Combining Monte Carlo and molecular dynamics simulations for enhanced binding free energy estimation through Markov State Models

Joan F Gilabert^{1,4}, Oriol Gracia Carmona^{1,4}, Anders Hogner² and Victor Guallar^{1,3}

¹Barcelona Supercomputing Center, Jordi Girona 29, E-08034 Barcelona, Spain

²Medicinal Chemistry, Research and Early Development Cardiovascular, Renal and Metabolism, BioPharmaceuticals R&D, AstraZeneca, Gothenburg, Sweden

³ICREA, Passeig Lluís Companys 23, E-08010 Barcelona, Spain

⁴These authors contributed equally to this work

Abstract

We present a multi-step protocol, combining Monte Carlo and molecular dynamics simulations, for the estimation of absolute binding free energies, one of the most significant challenges in computer aided drug design. The protocol is based on an initial short enhanced Monte Carlo simulation, followed by clustering of the ligand positions, which serve to identify the most relevant states of the unbinding process. From these states extensive molecular dynamics simulations are run to estimate an equilibrium probability distribution obtained with Markov State Models, which is subsequently used to estimate the binding free energy. We tested the procedure on two different protein systems, the Plasminogen kringle domain 1 and Urokinase, each with multiple ligands, for an aggregated molecular dynamics length of 760 μ s. Our results indicate that the initial sampling of the unbinding events largely facilitates the convergence of the subsequent molecular dynamics exploration. Moreover, the protocol is capable to properly rank the set of ligands examined, albeit with a significant computational cost for the, more realistic, Urokinase complexes. Overall, this work demonstrates the usefulness of combining enhanced sampling methods with regular simulation techniques as a way to obtain more reliable binding affinity estimates.

Introduction

The precise estimation of protein-ligand binding affinities through computational methods is one of the central tasks in computer-aided drug discovery. Binding affinity is typically the first parameter to be optimized when searching for a potential drug, and having a fast and accurate in-silico method would significantly enhance drug-discovery pipelines.¹ Many different approaches have been developed to estimate binding free energies, which are typically classified according to a tradeoff between computational cost and accuracy. The gold standard

of such approaches is the Free Energy Perturbation (FEP) technique²⁻⁴. FEP strategies take advantage of thermodynamic cycles to estimate binding free energies using non-physical transformations, circumventing free energy barriers and accelerating the calculations. They are limited, however, in the size of the perturbation that can be introduced, despite recent work to overcome this problem.⁵ Nonetheless, it has been shown⁶ that, when applying FEP to systems with small perturbations (small fragments to a ligand core) a limit of 1 kcal/mol accuracy can be reached, a level of accuracy that can positively impact drug development pipelines.⁷

When dealing with ligands with a significantly different structure, pathway-based approaches become useful. These methods, however, are computationally more expensive, since they require simulating the whole binding and unbinding process. For this reason, these techniques often use enhanced sampling techniques such as metadynamics,⁸ umbrella sampling,⁹ or steeredMD;¹⁰ many more methods exist, but a comprehensive review is out of the scope of this paper. The output of these simulations is typically analyzed using a potential of mean force (PMF), which is the projection of the complete free energy surface onto a few variables, thus simplifying the interpretation of the results.¹¹

While pathway-based methodologies can work with significantly different ligands, their most important drawback is the need for extensive sampling, which, together with the advent of highly parallel hardware (typically supercomputing facilities), has further popularized the use of ensembles of short simulations in place of a reduced number of long trajectories. The analysis of such ensembles is often done using Markov State Models (MSMs).^{12,13} MSMs allow the recovery of both kinetic and thermodynamic information by decomposing the conformational space into discrete states and estimating the transition probability between them. From the transition matrix, one can then extract the equilibrium probability of each state, which extends the timescales described beyond the length of any of the particular simulations used. One can find in the literature several approaches that exploit Markov State modelling to obtain kinetics measurements.^{14,15} However, the estimation of binding free energies using such approaches is not commonly described. Mainly, one can find the pioneering work by Buch et. al.,^{16,17} where they used almost 500 trajectories of unbiased molecular dynamics and constructed an MSM to estimate both kinetic and binding free energy measurements for the trypsin-benzamidine complex. In addition, Bernetti et. al.¹⁸ combined MSMs with Path Metadynamics to describe the association of alprenolol to β -AR as well as accurately predict its binding free energy. The most common applications of pathway-based methods appears to be in the study of protein-ligand association, for example using well-tempered metadynamics to study the binding to GPCRs,¹⁹ or bias-exchange metadynamics to study the binding to HIV-1 protease.²⁰ The potential of such approaches to investigate association mechanisms in a single protein-ligand is clear, however studies that involve a larger benchmark are difficult to find. A funnel metadynamics study estimated the binding affinity of 2 ligands to the protein human peroxiredoxin 5.²¹ Furthermore, a manuscript demonstrating the estimation of the binding free energy of six fragments to the protein FKBP using long equilibrium simulations has also been published.²² To the best of our knowledge, no work has been done where absolute binding free energies are estimated for different receptors and multiple ligands using unbiased simulations. This lack of additional

studies points out the difficulty in terms of sampling and computational power required to produce such results.

In a previous work,²³ we introduced a protocol which combined a short enhanced sampling simulation with longer Monte Carlo (MC) simulations using PELE.²⁴ The automated protocol provided an estimation of binding free energies of structurally diverse ligands in only a few hours of wall clock time (3 to 8 hours, depending on the target studied and when using 128 computing cores). The enhanced sampling step, which forced the ligand escape from an initial docked structure providing initial seeds for the production run, was critical in achieving such an improved level of sampling. In this work, we explore if similar results can be observed when using MD simulations for the production runs instead of PELE. Do we observe a significant enhancement of the performance? And if so, does it allow for computing several ligands in a reasonable amount of computing time (for example, when using ~1-5 μ s per ligand). To this end we present an extension of the recently developed AdaptivePELE,²⁵ allowing PELE and MD simulations to be run using the same code. We report partial success in the binding energy predictions, as well as a detailed analysis of the failures that we believe will be useful to the community of binding free energy calculations practitioners. Overall, converging MD absolute binding free energies still represents an excessive computational cost to be used in routine drug design projects.

Methods

PELE and adaptivePELE

The protocol developed in this work is composed of subsequent runs of MC and MD methods. MC simulations use the Protein Energy Landscape Exploration (PELE), which is a technique developed for the sampling of protein-ligand interactions. Each PELE MC step consists of a random translation and rotation of the ligand, followed by backbone protein perturbation using normal modes, side chain prediction and global minimization techniques. The final result is accepted or rejected using a Metropolis criterion where energy calculations use the OPLS2005 force field with an implicit solvent (currently the VDGBNP²⁶ or OBC²⁷, please refer to our recent book chapter for a comprehensive explanation of the method²⁸). In this work, PELE is called via the AdaptivePELE enhanced sampling method. In an AdaptivePELE simulation, several iterations are run, consisting of a short PELE simulation, then the output is clustered and some of the clusters are selected to start the next epoch. Clustering of all accepted MC steps is done with the leader algorithm²⁹ using the ligand RMSD (after protein superposition). Selection of the following initial structures is then based on different criteria, such as the clusters' population or the value of some metric such as SASA (solvent accessible surface area) or interaction energy. AdaptivePELE showed an order of magnitude speed-up with respect to PELE simulations, allowing us to map complex unbiased binding mechanism in less than one hour.²⁵

Molecular dynamics platform

All MD simulations were run using an extension to the previously published AdaptivePELE methodology. This new platform is aimed to simplify the use of MD simulation by leveraging the

well-designed OpenMM³⁰ API and automatizing the process of input preparation, parametrization and equilibration. First, the input structure (or multiple structures if required) are processed to fix common problems such as correcting alternative positions or identify disulphide bonds. Currently, only Amber force fields are used, by default Amber14³¹ for protein, and the ligands are parameterized with GAFF³² using the Ambertools suite.³³ After the system has been fully parameterized, it is solvated using a cubic box with TIP3P³⁴ waters, and ions are added to neutralize, also using the Ambertools suite. Input structures are then minimized for 2000 steps of L-BFGS optimization algorithm implemented by OpenMM, followed by 400 ps of NVT equilibration at 300 K with restraints to the protein and ligand heavy atoms of 5 kcal mol⁻¹ Å⁻², and finally by 4ns of NPT equilibration at 300 K and 1 atm with restraints to protein alpha carbons and ligand heavy atoms of 0.5 kcal mol⁻¹ Å⁻². In all simulations a 2 fs time step is used. The nonbonded interactions are treated using the particle mesh Ewald³⁵ method with a 9 Å cutoff and the length of the bonds containing a hydrogen atom is constrained. The values provided here are the defaults established in the protocol, although a more advanced user can specify different quantities for the strength of constraints, minimization steps, equilibration length or temperature. Finally, the production simulations are run, either using CPUs or GPUs. The software is prepared to be run in parallel both in shared-memory architectures and in distributed memory architectures, which for example allowed us to produce on the order of 800ns of simulation time per node in the Power9-NVIDIA V100 heterogeneous cluster belonging to the MareNostrum supercomputer for the URO receptor. For a more detailed description, one can check the repository where the source code of the project is stored.³⁶

Binding free energy protocol description

The protocol for the calculation of binding free energies is based on a previous study,²³ one should refer to said publication for extra details. In short, the protocol consists of five steps (summarized in Figure 1):

Figure 1. Schematic summary of the binding free energy protocol. The top left image depicts a summary of the ligand exit simulation, with structures colored from red to blue as the simulation advances. The top right image depicts the center of mass of the selected initial structures from the starting simulations in blue dots and the simulation box in grey. The bottom left image shows a schematic representation of the MSM clusters colored according to their equilibrium probability

1. **Initial MC equilibration.** The protocol starts with a short equilibration, ~ 500 PELE steps, to remove initial small structural inaccuracies caused by the crystallographic model or the docking procedure. The equilibrated structures are clustered in 10 groups, and each cluster center is used for the next step.

2. **Ligand exit simulation.** Short AdaptivePELE simulations are then used to map the ligand unbinding process. To reduce even further the computational time required for this step, we use a biasing reward function in the spawning step, where 75% of the cores are assigned to structures with the largest SASA and the rest to the less-visited clusters. In Figure 2A, a plot from a ligand exit simulation is shown. There, it can be seen how the SASA of the ligand gradually increases and with roughly 10 PELE steps a fully solvated state is reached (more challenging systems might require from 10 to 40 steps, overall less than half an hour CPU wall clock).

3. **Clustering.** In this step, the outcome of the unbinding simulation is clustered to obtain a small ensemble of structures that we assume will be representatives of the binding/unbinding of the protein-ligand complex. The ensemble will serve as the starting point for the production MD simulations, with the assumption that having multiple trajectories starting from different states of the binding process will increase the likelihood of observing rare events and therefore improve the convergence of the estimation and reduce the computational time required.

4. MD simulations. Continuing from the clustering step, a set of parallel MD simulations are run. In this work we benefit from the high parallelism of this task to run up to 40 replicas simultaneously, which allows us to obtain several independent trajectories starting from every starting state, improving, in turn, the overall sampling. In the simulations, the ligand is free to explore a region around the binding site, of spherical or cylindrical shape, which we refer to as the “simulation box”, the ligand is restrained to this volume by the use of a flat bottom potential, so the dynamics when the ligand is inside the box are not affected. This “box” (depicted in the top right image of Figure 1) lets the ligand form a bound complex with the protein as well as completely unbind (exploring the bulk solvent) while preventing it from exploring other regions of the protein, which we assume that will contribute less to the free energy of binding and hinder convergence.

5. MSM build-up and free energy estimation. The production simulations are analyzed using the typical pipeline of a Markov State Model construction.³⁷ The trajectories are post-processed by calculating the ligand’s center of mass at each snapshot, which is then used as input for K-Means clustering. The coordinates of the center of mass are used to clustering due because it is easily generalizable and is not system-dependent, but other features might result in better discretizations. Then trajectories’ snapshots are assigned to the closest cluster, and a count matrix is generated by using a sliding window approach with a lag-time τ , from which a transition matrix is estimated. Both clustering and MSM estimation are done with the PyEMMA package.^{38,39} The selection of the lag-time was based on the implied time scales plot (ITS), see Tables S7-S11. The estimated MSM provides us with an equilibrium probability distribution, π , which tells us the probability of each state after a time long enough that fluctuations of the system have relaxed. This probability allows estimating the binding free energy as described in previous works,^{16,17} where the binding free energy is obtained by dividing the partition function into two regions, bound and unbound. The unbound region is defined as the region where the resulting PMF is flat, and no protein-ligand interaction exists. As shown in reference 17, from the first eigenvector of the MSM transition matrix, a three-dimensional PMF, W , is constructed according to the equation:

$$W_{i\Box} = -k_B T \ln(\pi_i) \quad (1)$$

where the index i refers to a particular state, k_B is the Boltzmann constant, T is the temperature, which we set to 300 K, π_i is the probability of state i and V_i is the volume of the state i . The volume of each state is calculated by discretizing the space spanned by the ligand center of mass. We draw a cube using the maximum and minimum coordinates in each dimension and partition it into small cubes that are assigned to the closest cluster. The sum of the volumes of the cubes belonging to each cluster is taken to be the volume of the state. We would like to emphasize that the resulting PMF is not exactly a potential of mean force, but an approximation obtained from the MSM construction. As most previous authors, we assume that error associated with this approximation is small and still provides the versatility of tools of the MSM framework. Finally, the binding free energy, ΔG , is calculated as:

$$\Delta G = -\Delta W - k_B T \ln \left(\frac{V_b}{V^o} \right) \quad (2)$$

where V^o is the standard volume, corresponding to $V^o = 1661 \text{ \AA}^3$ and ΔW is the depth of the PMF, that is, the difference between the value of W in the unbound region and the minimum of the PMF, which is arbitrarily set to 0 in the binding site. The value of ΔW is calculated as the mean of the 10% of clusters with the highest value of W , to reduce inaccuracies in the binding free energy estimation. Finally, V_b is the volume corresponding to the bound configurations calculated as:

$$V_b = \sum_i V_{i\Box} \exp(-\beta W_{i\Box}) \quad (3)$$

To facilitate the analysis we typically project the obtained PMF onto the euclidean distance between the cluster center and the active site center, obtaining a plot such as the ones shown in Figure 2B. For a trustworthy model, one should expect the minimum of the PMF at distances close to 0 $\text{\AA}$, as well as a flat profile for larger distances (corresponding to the unbound region).

Figure 2. A) Evolution of the ligand's SASA along the initial AdaptivePELE exploration, showing the unbinding process for the plasmin system and B) illustration of a well-behaved PMF projection example

Systems used

In order to test the speed-up provided by the initial enhanced sampling simulation we have chosen two systems that have shown reasonable accuracy in previous absolute binding free energy estimations.²³ The receptors used are the following:

1. Plasminogen kringle domain 1 (plasmin) is one of the kringle domains of plasminogen, the inactive precursor of the plasmin enzyme. It contains a lysine binding site which is

crucial for the fibrinolysis action of plasmin. Plasmin has been suggested as a promising target for conditions involving heavy blood loss.⁴⁰ This protein is the simplest component of our benchmark due to the small size (only 80 residues) and exposed binding site. Despite its simplicity, it is far from trivial, as shown in previous work,⁴¹ where attempts to estimate the binding affinity using FEP and MM/GBSA proved unsuccessful. We have used the same 16 ligands used in the work mentioned before, comprising a rather short (challenging) range of experimental binding affinities between -2.1 to -3.9 kcal/mol, see Table S1 for 2D structures. For simulations, pdb id 4CIK was used, see Figure 3A for the structure of this receptor in complex with ligand L01.

2. Urokinase(URO), is a serine protease responsible for the catalysis of the transformation of plasminogen to its active form plasmin.^{42–44} This process causes degradation of the extracellular matrix, along with potential tumor migration and proliferation, which turns urokinase into a drug target. Its binding site is partially solvent exposed, but with a “tight” binding where all ligands adopt similar binding modes as evident from the different X-rays structures. The four ligands selected for this receptor, with 2D structures shown in Table S2, share a common naphthalene-1-carboxamide core having experimental binding free energies, between -6.5 and -9.2 kcal/mol,⁴⁵ and crystallographic structures. For simulations, pdb id 4FU9 was used, see Figure 3B for the structure of this receptor in complex with ligand 39L.

Figure 3. Crystal structure of the two protein systems used in this work. A) Complex of plasmin receptor with the L01 ligand and B) Complex of URO receptor with the 39L ligand

Results

Binding free energy for plasmin system

The first test of the method is the prediction of the binding affinities for a series of 16 ligands to the plasmin receptor. To do so, for each ligand we have performed 80 simulations of 200 ns, starting from 40 structures selected from the previous AdaptivePELE exit simulation; a representative ligand migration profile is shown in Figure 2A. This benchmark amounts to a total aggregate simulation time of 256 μ s (16 μ s per ligand). The simulations were clustered into 1000 states, and a MSM was built using a lag-time of 50 ns; a representative PMFs is shown in Figure 2B.

Figure 4A shows the correlation between the estimated values versus the experimental ones. To evaluate the prediction accuracy, we have calculated the root-mean-square error (RMSE), 0.64 kcal/mol, the R^2 value, of 0.343, and the Spearman correlation of 0.66. The plot shows an overall correct trend, with an approximate linear behavior, except for two groups: i) Ligand L03, which is overestimated, and ii) the group of 4 weakest binders, which, while being mostly correctly ranked inter-group, also show an overestimated binding free energy. There is no obvious reason for the outlying L03 ligand, the results for the second group, however, suggest that there is a lower limit in the resolution provided by the methodology, which seems to fail for the ligands with free energies below the -3 kcal/mol mark. As L03 seems to be an outlier, we recalculated the three metrics after removing, obtaining a R^2 value of 0.375, an RMSE of 0.57 kcal/mol and a Spearman correlation of 0.63. Omitting L03 does not appear to have a significant impact on the accuracy. Overall, the performance for this benchmark system is notable, with low error, and good ranking of most of the ligands, as shown by the Spearman correlation value, even considering the small range of experimental affinities spanned by the ligands used, that makes it more difficult to obtain accurate estimations.

Figure 4. Correlation of the predicted binding free energy with respect to experimental values for the plasmin system A) with a previous PELE exploration and B) starting all simulations from solvent exposed structures, without a previous enhanced sampling simulation. The solid line represents $y=x$.

Figure 5. Analysis of the evolution of the prediction statistics (Spearman, RMSE and R^2) as the simulation length increases for the plasmin simulations. The dotted line marks the predictions using only the first 100 ns for each trajectory

The convergence evolution can be followed in Figure 5, where the R^2 , RMSE and Spearman values are plotted against the aggregated simulation time. The aggregated time is built by adding the same initial fraction in each of the 80 trajectories. Results indicate that at 100ns propagation time, dotted line in Figure 5, there is already a qualitative increase that allows a fair ranking of the ligands (notice that the x axis starts at 75 ns propagation time due to a MSM lag time of 50ns).

To further test the convergence of the methodology, we have also performed some additional replicas under slightly different conditions. First, we have run two new replicas for ligand L04 changing the initial seed for the random number generation. The goal of this test is to check the robustness of the sampling obtained. The results are shown in Table S3. All three replicas result in predictions in quantitative agreement with the experimental value, with small differences between them (less than 0.5 kcal/mol). Additionally, for the L06 ligand, we have run the whole protocol, generating a new set of initial structures, and starting the new simulation from them. The results from this additional replica (shown in Table S4) show excellent agreement between the two replicas, with less than 0.1 kcal/mol difference.

Effect of initial enhanced sampling simulation

In order to test the effect of the initial enhanced AdaptivePELE run, we have re-done the simulations shown in the previous section, starting from a single, fully solvated structure, without running a previous MC enhanced sampling simulation. Simulation parameters, such as length,

as well as MSM parameters, were kept the same. It should be noted that the comparison of this new set of simulations to the previous ones does not allow to establish the efficiency of the initial enhanced AdaptivePELE run in comparison to other seeding mechanisms, such as the ones used by Witek et al,⁴⁶ but is instead intended to analyse whether there is any gain from our approach.

Figure 4B shows the correlation between the estimated values versus the experimental values. It is readily seen when comparing to Figure 4A that the method loses all predictive power. The estimation provides an RMSE of 0.79 kcal/mol, an R^2 value of 0.178, and a Spearman correlation of 0.41. The plot shows no visible trend, with an almost uniform distribution of estimated values, and almost no ranking, as can be seen from the drop in Spearman correlation magnitude. One should note that, again, the L03 ligand results in a significantly overestimated binding affinity, which suggests that this molecule, in particular, is problematic for our setup. However, like in the previous set of simulations, omitting the L03 ligand does not change significantly the accuracy, with a R^2 of 0.141, an RMSE of 0.75 kcal/mol and a Spearman correlation of 0.36.

Binding free energy for the URO system

To further test the method capabilities in a more realistic target we applied it to a small series of 4 ligands to the urokinase-type plasminogen activator. To do so, and based on the plasmin results, we started with 100 ns simulations, and increased the number of trajectories to 120 simulations from the 40 structures selected in the previous AdaptivePELE run, thus adding 50% more simulations due to the larger system complexity. This benchmark amounts to a total aggregated simulation time of 48 μ s (12 μ s per ligand). Analogous cluster sizes and lag time were used.

In Figure 6A, the correlation of the predicted binding free energies with the experimental values is shown. It is evident from the plot that the method is not able to correctly predict the affinities for this system. The modest R^2 value of 0.454 obtained does not correctly reflect the poor predictive value since the predictions are anticorrelated. This anticorrelation is indeed reflected in the value of the Spearman correlation, which is -0.6. The calculated RMSE is 4.31 kcal/mol. The plot separates two groups, the two ligands with lower binding affinity are properly ranked, while the other two ligands are reversely ordered with respect to each other and the first two ligands. In previous works,^{23,47} we have observed how the reduced exploration volume causes the prediction to underestimate the binding free energies, although, in those works, the ranking of the ligands was accurate. We note that with only two ligands per subgroup, any discussion about ordering within said subgroups might seem spurious. However, there are marked differences between subgroups such as the ligand size or the net charge, which suggests a more detailed analysis of the simulations.

Figure 6. Correlation of the predicted binding free energy with respect to experimental values for the URO system A) with shorter simulations, B) with longer simulations and a cylindrical “box”, and C) with a spherical box and longer simulations

In Figure 7A, we plot the PMF projection for the 1UP and 2UP ligands, while in Figure 7B, the same is done for the 7UP and 39L ligands. The PMFs of ligands 1UP and 2UP represent simulations which fit into the expected energy profile, i) the most probable states are the ones more closely resembling the native state and ii) the PMF values for the solvent clusters (distance > 20 Å) are mostly flat, which is expected for states with no protein-ligand interactions. In Figure 7B, we can already see that the PMF projection for 39L introduces several alternative minima with close PMF values to that of the bound state. This trend is even worse for the 7UP ligand, where the absolute minima is far from the native bound pose, at around 10 Å. The clusters of this minimum can be separated into two groups, whose representative structures are shown in Figure S1. The two groups overlap into an almost identical pose where the ligand is bound inside a smaller cavity about 10 Å from the active site (the native pose is shown with

green colored carbons in Figure S1). The two representative structures, shown in cyan and purple in Figure S1, are flipped with respect to each other showing almost identical interactions: a hydrogen bond with Glu151 and π - π interactions between the fused rings and Tyr149 and Tyr28 (shown colored in green).

Structurally, both 39L and 7UP are large and positively charged ligands. By favouring interactions with the protein surface and having longer residence times in metastable states, the deviations from the expected PMFs could be explained. Thus, one would need more simulation time to obtain well converged binding free energies. Let's remind here that each trajectory is propagated only for 100ns, a timeframe that might be too short to escape from a metastable state. Furthermore, if the simulated time scales do not provide enough occurrences of the process to study (binding and unbinding), then the estimated affinities would be dominated by random noise, which would fit well the behavior observed in Figure 6A.

Figure 7. One dimensional projection of the PMF over the distance to the initial structure, which represents the bound conformation. In (A) ligands 1UP and 2UP (weaker binders) are shown, while in (B) ligands 39L and 7UP (stronger binders) are shown

We measured the overall sampling of transitions by counting the number of binding and unbinding events. A binding event is defined as the ligand going from distances larger than 6 Å to the reference structure to distances smaller than 2.5 Å. We consider the reverse process to be an unbinding event. It should be noted that the number of such approximate bindings or unbindings will be an upper limit to the number of true binding or unbinding occurrences since observing a transition from a bound to a fully solvated pose will be orders of magnitude more difficult than a change of 4 Å in distance. Nonetheless, this measure gives us a useful approximation to the degree of metastability suffered by the simulation. The results are shown in

Table 1. Two facts stand out: i) even with our lax definition of unbinding, no such events are observed and ii) for the ligand with the most substantial deviation from our assumptions, 7UP, only three binding events are observed, which is in good agreement with the simulation analysis performed in the previous section. If we repeat the same analysis for the plasmin simulations, we can see (Table S5) that all ligands have more than 20 binding events, with many counting over 40, while all but two have at least one unbinding event.

Table 1. Number of binding and unbinding events found for the four ligands of the URO system

Ligand	1UP	2UP	7UP	39L
Number of binding events	12	26	3	12
Number of unbinding events	0	0	0	0

Exploring URO with longer simulations

As an attempt to improve the estimation of binding free energies for the URO system, we designed a new set of simulations. In order to better identify those states that are more stable, we have performed 40 simulations of 1 μ s, starting from the 40 structures selected from the previous AdaptivePELE exit run. This new set of simulations amounts to a total aggregate simulation time of 160 μ s (40 μ s per ligand). Besides, in these simulations, we have modified the simulation box, previously of spherical shape into a cylinder (see Figure S2 for a depiction of the new simulation box). This setup is similar in nature to the funnel metadynamics approach,^{48,49} and the setup for fully flexible docking method introduced by Bertazzo et al.⁵⁰ allowing complete exploration of the binding/unbinding process but restricting the (impossible) exploration of all solvated states and reducing the access to additional metastable surface states. The objective of this additional run is two-fold: i) check whether the longer time scales are sufficient to observe more transitions, and ii) avoid excessive sampling of secondary minima off-site, such as the ones mentioned in the previous section, which should improve the convergence of the estimation, particularly for the larger ligands. The selection of the cylinder parameters (radius and axes) was done manually, although in future versions of this approach we hope to develop an automatic procedure.

In Table 2, the results from this new set of simulations (labeled as “long, cylinder”) versus those shown in the previous section (labeled as “short, sphere”) as well as those of a new set (“long, sphere”) that will be introduced later in this section, are shown. Additionally, in Figure 6B the correlation of the predicted binding free energies with the experimental values is shown. The obtained R^2 value is 0.761, with a Spearman correlation value of 0.6 and an RMSE = 5.64 kcal/mol. Two facts are apparent: i) in the new simulations, the ligands are well separated into weak and strong binders, in opposition to the shorter runs. However, the error in the estimation goes up, from 4.31 kcal/mol to 5.64 kcal/mol, confirming our previous observations of a reduced energy values when limiting the bulk solvent exploration.²³ Furthermore, ii) there is an increase

in the number of observed transitions, we observe two unbinding instances for 1UP, where previously none were observed, and we see seven additional binding events for 39L, and six additional binding events for 7UP, but still no unbinding ones for any of them. The exception seems to be ligand 2UP, where seven fewer bindings are observed, for a total of 19 in the longer simulations. The decrease in transitions in ligands 1UP and 2UP could be easily explained by the fact that the new initial ensemble is smaller, 40 simulations instead of 120. For ligands requiring significantly more propagation time for binding, as in 7UP and 39L, the reduced ensemble is compensated by the increased simulation time producing more overall binding events. In Figure 8, we plot the PMF projections for these new simulations, to be compared with Figure 7. It can be seen that the models estimated from the long simulations show a cleaner profile, with a distinguishable minimum at distances close to zero and begins flattening at distances larger than 15 Å.

Table 2. Comparison of the results obtained for the simulations using a spherical simulation box and 100 ns simulations, the ones using a cylindrical box and 1 μ s, and the ones using a spherical box and 300 ns.

Ligand	Type of simulation	Predicted ΔG (kcal/mol)	Number of bindings	Number of unbindings
1UP	Short, sphere	-4.42	12	0
1UP	Long, cylinder	-1.85	11	2
1UP	Long, sphere	-4.11	3	0
2UP	Short, sphere	-4.17	26	0
2UP	Long, cylinder	-1.94	19	0
2UP	Long, sphere	-3.85	12	0
7UP	Short, sphere	-3.51	3	0
7UP	Long, cylinder	-3.25	9	0
7UP	Long, sphere	-4.682	0	0
39L	Short, sphere	-3.51	12	0
39L	Long, cylinder	-2.96	19	0
39L	Long, sphere	-3.87	3	0

Figure 8. One dimensional projection of the PMF over the distance to the initial structure of the longer simulations and cylindrical “box”. In (A) ligands 1UP and 2UP (weaker binders) are shown, while in (B) ligands 39L and 7UP (stronger binders) are shown

In an attempt to fully distinguish the effect caused by the cylinder box and the longer simulations, we performed a third set of simulations with the URO system, with 40 trajectories of 300 ns, starting from the 40 structures selected from the previous AdaptivePELE exit run. This new set of simulations amounts to a total aggregate simulation time of 48 μ s (12 μ s per ligand). The box used in this last attempt was the exact same sphere that was chosen in the first set of simulations with the URO receptor.

The analysis of this new set followed the same pattern of the previous two, with the same MSM parameters. In Figure 6C, the correlation of the predicted binding free energies with the

experimental values is shown. The obtained R^2 value is 0.326, with a Spearman correlation value of 0.8 and an RMSE = 4.08 kcal/mol. The correlation plot suggests a roughly correct ordering, but with very little difference between the predicted values for the four ligands and overlapping values if the error bars are considered.

Discussion

In this work, we present a multi-step approach leveraging enhanced MC sampling methods and high-performance GPU-accelerated MD simulations to improve the efficiency of absolute binding free energy estimation through MSM. To do so, we developed an extension to the AdaptivePELE method to create an easy-to-use platform to run MD simulations, which automatizes some time-consuming initial steps.

Enhanced sampling methods are becoming more and more popular due to the limited efficiency of unbiased molecular simulations. In this work, we make use of one such method, AdaptivePELE, to obtain a quick approximation to the unbinding mechanism, with the assumption that by starting multiple simulations from several of the states discovered in this initial run, the efficiency of the unbiased (production) simulations in sampling the binding space will increase. In a previous study, we already detected such an improvement when performing the production runs using standard PELE.²³ Here we seek a similar outcome when using molecular dynamics instead of PELE. To this aim we studied two systems with different complexities. The plasmin receptor, a rather small protein-ligand complex, allowed us to analyze the effect of the enhanced initial simulation. With the second system, the URO receptor, we estimated the computational cost when applying the method into a more realistic drug target and ligands.

By comparing the two panels of Figure 4, it is clear that our approach improves the performance of the binding free energy estimations. Albeit far from perfect accuracy, the estimations show a clear positive trend when comparing the prediction with the experimental values, which is not true when the initial exploration is avoided, providing results which are indistinguishable from random. While the Spearman and R^2 plot, Figure 5, indicates an increasing tendency with non-converged profiles, it also shows a good compromise between computational effort and screening potential, pointing to the potential of the method in medium throughput analysis for systems of this complexity. We should remind that the plasmin system, while being a fairly simple protein-ligand complex, has shown to be significantly challenging when using Glide docking or FEP techniques.⁴¹

Figure 9. Analysis of the evolution of the prediction statistics (Spearman and R^2) as the simulation length increases for the long, cylinder “box” ensemble. The solid vertical line marks where the two sets of simulations of the URO system have the same aggregated simulation time, while the dotted line marks the predictions with the same simulation length as the short simulations (100 ns). The starting point is based on adding the first 75 ns segments of all trajectories (with subsequent increments of 25ns).

On the other hand, the strategy of combining enhanced sampling with unbiased simulations does not appear to introduce sufficient efficiency to predict affinities for the URO receptor in a fast manner. This protein, with ligands that are much stronger binders than those of the plasmin receptor, still highlights the sampling problems that are so well-known, especially related to metastable trapping. Only when performing significantly larger simulation times we can correlate ligand's affinities. In Figure 9 we show the evolution of the Spearman and R^2 indicators when increasing the length of the simulations. As in the plasmin system, aggregated time is built by adding the initial fraction in each trajectory. Our results indicate that longer simulations, on the order of 500 ns per trajectory and an aggregated time $\sim 20 \mu s$, are necessary to achieve sufficient converged results. While these values will most likely be system dependent, the PMF projection seems to be a good indication for trusting the results, as shown in our previous study as well.²³

Similar to the funnel metadynamics approach,^{48,49} our results indicate that a smart restriction of the search space largely facilitates convergence. Figure S4A shows the evolution of the Spearman and R^2 indicators when using the original spherical restriction for the same 40 initial trajectories (and up to 300 ns). The results with this larger solvent “box” are significantly worse than with the customized cylinder one at early simulation times. Importantly, while later on correlation values seem to improve (see also Figure 6C), the PMF for the 39L and 7UP still show wrong predictions, with absolute minima located far from the crystal bound pose (Figure S4B). Thus, the analysis should include not only proper R^2 and Spearman values (misleading in this case), but also an inspection of PMF projections.

Comparing the number of binding and unbinding events between the three sets of simulations run with the URO receptor suggests that having a larger ensemble of parallel simulations can be beneficial for exploratory purposes, as evidenced by the larger amount of binding events of the “short, sphere” simulations vs the “long, sphere” one, that use the same amount of aggregated sampling but distributed between 120 and 40 trajectories of 100 and 300 ns, respectively.

In a recent paper,⁵¹ while comparing the efficiency and accuracy of estimations from metadynamics versus unbiased MD simulation, it was reported that 7 of the unbiased MD replicas run up to 2 μ s without showing an unbinding event. One realizes that even with the impressive technical developments that allow performing such long simulations, the task of estimating protein-ligand binding affinities remains a wide-open problem. Even when it is possible to obtain good results, the amount of time that needs to be invested makes such strategies unviable for a high (or even medium/small) throughput use in drug design projects. Thus, using molecular dynamics as the propagator for the production run, at this point, seems to be less efficient than using PELE, see Table S6 for a comparison of both methods' performance. It is possible, however, that more sophisticated uses of enhanced sampling methods, as well as complementary methodologies such as alchemical binding free energy methods, takes us closer to this challenging goal. We believe that the small benchmarks we have used in this work can be a suitable and demanding trial for such methods.

Overall, we have developed a highly automatic and user-friendly molecular dynamics package and applied it to a problem of capital importance, such as computational protein-ligand binding free energy estimation. In doing so, we also provided the most extensive absolute protein-ligand binding free energy calculations using unbiased MD and MSM. The developed methodology showed considerable success when dealing with fragment-like ligands; however, it suffered the well-known limitations of metastable trapping and insufficient sampling of rare events for ligands with higher affinity. As future work, we are currently working on swapping the unbiased simulations by enhanced sampling, with the hope of overcoming the limitations mentioned above and work towards fast and reliable binding free energy estimations.

Acknowledgments

This work has been funded by Spanish projects CTQ2016-79138-R and RTC-2017-6295-1.

Supporting information

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

References

- (1) Jorgensen, W. L. The Many Roles of Computation in Drug Discovery. *Science* **2004**, 303, 1813–1818.
- (2) Zwanzig, R. W. High-Temperature Equation of State by a Perturbation Method. II. Polar Gases. *J. Chem. Phys.* **1955**, 23, 1915–1922.
- (3) Jorgensen, W. L.; Ravimohan, C. Monte Carlo Simulation of Differences in Free Energies of Hydration. *J. Chem. Phys.* **1985**, 83, 3050–3054.
- (4) Wang, L.; Wu, Y.; Deng, Y.; Kim, B.; Pierce, L.; Krilov, G.; Lupyan, D.; Robinson, S.; Dahlgren, M. K.; Greenwood, J.; Romero, D. L.; Masse, C.; Knight, J. L.; Steinbrecher, T.; Beuming, T.; Damm, W.; Harder, E.; Sherman, W.; Brewer, M.; Wester, R.; Murcko, M.; Frye, L.; Farid, R.; Lin, T.; Mobley, D. L.; Jorgensen, W. L.; Berne, B. J.; Friesner, R. A.; Abel, R. Accurate and Reliable Prediction of Relative Ligand Binding Potency in Prospective Drug Discovery by Way of a Modern Free-Energy Calculation Protocol and Force Field. *J. Am. Chem. Soc.* **2015**, 137, 2695–2703.
- (5) Chen, W.; Deng, Y.; Russell, E.; Wu, Y.; Abel, R.; Wang, L. Accurate Calculation of Relative Binding Free Energies between Ligands with Different Net Charges. *J. Chem. Theory Comput.* **2018**. <https://doi.org/10.1021/acs.jctc.8b00825>.
- (6) Mondal, D.; Florian, J.; Warshel, A. Exploring the Effectiveness of Binding Free Energy Calculations. *The Journal of Physical Chemistry B*. **2019**. <https://doi.org/10.1021/acs.jpcb.9b07593>.
- (7) Shirts, M. R.; Mobley, D. L.; Brown, S. P. Free-Energy Calculations in Structure-Based Drug Design. *Drug Design*. pp 61–86. <https://doi.org/10.1017/cbo9780511730412.007>.
- (8) Laio, A.; Parrinello, M. Escaping Free-Energy Minima. *Proc. Natl. Acad. Sci. U. S. A.* **2002**, 99, 12562–12566.
- (9) Roux, B. The Calculation of the Potential of Mean Force Using Computer Simulations. *Comput. Phys. Commun.* **1995**, 91, 275–282.
- (10) Nicolini, P.; Frezzato, D.; Gellini, C.; Bizzarri, M.; Chelli, R. Toward Quantitative Estimates of Binding Affinities for Protein-Ligand Systems Involving Large Inhibitor Compounds: A Steered Molecular Dynamics Simulation Route. *J. Comput. Chem.* **2013**, 34, 1561–1576.
- (11) Zuckerman, D. M. *Statistical Physics of Biomolecules: An Introduction*; CRC Press, 2010.
- (12) Husic, B. E.; Pande, V. S. Markov State Models: From an Art to a Science. *J. Am. Chem. Soc.* **2018**, 140, 2386–2396.
- (13) Bowman, G. R.; Pande, V. S.; Noé, F. *An Introduction to Markov State Models and Their Application to Long Timescale Molecular Simulation*; Springer Science & Business Media, 2013.
- (14) Stanley, N.; Esteban-Martín, S.; De Fabritiis, G. Kinetic Modulation of a Disordered Protein Domain by Phosphorylation. *Nat. Commun.* **2014**, 5, 5272.
- (15) Casasnovas, R.; Limongelli, V.; Tiwary, P.; Carloni, P.; Parrinello, M. Unbinding Kinetics of a p38 MAP Kinase Type II Inhibitor from Metadynamics Simulations. *J. Am. Chem. Soc.* **2017**, 139, 4780–4788.
- (16) Doudou, S.; Burton, N. A.; Henchman, R. H. Standard Free Energy of Binding from a One-Dimensional Potential of Mean Force. *J. Chem. Theory Comput.* **2009**, 5, 909–918.

- (17) Buch, I.; Giorgino, T.; De Fabritiis, G. Complete Reconstruction of an Enzyme-Inhibitor Binding Process by Molecular Dynamics Simulations. *Proc. Natl. Acad. Sci. U. S. A.* **2011**, *108*, 10184–10189.
- (18) Bernetti, M.; Masetti, M.; Recanatini, M.; Amaro, R. E.; Cavalli, A. An Integrated Markov State Model and Path Metadynamics Approach To Characterize Drug Binding Processes. *J. Chem. Theory Comput.* **2019**, *15*, 5689–5702.
- (19) Provasi, D.; Bortolato, A.; Filizola, M. Exploring Molecular Mechanisms of Ligand Recognition by Opioid Receptors with Metadynamics. *Biochemistry* **2009**, *48*, 10020–10029.
- (20) Pietrucci, F.; Marinelli, F.; Carloni, P.; Laio, A. Substrate Binding Mechanism of HIV-1 Protease from Explicit-Solvent Atomistic Simulations. *J. Am. Chem. Soc.* **2009**, *131*, 11811–11818.
- (21) Troussicot, L.; Guillièrre, F.; Limongelli, V.; Walker, O.; Lancelin, J.-M. Funnel-Metadynamics and Solution NMR to Estimate Protein–Ligand Affinities. *Journal of the American Chemical Society*. **2015**, pp 1273–1281. <https://doi.org/10.1021/ja511336z>.
- (22) 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. *J. Chem. Theory Comput.* **2017**, *13*, 3372–3377.
- (23) Gilabert, J. F.; Grebner, C.; Soler, D.; Lecina, D.; Municoy, M.; Gracia, O. G.; Soliva, R.; Packer, M. J.; Hughes, S. J.; Tyrchan, C.; Hogner, A.; Guallar, V. PELE-MSM: A Monte Carlo Based Protocol for the Estimation of Absolute Binding Free Energies. *J. Chem. Theory Comput.* **2019**. <https://doi.org/10.1021/acs.jctc.9b00753>.
- (24) Borrelli, K. W.; Vitalis, A.; Alcantara, R.; Guallar, V. PELE: Protein Energy Landscape Exploration. A Novel Monte Carlo Based Technique. *J. Chem. Theory Comput.* **2005**, *1*, 1304–1311.
- (25) Lecina, D.; Gilabert, J. F.; Guallar, V. Adaptive Simulations, towards Interactive Protein-Ligand Modeling. *Sci. Rep.* **2017**, *7*, 8466.
- (26) Zhu, K.; Shirts, M. R.; Friesner, R. A. Improved Methods for Side Chain and Loop Predictions via the Protein Local Optimization Program: Variable Dielectric Model for Implicitly Improving the Treatment of Polarization Effects. *J. Chem. Theory Comput.* **2007**, *3*, 2108–2119.
- (27) Onufriev, A.; Bashford, D.; Case, D. A. Exploring Protein Native States and Large-Scale Conformational Changes with a Modified Generalized Born Model. *Proteins* **2004**, *55*, 383–394.
- (28) Gilabert, J. F.; Lecina, D.; Estrada, J.; Guallar, V. Monte Carlo Techniques for Drug Design: The Success Case of PELE. *Biomolecular Simulations in Structure-Based Drug Discovery* **2018**, 87–103.
- (29) Strauss, D. J.; Hartigan, J. A. Clustering Algorithms. *Biometrics*. 1975, p 793. <https://doi.org/10.2307/2529577>.
- (30) Eastman, P.; Swails, J.; Chodera, J. D.; McGibbon, R. T.; Zhao, Y.; Beauchamp, K. A.; Wang, L.-P.; Simmonett, A. C.; Harrigan, M. P.; Stern, C. D.; Wiewiora, R. P.; Brooks, B. R.; Pande, V. S. OpenMM 7: Rapid Development of High Performance Algorithms for Molecular Dynamics. *PLOS Computational Biology*. **2017**, p e1005659. <https://doi.org/10.1371/journal.pcbi.1005659>.
- (31) Maier, J. A.; Martinez, C.; Kasavajhala, K.; Wickstrom, L.; Hauser, K. E.; Simmerling, C. ff14SB: Improving the Accuracy of Protein Side Chain and Backbone Parameters from ff99SB. *J. Chem. Theory Comput.* **2015**, pp 3696–3713. <https://doi.org/10.1021/acs.jctc.5b00255>.
- (32) Wang, J.; Wolf, R. M.; Caldwell, J. W.; Kollman, P. A.; Case, D. A. Development and Testing of a General Amber Force Field. *Journal of Computational Chemistry*. **2004**, pp 1157–1174. <https://doi.org/10.1002/jcc.20035>.

- (33) Wang, J.; Wang, W.; Kollman, P. A.; Case, D. A. Automatic Atom Type and Bond Type Perception in Molecular Mechanical Calculations. *J. Mol. Graph. Model.* **2006**, *25*, 247–260.
- (34) Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of Simple Potential Functions for Simulating Liquid Water. *The Journal of Chemical Physics.* **1983**, pp 926–935. <https://doi.org/10.1063/1.445869>.
- (35) Darden, T.; York, D.; Pedersen, L. Particle Mesh Ewald: AnN·log(N) Method for Ewald Sums in Large Systems. *The Journal of Chemical Physics.* **1993**, pp 10089–10092. <https://doi.org/10.1063/1.464397>.
- (36) AdaptivePELE. AdaptivePELE/AdaptivePELE <https://github.com/AdaptivePELE/AdaptivePELE> (accessed Sep 30, 2019).
- (37) Prinz, J.-H.; Wu, H.; Sarich, M.; Keller, B.; Senne, M.; Held, M.; Chodera, J. D.; Schütte, C.; Noé, F. Markov Models of Molecular Kinetics: Generation and Validation. *J. Chem. Phys.* **2011**, *134*, 174105.
- (38) Scherer, M. K.; Trendelkamp-Schroer, B.; Paul, F.; Pérez-Hernández, G.; Hoffmann, M.; Plattner, N.; Wehmeyer, C.; Prinz, J.-H.; Noé, F. PyEMMA 2: A Software Package for Estimation, Validation, and Analysis of Markov Models. *J. Chem. Theory Comput.* **2015**, *11*, 5525–5542.
- (39) Wehmeyer, C.; Scherer, M. K.; Hempel, T.; Husic, B. E.; Olsson, S.; Noé, F. Introduction to Markov State Modeling with the PyEMMA Software [Article v1.0]. *Living Journal of Computational Molecular Science.* **2019**. <https://doi.org/10.33011/livecoms.1.1.5965>.
- (40) Cheng, L.; Pettersen, D.; Ohlsson, B.; Schell, P.; Karle, M.; Evertsson, E.; Pahlén, S.; Jonforsen, M.; Plowright, A. T.; Boström, J.; Fex, T.; Thelin, A.; Hilgendorf, C.; Xue, Y.; Wahlund, G.; Lindberg, W.; Larsson, L.-O.; Gustafsson, D. Discovery of the Fibrinolysis Inhibitor AZD6564, Acting via Interference of a Protein-Protein Interaction. *ACS Med. Chem. Lett.* **2014**, *5*, 538–543.
- (41) Schmidt, T. C.; Eriksson, P.-O.; Gustafsson, D.; Cosgrove, D.; Frølund, B.; Boström, J. Discovery and Evaluation of Anti-Fibrinolytic Plasmin Inhibitors Derived from 5-(4-Piperidyl)isoxazol-3-ol (4-PIOL). *J. Chem. Inf. Model.* **2017**, *57*, 1703–1714.
- (42) Jankun, J.; Skrzypczak-Jankun, E. Molecular Basis of Specific Inhibition of Urokinase Plasminogen Activator by Amiloride. *Cancer Biochem. Biophys.* **1999**, *17*, 109–123.
- (43) Matthews, H.; Ranson, M.; Kelso, M. J. Anti-Tumour/metastasis Effects of the Potassium-Sparing Diuretic Amiloride: An Orally Active Anti-Cancer Drug Waiting for Its Call-of-Duty? *Int. J. Cancer* **2011**, *129*, 2051–2061.
- (44) Tang, L.; Han, X. The Urokinase Plasminogen Activator System in Breast Cancer Invasion and Metastasis. *Biomed. Pharmacother.* **2013**, *67*, 179–182.
- (45) Dunbar, J. B.; Smith, R. D.; Damm-Ganamet, K. L.; Ahmed, A.; Esposito, E. X.; Delproposto, J.; Chinnaswamy, K.; Kang, Y.-N.; Kubish, G.; Gestwicki, J. E.; Stuckey, J. A.; Carlson, H. A. CSAR Data Set Release 2012: Ligands, Affinities, Complexes, and Docking Decoys. *J. Chem. Inf. Model.* **2013**, pp 1842–1852. <https://doi.org/10.1021/ci4000486>.
- (46) Witek, J.; Mühlbauer, M.; Keller, B. G.; Blatter, M.; Meissner, A.; Wagner, T.; Riniker, S. Interconversion Rates between Conformational States as Rationale for the Membrane Permeability of Cyclosporines. *ChemPhysChem.* **2017**, pp 3309–3314. <https://doi.org/10.1002/cphc.201700995>.
- (47) Takahashi, R.; Gil, V. A.; Guallar, V. Monte Carlo Free Ligand Diffusion with Markov State Model Analysis and Absolute Binding Free Energy Calculations. *J. Chem. Theory Comput.* **2014**, *10*, 282–288.
- (48) Limongelli, V.; Bonomi, M.; Parrinello, M. Funnel Metadynamics as Accurate Binding Free-Energy Method. *Proc. Natl. Acad. Sci. U. S. A.* **2013**, *110*, 6358–6363.
- (49) Brotzakis, Z. F.; Faidon Brotzakis, Z.; Limongelli, V.; Parrinello, M. Accelerating the Calculation of Protein-Ligand Binding Free Energy and Residence Times Using Dynamically Optimized Collective Variables. *J. Chem. Theory Comput.* **2019**, *15*, 743–750.

- (50) Bertazzo, M.; Bernetti, M.; Recanatini, M.; Masetti, M.; Cavalli, A. Fully Flexible Docking via Reaction-Coordinate-Independent Molecular Dynamics Simulations. *J. Chem. Inf. Model.* **2018**, 58, 490–500.
- (51) 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.

For Table of Contents Only