD and deep learning, and unlike other enhanced sampling methods, it stands out in simultaneously learning both a low-dimensional physically interpretable reaction coordinate (RC) and associated free energy. Here, we introduce a simple but powerful extension to RAVE which allows learning a position-dependent RC expressed as a superposition of piecewise linear RCs valid in different metastable states. With this approach, we retain the original physical interpretability of a RAVE-derived RC while making it applicable to a wider range of complex systems. We demonstrate how in its multi-dimensional form introduced here, RAVE can efficiently simulate the unbinding of the tightly bound benzene-lysozyme (L99A variant) complex, in all atom-precision and with minimal use of human intuition except for the choice of a larger dictionary of order parameters. These simulations had a 100 % success rate, and took between 3–50 nanoseconds for a process that takes on an average close to few hundred milliseconds, thereby reflecting a seven order of magnitude acceleration relative to straightforward MD. Furthermore, without any time-dependent biasing, the trajectories display clear back-and-forth movement between various metastable intermediates, demonstrating the reliability of the RC and its probability distribution learnt in RAVE. Our binding free energy is in good agreement with other reported simulation results. We thus believe that RAVE, especially in its multi-dimensional variant introduced here, will be a useful tool for simulating the dissociation process of practical biophysical systems with rare events in an automated manner with minimal use of human intuition.

## 1 INTRODUCTION

Modern all-atom simulation approaches to tackling open problems in the biological sciences must confront an inherent limitation stemming from the difference in timescales between the fast rovibrational molecular motions

relative to the (much) slower biochemical processes. With these fast molecular motions acting to constrain the timestep for integrating Newton’s equations of motion to small femtosecond values, it can be difficult to simulate timescales much greater than a few tens of microseconds despite intense developments in

computational hardware.<sup>1</sup> It is often the case, however, that fundamental biological processes reach timescales greater than milliseconds, one particular example being the (un)binding of specific ligand-protein complexes happening over multiple hours.<sup>2,3</sup> In order to deal with this restriction several methods have been introduced with the aim of reducing the timescale for sampling the slow biochemical processes while nonetheless being capable of recovering their original statistics.<sup>4-18</sup> Included among these are several interesting recent approaches that leverage deep learning to generate an optimum reaction coordinate (RC) that can then be used within a pre-existing enhanced sampling or markov state model framework.<sup>11-15</sup> One of these is a method we very recently proposed, named “Reweighted autoencoded variational Bayes for enhanced sampling (RAVE)”. The major distinguishing feature of RAVE is that the RC is learnt together with its Boltzmann probability distribution,<sup>18</sup> which can then serve as the ideal bias potential and be leveraged outside pre-existing biasing frameworks such as metadynamics or umbrella sampling.<sup>4-6,10</sup> In the original proof-of-concept paper, RAVE was applied to model potentials including a fullerene-nanopocket<sup>19-23</sup> unbinding test case where it was demonstrated that sampling in simulations could indeed be enhanced with the simultaneous on-the-fly learning of the RC and bias potential.<sup>18</sup> We demonstrated that RAVE could reproduce the dissociation free energy profile for the unbinding of a fullerene from a nanopocket in much less computational time than using the popular metadynamics and umbrella sampling methods.<sup>24</sup> These initial investigations were thus suggestive that RAVE could find important applications in simulations of ligand-protein complexes that now have a critical role in aiding drug design.<sup>25,26</sup>

Here we introduce a simple but powerful multi-dimensional extension to RAVE that makes it possible to obtain accurate absolute binding free energies  $\Delta G_b$  in realistic ligand-protein complexes. As our test-case, we choose benzene (un)binding from the L99A variant of the T4 bacteriophage lysozyme protein (T4L), which is a popular and challenging ligand-

protein complex for experimental and simulation studies.<sup>25,27-35</sup> For example, Deng and Roux investigated the binding affinities of the T4L protein with various aromatic ligands.<sup>25,30</sup> Miao and co-workers applied their Gaussian Accelerated Molecular Dynamics (GaMD) method on the complex.<sup>31</sup> Wang et al.<sup>33</sup> also looked into the binding free energies but from the perspective of using association and dissociation rate coefficients to calculate the  $\Delta G_b$ . Most recent, several (un)binding paths were simulated in order to make accurate kinetics predictions,<sup>34,35</sup> these (un)binding paths differing from one another in that different helix-helix distances were modulating the entry and exit of benzene from the binding pocket.

Using our multi-dimensional extension, we will show how RAVE can learn position-dependent RCs which lead to sampling benzene unbinding from the buried binding pocket of the T4L protein in 100% of our short independent MD simulations (that is 20 out of 20 simulations). The first unbinding event in these short independent runs occurred within 3-50 nanoseconds, corresponding to a seven order of magnitude speed-up relative to the actual process expected to take around 100 milliseconds. Furthermore, back and forth movement between the deep initial bound state and intermediate metastable states within the binding pocket was often observed prior to unbinding. This rate of success and speed-up, together with the reversible nature of the sampling within the buried binding pocket and elsewhere inside the protein, are to our knowledge unprecedented in systems that are this complex. In addition, we also show that RAVE can handle the calculation of binding free energies quite well. Multiple runs gave estimates of the binding free energy within reasonable error bars in agreement with other sampling methods. Lastly, using our simulations we can construct well-converged free energy profiles demonstrating the interdependence between ligand movement, protein fluctuation and water movement. It is our expectation that RAVE, especially with its multi-dimensional extension introduced here, will find use in automating the calculation of important quantities which are of practical in-

terest.

## 2 THEORY

### 2.1 RAVE

RAVE has been introduced in detail in the original proof-of-concept publication<sup>18</sup> and here we summarize its central features. To begin let us assume a molecular system with  $N$  atoms at temperature  $T$ , and under some other generic thermodynamic conditions. Our central objective is to sample the system's Boltzmann-weighted probability distribution using all-atom MD. To achieve this objective, the first step in RAVE is to launch an unbiased MD simulation from a point in configuration space that often, for ligand-protein unbinding processes, will correspond to the ground state if known or to a metastable state. Unless  $T$  is high, the simulation will be trapped in this state sampling the fast internal degrees of freedom as well as the a priori unknown RC, both according to their Boltzmann-weighted probabilities. There is useful information contained in this trapped simulation that can be leveraged to enhance fluctuations. The slow degree-of-freedom describing escape from the metastable state, for instance, stands apart as a distinct signal or feature relative to the fast internal oscillations that together amount to background noise. The probability distribution  $P$  of the simulation data, when projected onto the slow coordinate, provides a natural biasing potential that helps enhance the fluctuations by making motion along the RC more diffusive. Our central motivation in constructing RAVE was that both the slow degree-of-freedom as well as its probability distribution  $P$  are learnable via the use of certain forms of deep learning.<sup>18</sup> RAVE will then proceed from this initial unbiased simulation stage in an iterative fashion, each iteration aiming to simultaneously construct, from what will then be a biased simulation, a better biasing potential along a more refined RC. In this sense, the biasing along the approximate RC obtained from deep learning drives forward the enhanced exploration of the system's con-

figuration space, leading to the generation of new data that we can again analyze with deep learning. The success of this data generation scheme will depend on the RAVE protocol determining an appropriate RC to describe the slow coordinate, but the iterative protocol per construction provides a self-consistent check to help construct the correct solution. One can just screen through the set of available RAVE RCs, including any spurious solutions, since an appropriate RC should lead to systematically greater exploration of configuration space until ergodicity is achieved. It is the method's inherent new data generation scheme together with its simultaneous learning of both the RC as well as the bias potential that makes RAVE distinct relative to other deep learning based methods.<sup>12-15</sup>

Let us formalize this intuitive description. To start RAVE the user first defines a  $k$ -dimensional vector  $\mathbf{s}$  whose components are  $k$  order parameters  $(s_1, s_2, \dots, s_k)$  expected to have an important role in the process being studied. These are functions of the  $3N$  atomic coordinates  $\mathbf{x}$ ,  $s_i = s_i(\mathbf{x})$  where  $i = 1, 2, \dots, k$ , and represent a valid reduced dimensional description. One should think of a good order parameter vector  $\mathbf{s}$  as containing components of a basis set that together can be combined into a RC that describes the slow degree-of-freedom associated with the process. In the case of ligand-protein complexes, which are the subject of this work, the vector components are distances between the ligand center-of-mass and protein residues, protein inter-residue distances and protein residue hydration states.<sup>36,37</sup> It is standard practice in the enhanced sampling literature to introduce such order parameter descriptions.<sup>16,36-41</sup> Although the need to pre-select a set of order parameters might seem like a drawback of RAVE and these other methods,<sup>16,38-40</sup> in this work we will demonstrate how RAVE allows us to expand an initial minimal list of order parameters with a self-consistent test until the set of order parameters required to construct the RC is complete. This is akin to starting with a small basis set in quantum mechanics, and gradually expanding the list if needed.

With the order parameters in hand we can launch a brief unbiased molecular simulation from the ground state configuration. This generates a time-series ( $\mathbf{s}^1, \mathbf{s}^2, \dots, \mathbf{s}^n$ ) tracking the time evolution of  $\mathbf{s}$ , where  $n$  is the total number of timesteps in the simulation, and  $\mathbf{s}^i$  denotes values of the order parameter vector at time-step  $i$ . This time-series dataset can then be used to extract, *at the same time*, both (a) a latent variable  $z$  that describes a low-dimensional manifold capturing the interesting features of the data and (b) probability distribution  $P(z)$  of the data along this variable  $z$ . Since for our purposes the task of capturing the interesting features amounts to finding a low-dimensional representation of the data that disentangles the slow molecular motions from the rovibrational oscillations amounting to random noise, RAVE uses an unsupervised machine learning approach called variational autoencoder (VAE)<sup>42-44</sup> whose training protocol per construction compresses high-dimensional data into low-dimensional representations. A crucial point worth noting is that in traditional implementations of the VAE,  $z$  is described as tens of thousands of neural network parameters lacking clear interpretability. RAVE attempts to solve the interpretability problem by shifting the emphasis to the probability distribution  $P(z)$  of the latent variable, rather than the exact variable itself. By screening through trial RCs  $\chi$  expressed as linear combination of the order parameters ( $s_1, s_2, \dots, s_k$ ), RAVE determines the best RC as the  $\chi$  whose probability distribution  $P(\chi)$  is closest to the probability distribution learnt in VAE, namely  $P(z)$ , according to the Kullback-Leibler (KL) divergence metric:

$$\mathcal{D}_{KL}(P(z)||P(\chi)) = \sum_i P(z_i) \log \frac{P(z_i)}{P(\chi_i)} \quad (1)$$

Eq. (1) sums over the discretized distributions containing the same number of bins. In the limit when  $P(\chi)$  approaches  $P(z)$ , the KL divergence will tend to zero. The distribution  $P(\chi)$  minimizing Eq. (1), in addition to determining the RC  $\chi$ , immediately leads to a bias potential,  $V_b$ , equal to the inverted free energy,

$F$ :

$$V_b(\chi) = -F(\chi) = k_B T \log P(\chi) \quad (2)$$

At this point RAVE has determined, from an unbiased simulation, both an approximate reaction coordinate as well as a bias potential, which it uses to launch a short biased MD simulation generating a new dataset ( $\mathbf{s}^1, \mathbf{s}^2, \dots, \mathbf{s}^n$ ) containing larger fluctuations in configuration space, assuming that the RC identified indeed has sufficient overlap with the true RC. From this new dataset we can then extract a more accurate (in terms of sampling the tails) distribution  $P$  along a refined hidden latent variable  $z$  whose unbiased probability distribution we compare to that of trial RCs  $\chi$  again expressed as a linear combination of order parameters ( $s_1, s_2, \dots, s_k$ ). In order to reweight out the effect of the biasing, we calculate the KL Divergence as:

$$\mathcal{D}_{KL}(P(z)||P(\chi)) = \sum_i P^u(z_i) \log \frac{P^u(z_i)}{P^u(\chi_i)} \quad (3)$$

where  $P^u(z)$  and  $P^u(\chi)$  are the unbiased probabilities reweighted from a biased MD simulation. For this RAVE associates a weight  $w = e^{\beta V_b}$  to each point sampled during the biased simulation, where  $\beta = \frac{1}{k_B T}$  is the inverse temperature. This can then be used to recover the unbiased probabilities through the simple reweighting formula from importance sampling:

$$P^u(\chi) = \frac{\langle w \delta(\chi - \chi(t)) \rangle_b}{\langle w \rangle_b} \quad (4)$$

In Eq. 4 the subscript  $b$  indicates averaging while sampling from the biased simulation. From this point onwards all that is left is to launch another molecular simulation using the current biasing parameters and iterate between rounds of MD and VAE until the desired thermodynamic variables as well as the RC are converged.

## 2.2 Multi-Dimensional RAVE and the “washing out” trick

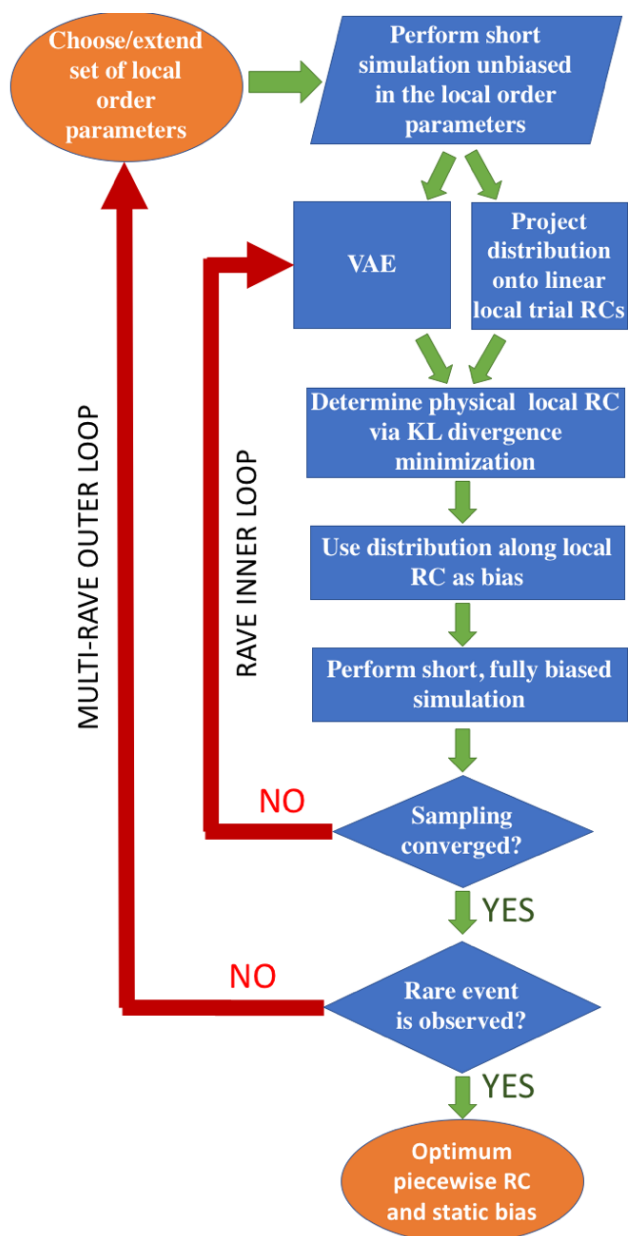
The original RAVE protocol was designed to determine a one-dimensional RC given a set of user-defined input order parameters.<sup>18</sup> Furthermore, the one-dimensional RC was restricted to be a linear combination of these order parameters, as described in Sec. 2.1. It can be impractical, however, to use just a single linear RC for problems such as ligand-protein unbinding, where the reaction path involves movements between multiple metastable states. Of course, in principle, we could extend the space of available RCs to include non-linear combinations of the chosen order parameters. These non-linear RCs together with the associated bias could handle realistic biochemical problems of arbitrary complexity. Unfortunately, the use of non-linear functions is not without their own complications. For instance, when using Eqs. (1) and (3), different non-linear combinations can lead to almost indiscernible values of KL divergence, increasing the likelihood of finding spurious solutions for the optimum RC and bias. Furthermore, the natural limit of taking non-linear combinations of order parameters is to directly use the deep neural network VAE bottleneck variable itself as the RC, similar to what has been done in recent autoencoder based work.<sup>12,13,45</sup> However, this at odds with our key intention of maintaining physical interpretability of the RC.

Here we develop an alternative approach based on representing the RC as a combination of multiple position-dependent piecewise linear functions. Although the true RC is not embedded in a linear subspace spanned from the user-defined input order parameters, we divide the sampled regions of configuration space into several sections such that within each section the local component of the RC is embedded in a linear subspace of the order parameters. Each linear RC component and its associated Boltzmann probability distribution can then be identified using the original RAVE protocol in one of two different procedures. These can be either (a) in a supervised piecewise manner where the network is trained only using data from spe-

cific parts of the configuration space, or (b) in a relatively unsupervised manner where the full dataset can be used, but with a slight modification that we introduce next. This allows enforcing that each local linear component is optimized strictly from only the local features of the configuration space. We adopt a minimalistic approach here, where each subsequent local component is introduced only when no further enhancement in ergodicity can be achieved with the components at hand. Notice that each local component that is introduced is free to learn a different direction in the configuration space of our molecular system such that biases along them act about multiple and ideally independent dimensions. It is for this reason that we call this extension of the protocol multi-dimensional RAVE (multi-RAVE). We provide a flowchart depiction of multi-RAVE in Fig. 1.

We now formalize this intuitive picture. To initiate multi-RAVE we first launch a short unbiased molecular simulation from an initial configuration that corresponds to the ground state configuration. Using the dataset containing the time evolution of  $\mathbf{s}$  we now subject the simulation data to the original RAVE protocol under the specific constraint that the first RC component, which we call  $\chi_1$ , be a linear combination of the user-input order parameters  $\mathbf{s}$ , i.e.  $\chi_1 = \mathbf{c} \cdot \mathbf{s} = c_1 s_1 + c_2 s_2 + \dots + c_k s_k$ . Nothing until now has distinguished multi-RAVE from the original protocol, since in essence we are performing RAVE on a single linear reaction coordinate. As described in Sec. 2.1 and in the original publication,<sup>18</sup> RAVE will proceed to iterate between rounds of MD and VAE in order to improve the choice of the coefficients  $\mathbf{c}$  and the associated unbiased probability distribution. Should the optimized one-dimensional RC  $\chi_1$  lead to sampling the desired rare event, then nothing else remains to be done. If, however, after initial enhancement in ergodicity the coefficients converge but the system still does not exhibit the desired rare event we set out to study, we proceed to introduce a new round of RAVE on potentially new order parameters.

Essentially, the first set of RAVE has brought the system to an intermediate where either (a) new order parameters are needed, and/or (b)



**Figure 1:** Flowchart illustrating the multi-RAVE protocol. The portions of the flowchart in blue represent the original RAVE protocol while the outermost loop represents the multi-dimensional extension introduced in this work.

previous order parameters suffice, but the RC undergoes a change which cannot be captured by the constraints of a linear formalism, which is the case with the benzene-T4L complex we will focus on. See Figs. 2b and 2c for an illustrative schematic of case (b). To tackle this issue, multi-RAVE now introduces a second linear RC component,  $\chi_2$ , that will be optimized within a region of configuration space not sampled during biased MD simulations with RC  $\chi_1$  and associated bias  $V_b(\chi_1)$ . To begin, however, the user will first redefine the vector  $\mathbf{s}$  whose components are now the  $d$  order parameters considered useful in constructing this new local linear RC component. The purpose of this redefinition is to capture a direction not expressible through the original  $k$ -dimensional subspace from which we optimized  $\chi_1$  but that describes the local behavior of the slow degree of freedom. The procedure for choosing the  $d$  new order parameters will be as usual context specific. For ligand-protein unbinding, which is the interest of this publication, these are mostly the distances between the ligand and center-of-mass and protein residues. Multi-RAVE now launches an MD simulation that is biased along  $\chi_1$  using the current converged estimate of  $V_b(\chi_1)$  and tracks the time evolution of the redefined order parameter vector  $\mathbf{s}$ .

The multi-RAVE protocol for optimizing the second set of biasing parameters  $\chi_2$  and  $V_b(\chi_2)$  is similar to performing the original RAVE protocol but with a simple, essential difference that we label as the “washing out trick”. It is our intention to learn, through  $\chi_2$ , the features that have not already been captured by  $\chi_1$ . We thus need a computationally easy way to turn off the features that we have already captured. Multi-RAVE implements this task in an automated manner by optimizing  $\chi_2$  without reweighting for the effect of the bias  $V_b(\chi_1)$  along  $\chi_1$ , even though that sampling was performed using a bias. Effectively this amounts to tempering the probability distribution along  $\chi_1$  so that any associated features are washed out. The net effect of this procedure is that the sampling incorporated into  $\chi_1$  appears featureless so that, during the optimization of  $\chi_2$  and  $V_b(\chi_2)$ , there is nothing to be learned from previously explored

configuration space regions.


Let us be more concrete about the protocol for learning  $\chi_2$  and  $V_b(\chi_2)$ . As we have mentioned multi-RAVE generates a time series  $(\mathbf{s}^1, \mathbf{s}^2, \dots, \mathbf{s}^n)$  using the converged estimate of  $V_b(\chi_1)$ . With the VAE we extract  $z$  and  $P(z)$  that will be our benchmark for screening through  $\chi_2 = \mathbf{c} \cdot \mathbf{s} = c_1 s_1 + c_2 s_2 + \dots + c_d s_d$  and corresponding  $P(\chi_2)$  via the KL divergence metric in Eq. (1). If we wished to ensure that the probabilities were unbiased we would apply the reweighting formula:

$$P(\chi_2) = \frac{\langle w_1 \delta(\chi_2 - \chi_2(t)) \rangle_b}{\langle w_1 \rangle_b} \quad (5)$$


where in Eq. (5)  $w_1 = e^{\beta V_b(\chi_1)}$ . Multi-RAVE simply sets  $w_1 = 1$  to turn off the reweighting and effectively wash out any features learnt so far (see Fig. 2a vs. 2c). We thus arrive at a component of the RC  $\chi_2$  and its probability distribution  $P(\chi_2)$  capturing the features of the configuration space previously uncaptured by  $\chi_1$ . These directly give us a second set of biasing parameters  $\chi_2$  and  $V_b(\chi_2)$  that are optimum given the quality of the sampling so far. We continue by iterating between the rounds of MD and VAE as per the RAVE protocol making sure to apply the washing out trick to the first set of biasing parameters  $\chi_1$  and  $V_b(\chi_1)$  but not the second, until the second RC and associated bias also converge. Without loss of generality, the same treatment then applies to all additional components  $\chi_i$  and their associated biases  $V_b(\chi_i)$  for  $i \geq 2$ . We proceed to introduce local components until their converged biasing parameters lead to sampling the full rare event. Until this is accomplished, the protocol is to simply for a given RC-component number  $i$ , ignore all biases for components 1 to  $i-1$  by setting corresponding weights  $w_i$  to 1 while optimizing the  $i$ th component and bias.

### 2.3 Binding Free Energy Calculation

While the protocol we have described here applies to generic systems, our focus in this work is on the calculation of absolute binding free



**Figure 2:** (a) A three-centered model potential with labels for the reactant (R), intermediate (I) and product (P) states. (b) Movement from left to right between the reactant to the product via an intermediate state cannot be captured with a single linear RC. (c) The three-centered model potential when the bias acting along the first RC component is not accounted for with the reweighting formula given in Eq. 4.



**Figure 3:** (a) Snapshot of the benzene-T4L complex in the bound state configuration from which all trajectories are initiated. (b) Snapshot of the benzene-T4L complex in the intermediate metastable state from which it escapes into the solvent through the helices 6, 7, 8 and 9.

energies of ligand-protein complexes, which is a challenging and important problem.<sup>25,30</sup> Without loss of generality, we give illustrative examples for the benzene-T4L complex studied in this work and described in detail in Sec. 3. When the multi-RAVE protocol has converged in its estimate of a piecewise linear RC and associated biases, we launch independent MD simulations from the bound pose of the ligand-protein complex using these optimal biasing parameters. Unlike other autoencoder based enhanced sampling work, no additional biasing as in metadynamics or umbrella sampling is required.<sup>12,15</sup> For the benzene-T4L complex studied here, the learnt piecewise RC and associated bias led to unbinding in simulation times between 3 to 50 nanoseconds when the natural process occurs on 100 ms or slower timescales,<sup>31,33-35</sup> which is reflective of the caliber of the biasing parameters that multi-RAVE has generated.

With 20 out of 20 MD simulations sampling the benzene-T4L unbinding rare event, ligand-protein binding free energy,  $\Delta G_b$ , can then be directly calculated from the free energy profiles estimated from these runs once these have converged. In order to calculate the binding free energy we have used the potential of

mean force (PMF) along the distance between the benzene center-of-mass and Tyr88. At the benzene-T4L bound position benzene is proximal to Tyr88. We then defined the reactant state along this PMF to include all values less than 0.8 Å and the product state to be all values greater than 1.8 Å. We include a standard correction term<sup>30,33</sup> that accounts for the different concentration of benzene used in experiments versus simulations (1 M vs. 5 mM in the present work). One crucial caveat needs to be addressed for calculation of binding free energies from simulations such as ours - once the ligand leaves the protein, it is free to explore the full solvent making it extremely rare for it to bind back. In principle, we could train RAVE to address this entropic portion of the configuration space as well. However, here for computational ease, we apply a soft restraining potential that brings the ligand back into the protein once it unbinds. The effect of this restraining potential is reweighted out in the calculation of  $\Delta G_b$ . See Supplemental Information (SI) for further details of this and other calculations in this work.



## 2.4 Simulation and Neural Network Details


The MD simulations in this work have been performed using the software GROMACS version 5.0<sup>46</sup> patched with PLUMED version 2.3.<sup>47</sup> These simulations were all performed in the constant number, pressure and temperature (NPT) ensemble, with  $\sim 10,000$  water molecules in a periodic box where all side lengths are seven nanometers. The pressure of the simulation was kept at 1.0 bar and the temperature at 298 K. Constant pressure was maintained using Parrinello-Rahman barostat<sup>48</sup> while the temperature was maintained constant with the v-rescale thermostat.<sup>49</sup> In addition, the CHARMM22\* force-field<sup>50</sup> was used to describe the system. All simulations were run for  $\sim 300$  picoseconds initially, although in later RAVE rounds the MD simulation times were increased to  $\sim 2$  ns (see Sec. 3.2 for additional information).

The current VAE implementation uses a neural network architecture almost identical to the set-up used in the original RAVE publication.<sup>18</sup> The input and reconstruction spaces are 2-dimensional (i.e. 2-dimensional order parameter vectors) while the probabilistic encoder and decoder are each 3 layers of 512-dimensional vectors. The latent variable representation of the RC is 1-dimensional. The final layer of the encoder did not have an activation function, meaning it was defined as a linear transformation. The output layer of the decoder used tanh as the activation function. We refer the readers to the original RAVE publication for a schematic illustration of the VAE architecture as well as additional details regarding the architecture.<sup>18</sup>

In order to implement the aforementioned VAE architecture and to train it we have used the high level deep learning library called Keras.<sup>51</sup> The training of the neural network was performed using the RMSprop algorithm, which is a variation of stochastic gradient descent. The learning rate used with RMSprop was 0.0001 or 0.0002 depending on the size of the dataset. Training was performed over 1000-3000 epochs also depending on the size of the

dataset. In total, 34 rounds of MD-VAE were used to train the RC and the time-independent static biases along them.

## 3 RESULTS



**Figure 4:** Calculated binding free energy as a function of time. Six different trajectories, in color, are shown, while the time-dependent average value is given as the black line represents.

### 3.1 Binding Free Energy

We have performed in total 20 biased MD simulations, all of them initiated from the benzene-T4L bound configuration using the RC and bias identified from construction rounds of RAVE. 6 of these trajectories ran for longer than 100 ns in order to monitor the detailed convergence of the  $\Delta G_b$  estimate. As can be seen in Fig. 4, this estimate appears to be converging before the simulation time reaches 100 ns. We thus obtain, using 100 ns, a  $\Delta G_b$  estimate of  $-8.2 \pm 0.8$  kcal/mol, although similar results would be obtained with much a shorter simulation time (see the behavior of the time-dependent average given as the black curve in Fig. 4). This result is in good agreement with other calculations reported in the literature, for example  $-6.9 \pm 0.8$  kcal/mol from Mondal et al.<sup>34</sup>

based on long unbiased simulations and subsequent application of Markov State Modeling, and  $-5.96 \pm 0.19$  kcal/mol from the alchemical-based calculations of Deng and Roux<sup>25</sup>. Notice in Fig. 4 that all 6 trajectories underestimate these reported values suggesting that our  $\Delta G_b$  estimates might be a lower bound. It is possible that this arises from the slower rate of sampling the unbound state as opposed to the initial metastable well corresponding to the bound complex, as we have used a sub-optimal restraining potential to bring the ligand back into the protein (see SI). Whether this is true in general, however, will be the subject of future investigations.

### 3.2 Order Parameters and Reaction Coordinate

We show in Table 1 the T4L residues we used to define the distances relative to the benzene center-of-mass. Our RC components were eventually based on using just these distance-based order parameters. We also considered other order parameters such as protein inter-residue distances and protein residue hydration states, but these had close to 0 weight in the RC on a consistent basis. We will thus not describe them further here. To drive the ligand to unbind from the protein, four local linear RC components was found to be sufficient. These components were learnt sequentially using RAVE:

$$\chi_1 = -(0.17d_1 + 0.02d_2)$$

$$\chi_2 = -0.40d_3$$

$$\chi_3 = 0.04\chi_1 + 0.05d_5$$

$$\chi_4 = d_2$$

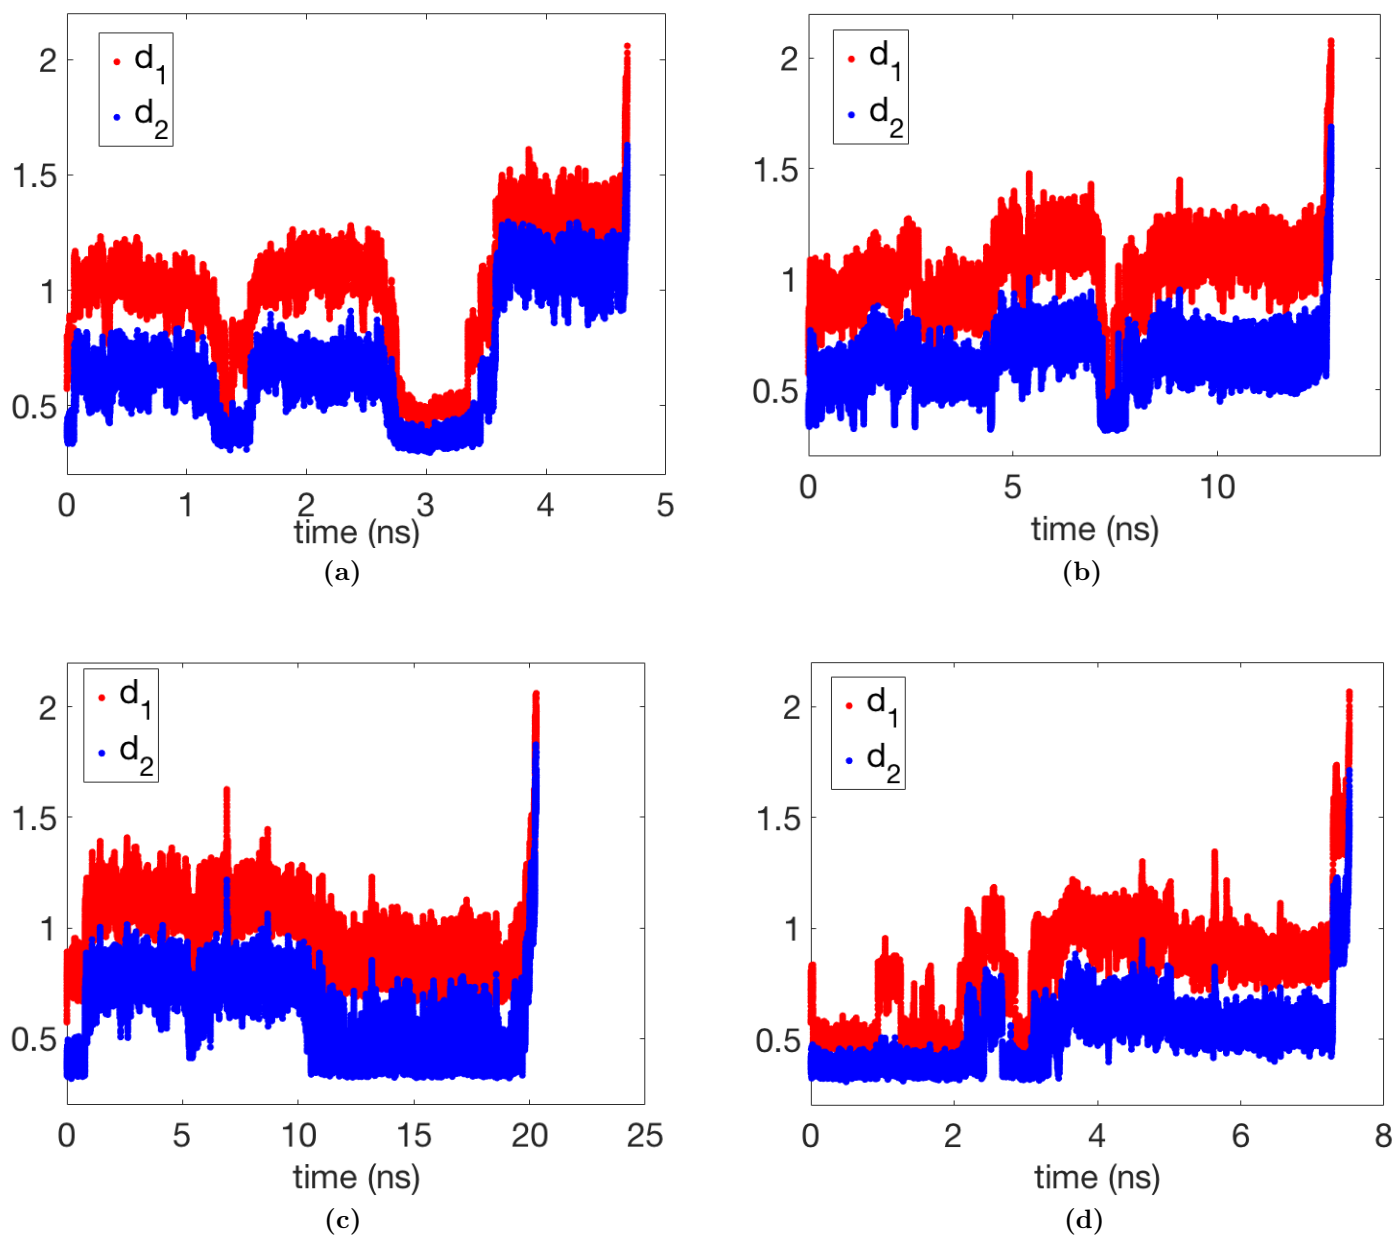
with  $d_i$ ,  $i = 1, 2, 3, 4, 5$ , defined as the distance between the center-of-mass of benzene and the protein residues according to Table 1. In terms of the snapshots in Figs. 3a and 3b,  $d_1$  measures the distance of benzene from helix 4,  $d_2$  from helix 5,  $d_3$  from helix 8,  $d_4$  from helix 7 and  $d_5$  from helix 6. Notice that  $\chi_1$  is composed of  $d_1$  and  $d_2$ , which have been defined using residues on helices that when in close contact with benzene corresponds to the initial reactant state.  $\chi_2$ , meanwhile, is defined in terms of  $d_3$  whose associated residue in close contact with benzene corresponds to intermediate metastable state.

We began learning the first of these four linear RC components using extremely short simulations of  $\sim 300$  picoseconds. These simulation times were enough to converge a first RC component but were not sufficient to learn a second RC component. For this reason we increased the simulation time for the remaining RC components to be trained to be  $\sim 2$  ns. We will investigate in future work the critical question of the simulation length which should be used to maximize the efficiency and accuracy of RAVE. Once we had learned linear RC components capable of driving exploration of configuration space from the tightly bound complex into higher energy regions and eventually into escape into the solvent we seized to continue training. Hence we did not attempt to train a RC and bias for capturing re-entry into the protein, and instead used simple restraining potentials to bring the ligand back.

As mentioned in Sec. 3.1, we ran 20 independent biased MD simulations, all of which led to benzene unbinding from the T4L buried binding site. Often, the unbinding event occurred in the first few ns of simulation although in one case benzene unbinding was observed after 50 ns. Prior to unbinding, however, we observed that benzene tends to show back and forth movement between being proximal to either helix 4 or helix 7, as shown in Fig. 5. This hopping between wells corresponds to sampling the initial bound state and an intermediate state, which suggests that RAVE can learn biasing parameters capable of reversibly sampling the reaction path until ligand escape into the solvent occurs. Taken together, these demonstrate the quality and reliability of the RC and bias so constructed.

### 3.3 Unbinding Mechanism

A range of specialized MD based enhanced sampling methods have been used to uncover several paths corresponding to benzene escape from the buried binding site of T4L.<sup>31,34,35</sup> Many of these methods have also attempted to quantify the actual dissociation and association rate constants. Here we do not attempt to quantify these constants (which will be the



**Figure 5:** (a-d) show four (out of twenty total) independent biased trajectories in terms of order parameters  $d_1$  (red) and  $d_2$  (blue). Clear back-and-forth movement between various metastable states can be seen. Also note the short but different first passage times to unbinding for different trajectories.

subject of future work). However, our biased trajectories do show clear back and forth movement between various metastable states (See Fig. 5 showing 4 out of 20 trajectories in terms of the  $d_1$  and  $d_2$  order parameters), which is a hallmark of lack of hysteresis in enhanced sampling simulations, and is typically rather difficult to achieve.<sup>10</sup> It is this easy inter-state movement which is what leads us to classify our unbinding as reversible, and gives us confidence to draw mechanistic conclusions on the basis of these trajectories. Our optimized piecewise linear RC and associated bias led to eventual benzene escape from the binding pocket and into the solvent through helices 6, 7, 8 and 9, via a metastable intermediate state corresponding to the benzene being in simultaneous close contact with helices 7 and 8. Figs. 3a and 3b show snapshots of the benzene-T4L bound pose and intermediate metastable state. The presence of these two states is also shown in Figs. 6a–6c. It is interesting to note that movement from the benzene-T4L bound position to the intermediate is correlated to slight increases in the “transient motions” of helices 7 and 8. The relevance of different short-lived protein “breathing motions” for making the deeply buried binding site accessible to ligands has been pointed by various other theoretical studies.<sup>32,34,35,52</sup> Escape of the benzene into the solvent appears to require these slight increases in helix-helix distance, as shown in Fig. 6b. A movie of the unbinding simulation showing these events is provided in the SI.

Interestingly, our initial benzene-T4L complex has two short-lived water molecules in the non-polar cavity between helices 8 and 10. The timescale for water exit tends to be much faster than benzene exit from the binding pocket, found to be less than 100 ns in our unbiased MD simulations. Similar waters have been reported in experiments for many variants of T4L and also from simulations.<sup>52</sup> As shown in Fig. 6c, the movement from the intermediate to the outside of the binding pocket occurs post-water exit. Sampling of the intermediate metastable state does take place while one water molecule either remains close to its initial position or enters the binding pocket before eventual escape

into the solvent. The exit of water molecules from the binding pocket is thus a relatively fast, but mandatory event for ligand unbinding to occur.


## 4 DISCUSSION

Over the past decade a tremendous number of machine learning approaches relevant to various aspects of molecular simulations have become available in the literature.<sup>12–15,53–58</sup> A major open problem in the field has been whether machine learning can be leveraged to make rare events less rare, therefore making amenable their accurate sampling via standard computational resources. Our recent publication marked a promising step forward in this direction.<sup>18</sup> Through the use of an iterative deep learning-MD scheme we demonstrated that significant enhancement in the sampling of model potentials could be achieved. In the current work we have extended the applicability of RAVE by showing how it can be used to simultaneously learn the reaction coordinate and calculate the absolute binding free energy in a much more challenging test case, namely the benzene-T4L system in explicit water. A simple but crucial methodological extension of RAVE has been introduced here, named multi-RAVE, that allows learning a RC with multiple components and explicit dependence on location in configuration space. It is this position dependence that allows multi-RAVE to construct an overall non-linear RC as a sum of piecewise linear components that nonetheless captures the slow non-linear degree of freedom, with the added benefit of allowing for clear physical interpretability.

RAVE shares some similarities with Diffusion Map-directed MD<sup>59</sup> and intrinsic map dynamics<sup>60</sup> which also perform sampling without pre-knowledge of low dimensional RC. Similar to RAVE, in these methods as well the exploration of the unknown configuration space is broken down into local components. In Diffusion Map-direction MD, for instance, localized slow diffusion coordinates are determined from an MD simulation. Intrinsic map dynam-

Table 1: T4L protein residues used to define the order parameter distances, relative to the benzene center-of-mass, used in the construction of the piecewise linear RC.

d1	d2	d3	d4	d5
Tyr88	Ala99	Leu133	Leu118	Val111



**Figure 6:** Contour plots of free energies, in kcal/mol, showing the escape pathway when projected onto different two-dimensional order parameters. The helix numbering follows the convention given in Fig. 3a. (a) Free energy profile given as a function of the main components of  $\chi_1$  and  $\chi_3$ . (b) Free energy profile given as a function of distance from binding site and helix7-helix8 breathing motion. (c) Free energy profile given as a function of distance from binding site and hydration state between helices 8 and 10. (d) Free energy profile given as a function of hydration state between helices 8 and 10 and helix7-helix8 distance.

ics, meanwhile, learns local parametrizations to the low-dimensional manifold that describes the slow process of interest. Both methods launch an MD simulation from the edge of the previously explored region in order to start a new iteration. RAVE is different from these methods because although it is learning local RCs from local regions of the configuration space, it is in principle independent of launching the simulation from the boundary between two regions due to the reweighting of just the “local” order parameters. That is, the final bias potential and RC learnt in RAVE allow complete movement from any point A to point B in configuration space in extra simulations with no further training or time-dependent biasing needed.

The set of RC components and bias potentials we used to describe the entire benzene-T4L unbinding process had four members (referred to in the text as  $\{\chi_i\}_{i=1}^4$  and  $\{V_b(\chi_i)\}_{i=1}^4$ ). Although we were able to get full unbinding using all biases at the same time, we found that the simulation time to unbind became much smaller (less than 10ns on average from earlier around 50ns or higher) if we switched once between the two sets of biases depending on which part of the landscape the trajectory commits to. This requires a labeling of the metastable states, and thus is a form of supervised learning. In future work we will address the question of how to achieve full unbinding without having the need to switch biases on and off depending on which metastable state the trajectory had committed to. That we were able to accomplish unbinding with only one switch is not trivial - this was only possible due to the quality of the RC and the bias learnt from deep learning. To use parlance from umbrella sampling, this is in a sense equivalent to sampling the full energy landscape of benzene-lysozyme unbinding with only two umbrella potentials. No such reported work exists, to the best of our knowledge.

In the work presented here we have shown how our recent RAVE algorithm<sup>18</sup> can be used to investigate the unbinding of realistic ligand-protein complexes. It is interesting to note that unbinding could be accomplished using geometric distances, without having to resort to protein-protein or solvation order param-

eters, whose weights were identified by RAVE to be consistently close to 0. We have noticed that the use of different distances could in principle lead to sampling different unbinding paths. Although we leave accurate kinetics of ligand-protein unbinding for upcoming work, that RAVE seems capable of reversibly sampling at least one path reversibly and in principle different paths as well seems to suggest promise for obtaining accurate kinetics, possibly through the use of an acceleration factor approach.<sup>61–63</sup> We also hope to be able to apply RAVE to other established benchmarks in the field such as the SAMPL challenges for blind prediction of host-guest binding affinities.<sup>64</sup>

To our knowledge the current work is the first to obtain reasonable absolute  $\Delta G_b$  estimates for the benzene-T4L complex using a PMF based approach. Wang et al.<sup>33</sup> attempted to obtain such an estimate but were not able to converge  $\Delta G_b$  despite long simulation times. In addition, Miao et al. obtained one binding event and thus produced estimates subject to high error.<sup>31</sup> It is often thought of as less convenient to use free energy surface or potential of mean force (PMF) based approach, as opposed to alchemical-based approaches when the bound pose corresponds to a ligand buried deep within the binding pocket of a protein instead of the protein’s surface. In addition, a big criticism of PMF based approaches has also been their sensitivity to the chosen RC to perform the sampling along. RAVE addresses all of these issues – it learns a RC on-the-fly and gives accurate free energy estimates in nominal simulation time.

An open-source software implementing RAVE in its original form as well as the multi-dimensional extension introduced here, will soon be released for use by the wider community. We hope the current work makes a strong case for the use of RAVE as a method that uses deep learning to construct both the reaction coordinate and associated free energy profile in complex molecular systems, a problem that has been one of the holy grails of the field.

## 5 ACKNOWLEDGMENTS

We would like to thank both Yong Wang and Kresten Lindorff-Larsen for sharing their GROMACS input files for the L99A lysozyme-benzene complex. We would also like to thank Pablo Bravo for his essential work in helping develop the original RAVE protocol. We thank Netnaset Woldegerima for her thoughtful insights as well both Yihang Wang, Zachary Smith and Freddy Alexis Cisneros for their careful reading of this manuscript. We also thank Deepthought2, MARCC and XSEDE (projects CHE180007P and CHE180027P) for providing the computational resources used to perform this work. PT would like to thank University of Maryland Graduate School for financial support through the Research and Scholarship Award (RASA).

## References

- (1) Shaw, D. E.; Maragakis, P.; Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Eastwood, M. P.; Bank, J. A.; Jumper, J. M.; Salmon, J. K.; Shan, Y.; Wriggers, W. Science **2010**, *330*, 341–346.
- (2) Copeland, R. A.; Pompliano, D. L.; Meek, T. D. Nature reviews Drug discovery **2006**, *5*, 730.
- (3) Copeland, R. A. Nature Reviews Drug Discovery **2016**, *15*, 87.
- (4) Laio, A.; Parrinello, M. Proceedings of the National Academy of Sciences **2002**, *99*, 12562–12566.
- (5) Torrie, G.; Valleau, J. Journal of Computational Physics **1977**, *23*, 187 – 199.
- (6) Bartels, C.; Karplus, M. Journal of Computational Chemistry **1997**, *18*, 1450–1462.
- (7) Tiwary, P.; van de Walle, A. Physical Review B **2011**, *84*, 100301.
- (8) Zheng, L.; Chen, M.; Yang, W. Proceedings of the National Academy of Sciences **2008**, *105*, 20227–20232.
- (9) Pfaendtner, J.; Bonomi, M. Journal of chemical theory and computation **2015**, *11*, 5062–5067.
- (10) Valsson, O.; Tiwary, P.; Parrinello, M. Annual review of physical chemistry **2016**, *67*, 159–184.
- (11) Galvelis, R.; Sugita, Y. Journal of chemical theory and computation **2017**, *13*, 2489–2500.
- (12) Chen, W.; Ferguson, A. L. arXiv preprint arXiv:1801.00203 **2017**,
- (13) Wehmeyer, C.; Noé, F. The Journal of Chemical Physics **2018**, *148*, 241703.
- (14) Mardt, A.; Pasquali, L.; Wu, H.; Noé, F. Nature communications **2018**, *9*, 5.
- (15) Sultan, M. M.; Wayment-Steele, H. K.; Pande, V. S. Journal of chemical theory and computation **2018**, *14*, 1887–1894.
- (16) Tiwary, P.; Berne, B. Proceedings of the National Academy of Sciences **2016**, 201600917.
- (17) Tiwary, P.; van de Walle, A. Multiscale Materials Modeling for Nanomechanics; Springer, 2016; pp 195–221.
- (18) Ribeiro, J. M. L.; Bravo, P.; Wang, Y.; Tiwary, P. The Journal of Chemical Physics **2018**, *149*, 072301.
- (19) Baron, R.; Setny, P.; McCammon, J. A. Journal of the American Chemical Society **2010**, *132*, 12091–12097.
- (20) Morrone, J. A.; Li, J.; Berne, B. J. The Journal of Physical Chemistry B **2011**, *116*, 378–389.
- (21) Li, J.; Morrone, J. A.; Berne, B. The Journal of Physical Chemistry B **2012**, *116*, 11537–11544.

- (22) Mondal, J.; Morrone, J. A.; Berne, B. Proceedings of the National Academy of Sciences **2013**, 110, 13277–13282.
- (23) Setny, P.; Baron, R.; Kekeneshuskey, P. M.; McCammon, J. A.; Dzubiella, J. Proceedings of the National Academy of Sciences **2013**, 110, 1197–1202.
- (24) Tiwary, P.; Mondal, J.; Morrone, J. A.; Berne, B. Proceedings of the National Academy of Sciences **2015**, 112, 12015–12019.
- (25) Deng, Y.; Roux, B. The Journal of Physical Chemistry B **2009**, 113, 2234–2246.
- (26) Casasnovas, R.; Limongelli, V.; Tiwary, P.; Carloni, P.; Parrinello, M. Journal of the American Chemical Society **2017**, 139, 4780–4788.
- (27) Morton, A.; Baase, W. A.; Matthews, B. W. Biochemistry **1995**, 34, 8564–8575.
- (28) Morton, A.; Matthews, B. W. Biochemistry **1995**, 34, 8576–8588.
- (29) Liu, L.; Baase, W. A.; Matthews, B. W. Journal of molecular biology **2009**, 385, 595–605.
- (30) Deng, Y.; Roux, B. Journal of Chemical Theory and Computation **2006**, 2, 1255–1273.
- (31) Miao, Y.; Feher, V. A.; McCammon, J. A. Journal of chemical theory and computation **2015**, 11, 3584–3595.
- (32) Wang, Y.; Papaleo, E.; Lindorff-Larsen, K. Elife **2016**, 5, e17505.
- (33) Wang, Y.; Martins, J. M.; Lindorff-Larsen, K. Chemical science **2017**, 8, 6466–6473.
- (34) Mondal, J.; Ahalawat, N.; Pandit, S.; Kay, L. E.; Vallurupalli, P. PLoS computational biology **2018**, 14, e1006180.
- (35) Nunes-Alves, A.; Zuckerman, D. M.; Arantes, G. M. Biophysical journal **2018**, 114, 1058–1066.
- (36) Tiwary, P.; Mondal, J.; Berne, B. J. Science advances **2017**, 3, e1700014.
- (37) Tiwary, P. J. Phys. Chem. B **2017**, 121, 10841–10849.
- (38) Best, R. B.; Hummer, G. Proceedings of the National Academy of Sciences **2005**, 102, 6732–6737.
- (39) Pérez-Hernández, G.; Paul, F.; Giorgino, T.; De Fabritiis, G.; Noé, F. The Journal of chemical physics **2013**, 139, 07B604.1.
- (40) Chen, M.; Yu, T.-Q.; Tuckerman, M. E. Proceedings of the National Academy of Sciences **2015**, 112, 3235–3240.
- (41) Tiwary, P.; Limongelli, V.; Salvalaglio, M.; Parrinello, M. Proceedings of the National Academy of Sciences **2015**, 112, E386–E391.
- (42) Kingma, D. P.; Welling, M. arXiv preprint arXiv:1312.6114 **2013**,
- (43) Rezende, D. J.; Mohamed, S.; Wierstra, D. arXiv preprint arXiv:1401.4082 **2014**,
- (44) Doersch, C. arXiv preprint arXiv:1606.05908 **2016**,
- (45) Hernández, C. X.; Wayment-Steele, H. K.; Sultan, M. M.; Husic, B. E.; Pande, V. S. Physical Review E **2018**, 97, 062412.
- (46) Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. SoftwareX **2015**, 1, 19–25.
- (47) Tribello, G. A.; Bonomi, M.; Branduardi, D.; Camilloni, C.; Bussi, G. Comp. Phys. Comm. **2014**, 185, 604–613.
- (48) Parrinello, M.; Rahman, A. Physical Review Letters **1980**, 45, 1196.



- (49) Bussi, G.; Donadio, D.; Parrinello, M. The Journal of chemical physics **2007**, 126, 014101.
- (50) Piana, S.; Lindorff-Larsen, K.; Shaw, D. E. Biophysical journal **2011**, 100, L47–L49.
- (51) Chollet, F. Keras.
- (52) Collins, M. D.; Hummer, G.; Quillin, M. L.; Matthews, B. W.; Gruner, S. M. Proceedings of the National Academy of Sciences **2005**, 102, 16668–16671.
- (53) Ma, A.; Dinner, A. R. J. Phys. Chem. B **2005**, 109, 6769–6779.
- (54) Geiger, P.; Dellago, C. The Journal of Chemical Physics **2013**, 139, 164105.
- (55) Häse, F.; Kreisbeck, C.; Aspuru-Guzik, A. Chemical science **2017**, 8, 8419–8426.
- (56) Kobayashi, R.; Giofré, D.; Junge, T.; Ceriotti, M.; Curtin, W. A. Phys. Rev. Materials **2017**, 1, 053604.
- (57) Nguyen, T. T.; Székely, E.; Imbalzano, G.; Behler, J.; Csányi, G.; Ceriotti, M.; Götz, A. W.; Paesani, F. The Journal of Chemical Physics **2018**, 148, 241725.
- (58) Han, J.; Zhang, L.; Car, R. arXiv preprint arXiv:1707.01478 **2017**,
- (59) Zheng, W.; Rohrdanz, M. A.; Clementi, C. The journal of physical chemistry B **2013**, 117, 12769–12776.
- (60) Chiavazzo, E.; Covino, R.; Coifman, R. R.; Gear, C. W.; Georgiou, A. S.; Hummer, G.; Kevrekidis, I. G. Proceedings of the National Academy of Sciences **2017**, 114, E5494–E5503.
- (61) Grubmüller, H. Physical Review E **1995**, 52, 2893.
- (62) Voter, A. F. Physical Review Letters **1997**, 78, 3908.
- (63) Tiwary, P.; Parrinello, M. Physical review letters **2013**, 111, 230602.
- (64) Muddana, H. S.; Varnado, C. D.; Bielawski, C. W.; Urbach, A. R.; Isaacs, L.; Geballe, M. T.; Gilson, M. K. Journal of computer-aided molecular design **2012**, 26, 475–487.