

Estimation of binding free energies with Monte Carlo atomistic simulations and enhanced sampling

Ph.D thesis in Computational and Applied Physics
Department of Applied Physics, June 2020

Candidate: Joan Francesc Gilabert Navarro

Director: Prof. Victor Guallar Tasies (Barcelona Supercomputing Center)
Tutor: Prof. Daniel López Codina (Universitat Politècnica de Catalunya)



Contents

Abstract	3
List of Abbreviations	5
Chapter 1 Introduction	7
1.1 Motivation	7
1.2 Objectives	9
1.3 Theoretical basis	10
1.3.1 Biomolecules: proteins and ligands	10
1.3.2 Protein-ligand binding	11
1.4 State of the art	13
1.4.1 Virtual Screening	14
1.4.2 End-point methods	16
1.4.3 Alchemical methods	16
1.4.4 Pathway-based methods	18
1.5 Molecular mechanics	20
1.5.1 Molecular dynamics	22
1.5.2 Monte Carlo methods	23
1.5.3 PELE	24
1.5.4 Enhanced sampling	26
1.6 Markov State Models	29
1.7 PMF, probability distribution and ΔG	31
1.8 Systems used	33
1.8.1 Trypsin	33
1.8.2 Plasminogen kringle domain 1	34
1.8.3 Urokinase-type plasminogen activator	34
1.8.4 Progesterone receptor	35
1.8.5 Mitogen-activated protein kinase 1	36
Chapter 2 Results	40
2.1 AdaptivePELE	40
2.1.1 AdaptivePELE algorithm	40
2.1.2 Binding free energies from AdaptivePELE simulations	43
2.2 Protocol for the estimation of binding free energies	52
2.3 PELE-MSM	57
2.3.1 Binding Free Energy for the Plasmin System	57
2.3.2 Binding Free Energy for the URO System	59

2.3.3	Binding Free Energy for the ERK2 System	63
2.3.4	Binding Free Energy for the PR System	65
2.3.5	Effect of the previous AdaptivePELE simulation	69
2.4	MD-MSM	72
2.4.1	Binding Free Energy for the Plasmin System	72
2.4.2	Binding Free Energy for the URO System	75
2.4.3	Binding Free Energy for the ERK2 System	83
2.4.4	Effect of the previous AdaptivePELE simulation	88
2.5	AdaptivePELE-MSM	89
2.5.1	Binding Free Energy for the URO System	90
2.5.2	Binding Free Energy for the PR System	93
Chapter 3 Discussion		96
3.1	Importance of scientific software	96
3.2	Successes and limitations of the protocol	97
3.2.1	Sampling limitations	99
3.2.2	Modelling limitations	101
3.3	Effectiveness of unbiased simulations	104
3.3.1	Role of enhanced sampling methods	105
3.3.2	Sampling efficiency of longer vs multiple shorter simulations	108
Chapter 4 Conclusions		111
Appendices		112
Appendix A PMF plots for PELE-MSM benchmark		113
A.1	PMF plots for plasmin	113
A.2	PMF plots for URO system	116
A.3	PMF plots for ERK2 system	119
Appendix B PMF plots MD-MSM		122
B.1	PMF plots for plasmin	122
B.2	PMF plots plasmin no exploration	127
Appendix C Structures ERK2 MD simulations		131
C.1	Structures ERK2 sphere	131
C.2	Structures ERK2 cylinder	133
Appendix D PMF plots for AdaptivePELE-MSM		135
D.1	PMF plots for URO system	135
D.2	PMF plots for PR system	141
Appendix E Thesis summary		149
E.1	Thesis summary	149
E.2	Resum de la thesis	150
Bibliography		152

Abstract

The advances in computing power have motivated the hope that computational methods can accelerate the pace of drug discovery pipelines. For this, fast, reliable and user-friendly tools are required. One of the fields that has gotten more attention is the prediction of binding affinities. In this context, this thesis is focused in the development of efficient methods for the estimation of protein-ligand binding free energies. We have developed a protocol that combines enhanced sampling with more standard simulations methods to achieve higher efficiency. Thanks to the modularity and automation that the protocol offers we were able to test three different simulation methods. Through the use of a diverse benchmark of protein systems we have established the cases in which the protocol is expected to give accurate results, and which areas require further development. Overall, we hope that the methodology developed in this work can contribute to the drug discovery process and motivate further advances in the field.

Keywords: Binding free energy, enhanced sampling, PELE, molecular dynamics, Markov State Models, AdaptivePELE

Acknowledgements

First of all I would like to thank my thesis director, Professor Victor Guallar Tasies, for giving me the opportunity to carry out my Ph.D. During these years I have learnt about enzymes, drug discovery, molecular simulations and plenty of other topics that would have been impossible if I had not been granted the freedom that I had. Furthermore, he taught me invaluable lesson as a scientist, how to thoroughly analyse data and the importance of being able to clearly communicate findings to others. I want to thank my UPC tutor Professor Daniel López Codina for facilitating my enrolment in the Ph.D programme and always being available to me.

A big thank you goes to all the people that I have worked with in the EAPM group. Special thanks to Dani Lecina and Jorge Estrada, whose help and mentoring were incomensurable during my early days as a master student. Also, I am grateful to Oriol Gracia, for many clarifying discussions and collaboration, a part of this thesis would probably not been possible without his work. Moreover, I am grateful to the rest of past and present colleagues: Gerard, Pep, Carles, Ferran, Emanuele, Suwipa, Marina, Martin, Alberto, Sergi, Isaac, Oliver, Jelisa, Martí, Ruben and Marc, coffe breaks and lunches will not be the same without you. I also want to thank Robert Soliva and Daniel Soler from Nostrum Biodiscovery, for their fruitful collaboration in the development of PELE-MSM.

All in life is not work, and I would like to thank all the people that has participated in the BSC Basketball games during all these years. I met great people from all departments there, and during my Ph.D I have always looked forward to Tuesday afternoons.

I want to give special thanks to my family and friends, for their support not only during this period but for all my life. To my parents, grandparents, my little brother and rest of family. They have always served as an example of hard work, which I tried to mirror. I would also like to thank them for their faith in me, which I will never be able to repay.

My last words of gratitude go to my girlfriend, Berta. Her encouragement and support has brought out the best version of me. I love you.

List of Abbreviations

ΔG_b^o	Standard free energy of binding
C^o	Standard concentration
k_B	Boltzmann constant
ANM	Anisotropic Network Model
CV	Collective Variable
ERK2	Mitogen-activated protein kinase 1
FAST	Fluctuation Amplification of Specific Traits
FEL	Free Energy Landscape
FEP	Free Energy Perturbation
GMRQ	Generalized Matrix Rayleigh Quotient
GPU	Graphical Processing Unit
iMapD	Intrinsic Map Dynamics
LIE	Linear Interaction Energy
MAB	Multi Armed Bandit
MC	Monte Carlo
MCMC	Markov Chain Monte Carlo
MD	Molecular Dynamics
MM	Molecular mechanics
PCA	Principal Component Analysis
PDB	Protein Data Bank
PELE	Protein Energy Landscape Exploration
Plasmin	Plasminogen kringle domain 1

PMF Potential of mean force
PR Progesterone receptor
RAVE Reweighted Autoencoded Variational Bayes for Enhanced sampling
REAP REinforcement learning based Adaptive samPling
RL Reinforcement Learning
RMSD Root-Mean-Square Deviation
RMSE Root-Mean-Square Error
SASA Solvent-Accessible Surface Area
TI Thermodynamic Integration
TICA Time-lagged Independent Component Analysis
TRP Trypsin
URO Urokinase-type plasminogen activator
VAC Variational Approach to Conformational dynamics
VAE Variational AutoEncoder
VS Virtual Screening
WE Weigthed Ensemble

Chapter 1

Introduction

1.1 Motivation

The process of developing a drug is long and costly. Even if we ignore the difficulty of characterising a disease and finding an appropriate target for treatment, an average of around 10-15 years are needed to bring a new drug to the market (Ciociola et al., 2014). The overall cost of this long process has been a subject of great debate, but studies estimate that it has an associated cost of somewhere between US\$161 million and US\$2 billion (Sertkaya et al., 2016).

The main reason why drug candidates take so long to be developed is that a successful drug requires many parameters to be tuned, as schematically illustrated in Figure 1.1. First, the molecule needs to bind the target of interest, and do so for enough time that the desired effect can be produced. Furthermore, the drug must be selective, avoiding interaction with other proteins that might cause side-effects and/or toxicity. These are only a small subset of properties to be tuned, many others such as solubility, permeability or bioavailability must be considered to ensure that the molecule reaches the target protein in the patient (Gleeson, 2008). For this reason, the process of drug discovery is typically divided in several iterations, where the number of considered molecules is reduced as the candidates fail to meet standard values for the properties described before.

Computational methods have found their way in multiple steps of the drug discovery process, due to their potential efficiency increase and cost reduction (Jorgensen, 2004). Different methods have impacted different stages of the pipeline. In the initial steps, the most commonly used techniques are those which have low computational cost, such as molecular docking and scoring. They are usually grouped under the virtual screening category (Patrick Walters et al., 1998), and simulate in a computer the screening of vast libraries of molecules to find good matches for the target in question. This procedure is particularly useful where there is not much previous knowledge about the receptor and serves as a quick filter for ligands that are not likely to have enough affinity for the target protein.

When a more detailed understanding of the chemical interplay between protein and ligand




Figure 1.1: Schematic illustration of some of the most important parameters to be optimised during the drug discovery process. The number represents the usual order in which each parameter is typically optimised.

is required, it is common to resort to more complex molecular mechanics (MM) methods. Monte Carlo (MC) and molecular dynamics (MD) simulations are the most popular version of such techniques. They use an atomic-resolution classical force field to estimate the potential energy of the protein-ligand complex, producing a more accurate approximation than virtual screening methods, but requiring a much higher computational time. These methodologies are able to sample extensively both the ligand and protein's degrees of freedom. In the limit of infinite sampling, this could be used to easily compute the binding affinity, as will be discussed in section 1.3, but, in the usual situation of limited sampling, simulations provide only an incomplete picture of the binding process, and thus, a less accurate affinity estimation. This limitation has led to the development of several techniques, known as enhanced sampling methods, that attempt to circumvent this problem by modifying the underlying potential energy surface, such as Umbrella Sampling (Torrie and Valleau, 1977) or Metadynamics (Laio and Parrinello, 2002).

The main difficulty in the estimation of binding free energies lies in the tremendous complexity of the protein-ligand dynamics. At the atomic level, the protein-ligand complex is typically formed by several thousand atoms, which can interact through different forces. Some of these interactions operate at small distances, while others have long ranges. The precise determination of the interplay between the multiple molecular interactions and the complete exploration of the highly dimensional space make of producing accurate and fast estimates of the free energy of binding a formidable task.

The most accurate techniques for the estimation of binding free energies are based on MM methods explained above. They are generally split into alchemical methods and pathway-based methods. Alchemical methods estimate the binding free energy through a nonphysical transformation, reducing the amount of simulation required, while pathway-based methods need to reproduce the entire binding or unbinding path. Both groups produce generally accurate predictions but are either limited by their computational cost or by restricted applicability for some systems (see sections 1.4.3 and 1.4.4 for more details of these techniques respectively).

Enhanced sampling methods can certainly alleviate some of these problems, reducing the simulation time required and allowing the study of problems on time scales that are not accessible to regular MM simulations. Nonetheless, commonly enhanced sampling often requires non-negligible theoretical expertise in order to produce reliable results and/or complicated initial setups. In order to have a significant impact in the drug discovery process, computational approaches need to be fast and accurate, but also user-friendly to some extent. Current developments are closing the gap on the first two challenges, aided by the increase of computational power available. However, no universally applicable method exists that can quickly and accurately predict the binding free energy of any protein-ligand system.

1.2 Objectives

The aim of this dissertation is to develop an efficient protocol for the estimation of protein-ligand binding free energies. The main goal is to produce a method that can benefit from the accuracy of MM techniques but at the same time reduce the simulation time required. The ultimate objective of such protocol would be to produce fast and reliable binding affinity estimates, so that it could be applied in the early steps of the drug discovery process, focusing the search for new drug molecules and improving its efficiency. In short, this thesis is dedicated to testing the following hypothesis:

Using enhanced sampling methods, the sampling of the protein-ligand binding and unbinding processes can be accelerated to a degree where fast absolute binding free energy estimation can be accomplished, allowing its routine use in real-life drug discovery projects.

The objectives are divided in three parts. First, to develop the full protocol, designing, testing and implementing its different steps. Combining several methods requires careful design to ensure seamless interconnection, which in turn has the advantage of being scalable, permitting its application to larger benchmarks and providing more robust results. Secondly, we compare the two main MM methods (MC and MD) as the simulation propagators. Both techniques have demonstrated their effectiveness in the study of protein-ligand interactions, but given their unique characteristics we attempt to study which one is better suited for the estimation of binding free energies. Finally, we explore whether the use of an enhanced sampling as the simulation propagator can significantly accelerate the full protocol run time while providing accurate results.

This work attempts to solve one of the most difficult open problems in computational biophysics, the estimation of protein-ligand binding free energies. This is not only an important academic challenge, but any advances achieved are likely to contribute to the search of new pharmacological treatments.

1.3 Theoretical basis

1.3.1 Biomolecules: proteins and ligands

Living organisms use proteins as their workforce. For example, a protein called trypsin actively participates in the digestion process, DNA polymerase reads and replicates the DNA sequences stored in the nucleus of our cells and another protein, collagen, forms fibrous tissues such as tendons, ligaments, and skin. The diversity of functions performed by proteins, and therefore their structures, is immense. In the human body, a minimum of 20.000 proteins exists, although some argue that this number may be as high as 100.000 (Ponomarenko et al., 2016). Remarkably, all this variety arises from the combination of only 20 amino acids. The different chemical and structural characteristics of the amino acids (depicted in Figure 1.2A) give in turn different properties to the resulting protein.

The sequence of amino acids (also known as residues) that forms a protein is determined from the DNA sequence that encodes it. In a ribosome, the residues are attached, forming a linear filament (which is known as the primary structure). This shape is unstable, and the secondary structure quickly emerges. The protein then forms a series of helices (α -*helix*) and planar sheets (β -*sheet*) connected by stretched regions known as *loops*. At this stage, the protein structure is partially folded, the newly formed domains participate in interactions that cause the molecule to fold completely and adopt its tertiary structure. Usually, this is the protein's active form, but sometimes several copies of the protein are required to carry out its function, forming what is known as the quaternary structure. This process is illustrated in Figure 1.2B for the case of hemoglobin, a protein that transports oxygen from the lungs to the rest of the body, and which is formed by two pairs of two different subunits, commonly known as α and β .

Proteins are not the only biomolecules present in a cell. A cell is a crowded aqueous environment, where in addition to proteins, many other molecules such as ions, nucleic acids, lipid membranes or small molecules are present. The latter category typically encompasses what we will refer to as ligands in this thesis, compounds with a molecular weight smaller than 900 Dalton (Dougherty and Pucci, 2014). Ligands are commonly found in protein binding sites, contributing to the protein function. Often, pharmaceutical drugs fall into this category, since their reduced size gives them important advantages, such as improve absorption and simpler administration (Robin Ganellin et al., 2013). Contrary to proteins, ligands are not constructed by the combination of a fixed set of monomers, but by covalently bonded atoms (commonly carbon, oxygen, nitrogen, sulphur and the atoms in the halogen group). As a consequence, the size of the space spanned by all theoretically possible ligands has been estimated to be at least of 10^{60} molecules (Dobson, 2004). This large space is one of the main reasons why creating a new drug is so difficult,




Figure 1.2: Protein building blocks and the different hierarchies of protein structure. Panel (aA): Natural amino acids classified according to their chemical properties. Source: https://en.wikipedia.org/wiki/File:Amino_acids.png. Panel (b): Process of protein folding illustrating the different structural structural hierarchies: primary, secondary, tertiary and quaternary. Source: https://en.wikipedia.org/wiki/File:225_Peptide_Bond-01.jpg

to advance one needs to reduce the amount of possible molecules to a tractable size, for example in the order of 10^4 , which leaves out almost all of the chemical space, probably missing out on possibly useful ligands.

1.3.2 Protein-ligand binding

Diseases may appear for multiple reasons. For example, mutations in a protein might cause misfolding, that is, the protein adopting a structure that prevents the molecule from performing its function. In other cases, the protein is overexpressed, working at a problematically excessive rate. Often, these malfunctions are targeted through protein-ligand interactions, where the ligand might cause the protein to stop working, or the opposite, helping the protein achieve its function, by modifying its structure. Several theories of the mechanism through which such interactions occur exist.

The oldest theory is known as the *Lock-and-Key* model (Fischer, 1894), which describes the protein and the ligand as two rigid molecules (the “lock” and the “key”, respectively). In the model, the protein-ligand complex is formed if, and only if, the shapes of the two molecules match, in analogy of a real life lock and key.

The assumption of rigidity imposes too great a restriction for the study of protein-ligand interaction. Further insights lead to the substitution of the lock-and-key model for other theories that account for the molecules’ dynamical nature. The two main alternatives are the *induced fit* (Koshland, 1958) and the *conformational selection* models (Frauenfelder et al., 1991; Tsai et al., 1999). In the former, the shape of the two compounds does not

need to perfectly match, instead a provisional complex is formed and then a change in the protein structure (and ligand) is induced to accommodate the latter.

On the other hand, conformational selection, sometimes also called *population selection*, *fluctuation fit* or *selected fit*, describes the binding process as the interaction between the fluctuating structural population of both the protein and the ligand, “selecting” from all the available conformations those that match. It should be noted, however, that the two competing theories are not mutually exclusive. Several examples can be found where the effects described by the conformational selection model are followed by those of induced fit (Grünberg et al., 2004; Wlodarski and Zagrovic, 2009), and extended models that combine the two have been developed (Csermely et al., 2010).

The protein-ligand complex is typically, but not always, a non-covalent interaction mediated by hydrogen bonds, hydrophobic forces, van der Waals forces, π - π interactions and/or electrostatic interactions. The combined strength of said interactions is defined as the protein-ligand binding affinity, given by the equilibrium dissociation constant K_d , or its inverse, the binding constant K_b . We start by modelling the binding process using the following chemical reaction



where R represents the receptor, L represents the ligand, and RL represents the receptor-ligand complex. At equilibrium, the reaction is characterised by its binding constant

$$K_b = \frac{C_{RL}^{eq}}{C_L^{eq} C_R^{eq}} \quad (1.2)$$

where C_L^{eq} is the ligand equilibrium concentration, C_R^{eq} is the receptor equilibrium concentration and C_{RL}^{eq} is the complex equilibrium concentration. As mentioned above, the dissociation constant measures the strength of the protein-ligand interaction. Often, however, instead of the dissociation constant, the binding free energy is reported. The two quantities are related by

$$\Delta G_b^o = -k_B T \ln(C^o K_b) \quad (1.3)$$

where ΔG_b^o is the standard binding free energy (which will also be referred to in this work as the *absolute protein-ligand binding free energy*), k_B the Boltzmann constant, T the temperature and C^o the standard concentration. The term “standard”, here, refers to the standard state, which is arbitrarily typically defined as $C^o = 1$ M. This definition is necessary to compare different binding free energies, since, contrary to equation 1.2, the definition of the standard binding free energy is independent of the concentration of the three species of molecules (General, 2010). A more detailed treatment of this relation can be found in a review by Zhou and Gilson (Gilson and Zhou, 2007).

An alternative, more visual, understanding of the binding free energy comes from the *free energy profile* of the reaction, shown in Figure 1.3a. A free energy profile plots the free

energy value along a *reaction coordinate*, a parameter that represents the progress along a reaction and is able to separate the different states. In the example shown here, there are four states. L-P, which represents the protein-ligand complex, is the most stable state and therefore has the minimum free energy value. L+P represents the free protein and ligand, with no interaction between them, which is why the value for the free energy is constant. Finally, the states labelled L-P* and L-P** represent hypothetical metastable intermediate states, that is, states that are not stable but are long-lived and relevant to the dynamics.

The presence of these metastable states is meant to indicate the complexity of the dynamics of protein-ligand binding, which may not be apparent from the reaction and the free energy profile. Importantly, these four states are located in local minima of the free energy and are separated by local maxima, free energy barriers that need to be overcome to transition between two states. Moreover, to obtain a more realistic representation of the system evolution one should study the *free energy landscape* (FEL), which is significantly more complex than the free energy profile, as schematically illustrated in Figure 1.3b.

The free energy profile can be understood as a section of the FEL, projected over a single coordinate, but the full dimensionality of the FEL is on the order of $3N$, where N is the number of atoms of the system; in practice the full dimensionality is less than $3N$, given that many coordinates are correlated due to the intramolecular bonds. As a consequence, the picture painted by the FEL is significantly more intricate, with many more metastable states represented as basins that are shallower than the absolute minimum. The presence of this multitude of valleys contributes to the metastable trapping of molecular simulations, reducing their efficiency. For this reason, the concept of the free energy landscape has been extensively employed in the design of enhanced sampling methods (see section 1.5.4 for more details).

1.4 State of the art of protein-ligand binding free energy calculations

The computational estimation of binding free energies is an active field of research, exploring multiple different approaches, as illustrated in a recent review article (Limongelli, 2020). All the research effort has led to significant advances but has also uncovered several important challenges, which are typically classified in two categories: i) those related to complexity of the protein-ligand interactions, and ii) those related to accuracy limitations of the models used.

Proteins are composed of several thousand atoms, which means that when they are studied by methods with atomic resolution several thousand degrees of freedom need to be taken into account. Thus, an exhaustive description of protein-ligand binding would require to explore a $\sim 3N$ -dimensional space, where N is the number of atoms of the complex, a space so large that it is impossible to explore using current techniques (Mobley, 2012). Nonetheless, even if one could somehow sample the full $3N$ dimensions, in the so called “infinite sampling” scenario, existing methods would still only provide a partially correct answer, as




Figure 1.3: Panel (a): Representation of the free energy profile of the protein-ligand binding reaction, with the binding free energy as the difference between the bound and unbound states. Several possible metastable states, L-P* and L-P**, are shown. Adapted with permission from: Limongelli V. Ligand binding free energy and kinetics calculation in 2020. *WIREs Comput Mol Sci.* 2020; e1455. <https://doi.org/10.1002/wcms.1455>. Panel (b): Representation of the free energy landscape of the protein-ligand system. Source: <https://tinyurl.com/v6lwg22>

most of the current techniques model the molecular interactions using classical physics, but said interactions are essentially quantum in nature. The combined inaccuracies of these two limitations establishes a compromise that the developers of new methodologies must consider, as using more accurate models limits the sampling capabilities that the technique can offer.

The trade-off mentioned above gives rise to a natural classification of the existing methods for binding free energy estimation according to two axes: cost (measured as the computational time required) and accuracy. In Figure 1.4, an schematic representation of the methods that will be discussed in sections 1.4.1 to 1.4.4 is shown. Each category is shown in a different colour, with an area that represents approximately where the methods would fall in terms of accuracy and computational cost. It should be noted that this illustration is only meant to simplify the introduction to the different techniques, and it should not be taken as a quantitative comparison between the categories.

1.4.1 Virtual Screening

The first class of methods we explore are the so-called *Virtual screening* methods (VS). The techniques included in this category are characterised by fast execution times and are designed to process virtual libraries of between tens of thousands and millions of molecules (Patrick Walters et al., 1998). Virtual screening is typically employed at the beginning of the drug discovery process, when even the most approximate methods can contribute to advance. In a VS pipeline, multiple filters are applied to subsequently less molecules. As




Figure 1.4: Schematic representation of state of the art methods according to the cost-accuracy trade-off. Each group of methods is coloured differently and assigned an area according to the approximate values of cost and accuracy that characterises them.

the steps become more detailed, the size of the input is reduced, until the final steps of docking and scoring, which take into account 3D structural information about the protein and ligand, are reached.

Docking programs attempt to find the best possible fit between two molecules, however, to do so in a fast manner, they simplify the problem by assuming that both protein and ligand are rigid, although by using many different ligand structures the assumption of ligand rigidity can be eliminated in practice, see for example the AutoDock Vina (Trott and Olson, 2010), DOCK (Lang et al., 2009) or rDock (Ruiz-Carmona et al., 2014) software. This is obviously limiting, as the dynamical nature of protein is key in the binding process. To alleviate this constraint, improved strategies exist, such as induced-fit docking and ensemble docking. Induced-fit docking, implemented for example in Glide (Halgren et al., 2004), partially introduces flexibility into the protein, producing better results at a higher computational cost. On the other hand, ensemble docking imitates the “conformational selection” mechanism of binding, and instead of attempting to fit the ligand into a single protein conformation, generates an ensemble of receptor structures and then tries to dock the ligand (Huang and Zou, 2007).

The output of a docking run is a list of poses, usually with around a thousand elements, which needs to be somehow ranked to be useful. This is the task of Scoring techniques (Huang et al., 2010). Scoring functions use different strategies to predict the binding affinity, some use a force-field based approach, calculating the energy of the complex from models of the molecular interactions, others use a knowledge-based or even consensus functions, combining different scoring functions for a more robust result.

1.4.2 End-point methods

Removing the requirement of molecular rigidity allows for a greater exploration of protein-ligand interaction at the expense of larger computational efforts. Most methods for the estimation of binding free energies assume this cost, and the computationally cheapest group of techniques is known as *end-point methods*. Looking back at the reaction shown in Equation 1.1, end-point methods simulate only the two extreme points (hence the name), the protein-ligand complex and the two molecules diffusing freely. The main examples of such methods are the Linear Interaction Energy (LIE) (Wang et al., 1999), MM/PBSA (Swanson et al., 2004) and MM/GBSA (Genheden and Ryde, 2015). The core estimation of this group of techniques can be expressed as

$$G = E_{internal} + E_{electrostatic} + E_{vdw} + G_{sol} - TS \quad (1.4)$$

where $E_{internal}$ is the energy associated with the internal degrees of freedom (bonds, angles and dihedrals), $E_{electrostatic}$ is the energy associated with the electrostatic interactions, E_{vdw} represents the van der Waals terms, G_{sol} is the solvation free energy and $-TS$ corresponds to the entropy contribution. The binding free energy is then calculated as

$$\Delta G = \langle G_{RL} - G_R - G_L \rangle \quad (1.5)$$

where the angular bracket denotes the average of the protein, ligand and protein-ligand complex contribution, which may be calculated from a single simulation of the ligand bound to the protein or from multiple simulations of the complex and the free protein and ligand. End-point methods occupy an intermediate stage in the accuracy-cost trade-off, being more expensive than docking but generally more accurate. However, while useful for quick exploratory studies, this group of techniques is generally not accurate enough for drug discovery projects.

1.4.3 Alchemical methods

The binding free energy is a state function, that is, the free energy difference between two states, in this case the protein-ligand complex and both molecules in solution (unbound conformation), depends only on the values of the initial and final states, and not on the intermediate values. This fact is already partially exploited in the methods explored in the last section, reducing the exploration to only the two end states. One can realise, as demonstrated by early works of Kirkwood (Kirkwood, 1935) and Zwanzig (Zwanzig, 1954), that it is possible to rigorously estimate the binding free energy through a nonphysical path, rather than following the binding or unbinding paths. In this way, simulations can be driven through a series of intermediate states that would not occur naturally but that are easier to sample, and as long as the initial and final states correspond to the physical bound and unbound conformations, the free energy of binding would converge to the correct value. Due to the use of the intermediate nonphysical states, the group of methods described in this section are known as *alchemical methods*.

The two main examples of alchemical methods are *Thermodynamic Integration* (TI) (Kirkwood, 1935) and *Free Energy Perturbation* (FEP) (Zwanzig, 1954). Both methods rely on a thermodynamic cycle similar to the one shown in Figure 1.5. The present cycle is an example of a FEP application, where ligand A is turned into a different ligand B through a nonphysical transformation. One can choose whether to traverse the cycle horizontally (calculating ΔG_{bind}^A and ΔG_{bind}^B) or vertically (calculating $\Delta G_1^{A \rightarrow B}$ and $\Delta G_2^{A \rightarrow B}$). Since traversing the whole cycle implies that the total difference is zero, the difference between the two directions is also zero, allowing to write the binding free energy difference between the two ligands as

$$\Delta\Delta G = \Delta G_{bind}^B - \Delta G_{bind}^A = \Delta G_1^{A \rightarrow B} - \Delta G_2^{A \rightarrow B} \quad (1.6)$$

and

$$\Delta G_i^{A \rightarrow B} = -k_B T \ln \langle e^{-\beta \Delta U_{A \rightarrow B}(x)} \rangle \quad (1.7)$$

where β is the inverse temperature, $\beta = \frac{1}{k_B T}$ and the i sub-index distinguishes which process is taking place, 1 for the change of ligand inside the protein and 2 for the change of ligand in water. Note that the result of FEP is not binding free energy of the two ligands, but the difference between them, which is why FEP is usually referred to as a *relative free energy method*, as opposed to *absolute binding free energy methods* that would provide the free energy estimation for both ligands separately. Relative free energies are often as useful as the absolute counterparts, particularly after other methods such as virtual screening have reduced the number of ligands to evaluate. Furthermore, the availability of implementations such as FEP+ from Schrödinger, with notable predictive performance and accuracy as high as 1 kcal/mol (Wang et al., 2015), or the recent open-source implementation based on Gromacs and PMX (Gapsys et al., 2020) with similar accuracy to that of FEP+, has converted FEP into the gold-standard method for relative free energy calculations.

Despite the popularity and success of FEP, there are some critical points that need to be considered. Mainly, to ensure good convergence, the perturbation introduced to the changing ligand has to be small. That is, ligand A and B have to be very similar, up to 10 heavy atoms (Wang et al., 2015). Importantly, the perturbation must typically avoid net charge modifications, although recent works appear to be able to relax this constraint (Chen et al., 2018). Moreover, FEP calculations are dependent on the initial pose of the ligand. If this pose is incorrect or the ligand can bind the protein with multiple binding poses then the result of the calculation is likely to be incorrect. Finally, the simulations between the intermediate states must be kept short to avoid long computational times, but if protein motion or water network rearrangements are important for ligand binding the simulations might be too short to observe them, limiting the accuracy of the estimation. All in all, while FEP is a successful tool, its limitations prevent it from being universally applied, requiring other complementary methods.




Figure 1.5: Illustration of a typical thermodynamic cycle used in FEP calculations. The relative binding free energy is estimated through the simulation of process 1 and 2, thus avoiding the simulation of the computationally expensive binding process of the two ligands. Extracted with permission from: Limongelli V. Ligand binding free energy and kinetics calculation in 2020. *WIREs Comput Mol Sci.* 2020; e1455. <https://doi.org/10.1002/wcms.1455>

1.4.4 Pathway-based methods

The term *Pathway-based methods* refers to a group of techniques which, contrary to the ones previously discussed, estimate the binding free energy by simulating the binding mechanism of the ligand to the protein. Examples of such techniques include *Umbrella sampling* (Torrie and Valleau, 1977), *Funnel Metadynamics* (Limongelli et al., 2013) or *Steered MD* (Izrailev et al., 1997). Pathway-based methods usually provide the estimation of the binding free energy through the calculation of a potential of mean force (PMF) (see section 1.7 for more details).

Since the simulation of the whole binding or unbinding process is computationally expensive, the methods mentioned make use of some biasing scheme to accelerate the sampling. Steered MD introduces an external force applied to the ligand that pulls it out of the binding site, mimicking an atomic force microscopy experiment. The binding free energy is then recovered using the Jarzynski non-equilibrium work theorem (Jarzynski, 1997). On the other hand, umbrella sampling accelerates the simulations by adding harmonic potentials along one or more reaction coordinates, that is, user-defined dimensions, expressed as function of some degrees of freedom relevant to the system that describe well the binding reaction.

Finally, funnel-metadynamics is a variant of the well known metadynamics method (Laio




Figure 1.6: Illustration of typical funnel metadynamics setup. The light blue area denotes the volume available for the ligand to explore. The narrow cylinder avoids the exploration of many homogeneous ligand solvated states, while the cone allows perfect exploration of the binding site. Extracted with permission from: Limongelli V. Ligand binding free energy and kinetics calculation in 2020. *WIREs Comput Mol Sci.* 2020; e1455. <https://doi.org/10.1002/wcms.1455>

and Parrinello, 2002), which biases the dynamics by depositing history-dependent potential along a few selected degrees of freedom, usually termed *collective variables* (CV). This has the effect of “filling” the already visited minima, enhancing the exploration and the probabilities of finding transitions between states. Funnel metadynamics introduces an extra potential that restricts the volume accessible to the ligand, depicted in Figure 1.6. The shape drawn by this potential, which gives name to the method and is shown in blue in the figure, if implemented using a flat-bottom potential. Thus, when the ligand is inside the delimited space it feels no other perturbation than that of classical metadynamics, exploring freely the binding site. However, after leaving the funnel, the ligand feels a repulsive force, drastically reducing the number of available solvent conformations and improving the convergence of the simulation.

Pathway-based methods have three important advantages when compared with alchemical methods. First, they do not require precise information about the binding mode, just with approximate information about the binding site the active conformation can be recovered from the state of lowest energy. Moreover, pathway-based methods can be used to study ligands with important structural differences without any convergence or performance penalty. Finally, in addition to the free energy estimation, additional kinetic and structural information can be recovered from the simulations, which might be useful for systems where knowledge is limited. On the other hand, some drawbacks can be readily noted. Importantly, pathway-based methods require higher computational times than those of alchemical techniques, given the additional simulation needed. Also, it should be noted that while precise information about the binding mode is not required, typically

pathway-based methods require the user to define a reaction coordinate of collective variable, which will be system-dependent. Moreover, a badly chosen option will hamper the convergence of the estimation or directly provide incorrect results.

All in all, the different weaknesses and strengths of the methods reviewed throughout section 1.4 show that it does not exist a universally applicable technique for the estimation of protein-ligand binding free energies. The problem is still highly system-dependent, and a combination of different techniques is needed to complement each other and cover their weak points.

1.5 Molecular mechanics

Until this point, this dissertation has focused on framing the problem of protein-ligand binding free energies and reviewing its most important advances. For the rest of this chapter, we discuss the methods that will constitute the foundations for the original work developed in this thesis.

Almost all the techniques reviewed in previous sections use, in one way or another, a *physics-based* approach. As the name indicates, these techniques use an approximation based on principles from physics to describe the behaviour of some system, in our case, molecular (atomistic) view of the protein-ligand complex. Physics-based approaches can be roughly classified in two categories: classical or quantum, although techniques that mix the two also exist. Here we will focus on methods based on classical physics, particularly on molecular mechanics (MM).

Molecular mechanics models the molecules at the atom level, assuming that each atom is a solid sphere, connected to other atoms via covalent bonds which are modelled as springs. The atoms possess a radius, called the Van der Waals radius that describes the minimum distance at which another atom can be found and some electrostatic charge. From the combination of all mentioned assumptions, the notion of a force field appears. A force field can be thought of a map from atom positions to an energy value, for example

$$\begin{aligned}
 E = & \sum_{bonds} K_r(r - r_{eq})^2 + \sum_{angles} K_\theta(\theta - \theta_{eq})^2 + \sum_{torsions,n} K_{\phi,n}[1 + \cos(n\phi_i - \delta_n)] + \\
 & + \sum_{i<j} k_e \frac{q_i q_j}{r_{ij}} + \sum_{i<j} 4\epsilon_{ij} \left[\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^6 \right]
 \end{aligned} \tag{1.8}$$

where the first two terms correspond to the energy of the covalent interactions, measured as an harmonic potential of the deviation from the equilibrium values. The third term, models the steric interaction between atoms separated by 3 covalent bonds (also known as 1, 4 pairs), as well as the improper dihedrals, used to enforce planarity where required. The final two terms constitute the so-called non-bonded terms, electrostatic and van der Waals. The former is modelled by a Coulomb potential, while the latter typically uses a 6-12 Lennard Jones potential, which is strongly repulsive at very short distances,




Figure 1.7: Illustration of the main interactions present in a classical force field. Graphs depict the functional shape of some of these terms, while the arrows highlight the extent of the interactions, and which atoms participate. Source: https://upload.wikimedia.org/wikipedia/commons/5/5c/MM_PEF.png

but as the distance increases the attractive term overcomes the repulsion. There are multiple parameters that need to be adjusted in each term, and their values will depend on which force field is used. Prominent examples include AMBER (Maier et al., 2015), OPLS (Jorgensen and Tirado-Rives, 1988), CHARMM (Vanommeslaeghe et al., 2010) or GROMOS (Reif et al., 2012).

In Figure 1.7 the terms described in equation 1.8 are depicted graphically. Missing from said equation is the contribution of the solvent, which can be considered using two models: explicit or implicit waters. In explicit solvent models, the water molecules are incorporated to the simulations, just like the protein and ligand molecules, and their energy contribution is taken into account similarly. There several such models, such as the TIP3P (Jorgensen et al., 1983) or SPC (Berendsen et al., 1981), which differ in parameters like the angle formed by the three atoms or their electric charges. Considering every single water molecule explicitly results in an elevated computational cost, which prompted the creation of simplified representations of the water interactions. Known as continuum models or implicit solvent models, these techniques use an average interaction field instead and eliminate the water molecules from the simulation. Implicit solvent models incorporate this average interaction through an extra term to equation 1.8

$$\Delta G_{solv} = \Delta G_{pol} + \Delta G_{np} \quad (1.9)$$

ΔG_{pol} is known as the polar term, represents the reversible work done to create the charge

distribution of the solute, while ΔG_{np} , the non-polar term, measures the reversible work needed to create the empty space in the solvent needed to fit the solute molecules and the repulsive forces due to van der Waals effects. Approaches based on such approximation, like VDGBNP (Zhu et al., 2007) or OBC (Onufriev et al., 2004), are faster than explicit solvent modelling, but are limited in accuracy when local fluctuations are important, such as in the formation of hydrogen bond interactions between water and solute molecules.

The combination of a force field and a water model provides the energy of a molecular conformation, mapping the atomic coordinates to an energy value. From statistical mechanics, we know that the probability of a given state, x , can be obtained from the Boltzmann distribution

$$P(x) \propto e^{-\beta E(x)} \quad (1.10)$$

To precisely calculate this probability one would need to evaluate a normalisation factor, the partition function, which requires an energy calculation for every possible conformation of the system in question. For a protein-ligand system this is impossible, due to the immense number of degrees of freedom needed to describe it. Nonetheless, one can design strategies to sample the Boltzmann distribution without needing to enumerate the full space of possible conformations. If sufficient samples are drawn then it is possible to estimate probabilities and gain useful information about the system. In the following section we will provide an overview of the two most popular strategies to sample the Boltzmann distribution in protein-ligand systems, *Monte Carlo methods* (MC) and *Molecular dynamics* (MD).

1.5.1 Molecular dynamics

The idea behind molecular dynamics is rather simple. In essence, MD numerically integrates Newton's equations of motion under a potential described by the force field. In practice, however, its application is not so simple. To begin, the equations of motion need to be solved for every atom imposing steep computational demands. Also, the time step used in the integration scheme needs to be small in order to reduce the integration error. In atomistic systems, the time step should be in the order of 1 fs, which is about 10 times smaller than the periodic temporal scale at which the fastest motions, the stretching of the bonds involving a hydrogen atom (such as O-H or N-H), occur. If these bonds are constrained using techniques such as SHAKE (Ryckaert et al., 1977) the time step can be doubled (adjusting it to 1/10 of the fastest bending period).

An MD simulation starts from a 3D structure of the molecules of interest, either from an experimental source (X-Ray crystal or NMR) or a model constructed, for example, through homology modelling. These structures typically have small artifacts that are corrected by running a minimisation step, which also ensures that the forces in the atoms are not too large so that they will not increase dramatically with each time step, causing the simulation to "explode". When the system is in an energy minimum the equilibration process begins, which is necessary to slowly adapt our system dynamics to the desired

experimental (pressure, temperature, etc) conditions.. There is no established sequence for the equilibration procedure, as it depends on the statistical ensemble to simulate (NPT for constant pressure and temperature, NVT for constant volume and temperature or NVE for constant volume and energy are the most common). Typically one equilibration stage includes a thermalisation, that is, the increase of the systems temperature until the target value is reached, and checking whether the thermostat works properly, if the simulation is carried out at constant temperature. If the simulation is run at constant pressure, a barostat needs to be added. Once the equilibration is finished, then the actual simulation is run, usually referred to as the production phase. The equations of motion are integrated until some stopping criteria is met (normally simulation time) and the coordinates are periodically saved.

MD is a highly popular tool due to the excellent spatial and temporal resolution it offers, permitting the observation of phenomena at atomic level and with a femtosecond time scale. It has been often referred to as a “computational microscope”, offering a complementary point of view that is not available in experimental studies. However, despite its popularity, the small time step needed for the integration prevents MD simulations from reaching the long time scales of some slow processes such as ligand binding or protein folding. For this reason, several enhanced sampling methods have been designed to increase the available time scales for MD simulations (for more details on such methods see section 1.5.4).

1.5.2 Monte Carlo methods

Monte Carlo methods were one of the earliest contributions that computing made to the world of science. During the development of the Manhattan project, in Los Alamos, in 1946, the Polish mathematician Stanislaw Ulam came up with a simple procedure to calculate the probability of winning in a game of solitaire without the need of complicated combinatorics calculations (Eckhardt, 1987). Together with John Von Neumann, they quickly applied the new strategy to solve problems like neutron diffusion, using the, at the time quite powerful, ENIAC computer. The core idea was to repeatedly draw random samples of the process of study, so that when enough samples have been collected one can use statistical methods to extract information about the system, even if its evolution is fully deterministic in nature.

The random and statistical nature of MC methods highlight their main difference to MD ones. While MD trajectories follow the time evolution of the system, the concept of time is absent from MC methodologies. Here, the protein-ligand is randomly perturbed, the term perturbation might refer to a wide range of movements such as small few atoms displacement or more complex side chain rotations, to name a few. After the new conformation is applied, its energy is calculated and the proposed move is accepted or rejected according to (typically) the *Metropolis criterion* (Metropolis and Ulam, 1949)

$$P_a = \min(1, e^{-\beta(E_n - E_o)}) \quad (1.11)$$

where P_a is the probability of accepting the move, E_n is the energy of the new configuration and E_o is the energy of the old configuration. After the energy calculation, a random number r is generated and the move is accepted if $r \leq P_a$. Note that if the move improves the energy $E_n < E_o$ it will automatically be accepted. If the attempt is rejected, the old conformation is considered again. The process is iteratively applied and the successive states are joined in a chain, which is why these kinds of methods are known as Markov Chain Monte Carlo (MCMC). If the moves applied satisfy some mathematical conditions, the theory of Markov processes guarantees that the simulation will eventually reach a *stationary or equilibrium distribution*, a probability distribution that is left invariant after random MC moves. In the case of protein-ligand systems, this distribution is the Boltzmann distribution (see equation 1.10). The most common of such conditions is known as *detailed balance* and mathematically is expressed as

$$\pi_i p_{ij} = \pi_j p_{ji} \quad (1.12)$$

where π_i is the probability of state i and p_{ij} is the probability to transition from state i to state j . This equation, based on the physical principle of *microscopic reversibility*, describes a system in which every transition is equilibrated by its reverse. However, some authors have proposed that the detailed balance condition is too strict and a weaker, but more difficult to implement, balance condition would be enough to guarantee Boltzmann sampling (Manousiouthakis and Deem, 1999).

The proposal move is the most important step in a MC simulation. In a protein system generating perturbations that lead to high acceptance is difficult, given that if the moves involve big translations there is a high probability of generating steric clashes. On the other hand, while small moves would avoid the clashes, the obtained sampling would be inefficient, requiring many steps to observe rare events such as ligand binding or conformational changes. Compared to MD, MC sampling has a theoretical advantage in its exploratory efficiency. Since there is no concept of time in a MC move, it is not necessary to surmount energy barriers pushed by thermal fluctuations, the move may take a more efficient path by “walking” around the barrier. In practice, this advantage is typically hindered by the ruggedness of the energy landscape. The binding mechanism is not a simple barrier to overcome, but a complex network of hills of variable height and thus while surpassing any particular barrier is easier than with MD, the overall exploration is as limited in practice. Nonetheless, MC methods are still highly useful in the study of protein ligand interactions with prominent examples being the MCPRO software (Jorgensen and Tirado-Rives, 2005), ProtoMS (Woods et al., 2017) or PELE (Borrelli et al., 2005).

1.5.3 PELE

The Protein Energy Landscape Exploration (PELE) method is a MC method based on the idea that efficient sampling can be obtained by carefully crafting a complex move to increase the acceptance probability, as opposed to most MC methods that attempt simple moves at a higher rate. A PELE move or step is divided into perturbation and relaxation

phases. This two stages design follows the idea found in other methods such as the *Monte Carlo with minimisation* approach (Li and Scheraga, 1987), the *Activation-relaxation technique* (Malek and Mousseau, 2000) or *basin hopping* (Wales and Doye, 1997). By introducing an energy minimisation step (the core of the relaxation stage) the mentioned methods can explore more efficiently the multiple energy minima of the system. The main difference between PELE and the mentioned methods is the introduction of protein structure prediction functions in each MC move, which helps to keep the acceptance probability high.

The different steps of the PELE algorithm are illustrated in Figure 1.8. An MC step starts with a random translation and rotation of the ligand, which is modelled by a rigid core connected to a set of rotatable side chains. Several perturbed poses for the ligand are generated, and the one with the lowest energy is chosen. Optionally, the direction of movement of the ligand might be reused for several steps, to help sample rare events if the system is near a transition state, this option is named *steering*. Additionally, often the space available to the ligand is restricted, for example to a sphere; this restricted space is usually called the *simulation box*, and if during the perturbation the ligands steps outside the move is automatically rejected. After the ligand, the protein structure might also be perturbed. This is done either using an Anisotropic Network Model (ANM) (Atilgan et al., 2001) or by using a set of predefined normal modes which can be extracted from a group (more than 2) X-ray structures or from a previous MD simulation (for example using principal component analysis). Typically, a combination of several of the lowest frequency modes are used, although the option to use a single mode or a random mixture also exists. This phase is usually applied every few steps to reduce the run time of the simulation, and to emphasise local rapid movements (ligand and side chain) over slower ones (backbone reorganisation). After the perturbation, energy minimisation is applied with the alpha carbons constrained, so the backbone can be reorganised following the perturbation.

After the perturbation, the relaxation phase begins by running side chain prediction techniques. The main goal of this step is to sample local interactions (hydrogen bonds, π - π stacking, etc.) and possible small openings that may facilitate the ligand binding. The side chain prediction acts on those residues that are close to the ligand (by default $\sim 6\text{\AA}$), and could also include the ligand itself and residues that have suffered a large increase in energy during the perturbation. The side chain prediction proceeds on one residue at a time, keeping the others fixed. Following the algorithm of Jacobson et al. (Jacobson et al., 2002) a rotamer library is considered, the possible conformations are clustered and the best one is selected. This is iterated a few times for each residue until energy changes are negligible or a maximum number of steps is reached. Afterwards, a partial minimisation of the side chain atoms is performed. However, if no suitable conformation is found the MC step is rejected.

Finally, an energy minimisation is applied to the whole system to relax the structure and improve the acceptance probability. Since the minimisation prevents the exploration of some regions, the detailed balance condition is broken. To alleviate this problem, the convergence criteria for the minimisation is significantly relaxed as discussed in Daniel Lecina's doctoral thesis (Lecina, 2017), allowing the partial recovery of detailed balance.




Figure 1.8: Schematic representation of the different stages in a PELE step. Image author: Ryoji Takahashi

PELE can use the OPLS2005 force field (Kaminski et al., 2001; Banks et al., 2005) or the AMBER99sbBSC0 (Pérez et al., 2007) with two options for the implicit solvent model, OBC (Onufriev et al., 2004) and VDGBNP (Zhu et al., 2007), however, some discrete water molecules can be placed if they are important for the protein-ligand dynamics and will be moved during the final minimisation step. Here we have reviewed the aspects of PELE that are most relevant for the work done in this thesis, but a more thorough view of PELE can be found in a recent book chapter (Gilabert et al., 2018).

1.5.4 Enhanced sampling

Both MD and MC methods are routinely applied in the study of protein-ligand systems, thanks to the extensive work done in providing user-friendly software packages, together

with more efficient software and hardware implementations. In particular, the hardware side has driven most of the improvements with the wide-spread use of Graphical Processing Units (GPU) to calculate the expensive but easily parallelisable operation involved in a MD simulation (Harvey et al., 2009; Kutzner et al., 2015; Lee et al., 2018). Other significant contributions involve the design of specialised machines like Anton (Shaw et al., 2008) or MDGRAPE (Susukita et al., 2003) which have achieved simulation lengths in the order of milliseconds (Shaw et al., 2010) and the development of distributed computing networks such as GPUGRID (Buch et al., 2010) or Folding@Home (Shirts and Pande, 2000) that take advantage of a wide range of volunteer devices that perform molecular dynamics simulations when idle. Despite the groundbreaking advancements, regular molecular simulations are often not powerful enough to obtain robust sampling of rare events, which has led to the development of *enhanced sampling* methods.

The term enhanced sampling refers to the set of methods designed to accelerate a molecular simulation either by modifying the potential energy surface of the system of interest or by forming a model that explains the dynamics and iteratively improving it. In the first group we find some methods that have been already discussed in section 1.4.4, metadynamics and umbrella sampling. As mentioned before, these methods add history-dependent biasing potentials in some user-defined coordinates making it easier for the simulation to overcome free energy barriers. Similarly, in this group of methods one can find metadynamics variants such as *well-tempered metadynamics* (Barducci et al., 2008), where the bias strength is gradually decreased to improve the convergence properties and avoid overfilling the energy minima, or *infrequent metadynamics* (Tiwary and Parrinello, 2013) which allows the extraction of kinetic information by reducing the frequency at which the bias is applied and assuming that the time spent on a transition state is much less than the waiting time between transitions. In the same group, one can find methods that instead of adding biasing interactions modify the underlying free energy landscape. Prime examples here are the *basin-hopping* (Wales and Doye, 1997) and *basin-sampling* methods (Wales, 2013). These techniques explore a modified version of the FEL, as illustrated in Figure 1.9. By running an energy minimisation after every step, the shape of the underlying landscape is simplified, removing barriers and improving the sampling efficiency.

A second group of methods, contrary to the ones outlined above, leaves the FEL untouched, but accelerates sampling by adaptively focusing on more interesting regions and forming a better model of the landscape. For example, *Adaptive Sampling* methods (Singhal and Pande, 2005; Doerr and De Fabritiis, 2014; Bowman et al., 2010) follow an iterative algorithm with three steps: i) run N short simulations, ii) build a Markov State Model (see section 1.6 for more details) of the simulations run so far, and iii) use the information provided by the MSM to select the most interesting states and start a new round of simulations from them. The criteria used to select those states is commonly referred to as *reward function*, drawing inspiration from Reinforcement learning (RL) techniques. In RL, an agent is trained to perform a task by maximising the reward obtained by its actions. The agent must navigate the *exploration-exploitation dilemma*, that is, if it is too greedy and just gets the highest rewards possible it will quickly get stuck in a local minimum, failing to obtain the maximum possible reward. This is why in complex problems a balance between exploiting current information and gaining new information is needed




Figure 1.9: Graphical representation of the effect of basin-hopping techniques on the FEL, reducing significantly the free energy barriers by the use of energy minimisation. Source: https://esa.github.io/pagmo2/_images/mbh.png

(Auer et al., 1995). The prime example of the application of ideas from RL in molecular simulations is the *Fluctuation Amplification of Specific Traits* (FAST) which runs iterations of simulations, clustering and selection based on the following reward function

$$r_i = \phi_i + \alpha\psi_i \quad (1.13)$$

where r_i is the reward of a particular state i , ϕ_i is a score design to favour exploitation towards states that optimise some structural metric, ψ_i is a score proportional to the exploration accumulated in each state, favouring the sampling of less populated states and α is a parameter that can be tuned to direct the simulation towards either side of the exploration-exploitation trade-off. Similarly, the *REinforcement learning based Adaptive samPling* (REAP) method (Shamsi et al., 2018) which uses a weighted combination of several reward functions that can change during the simulation, facilitating the sampling of rare events that need more than one order parameter to be successfully described.

More mainstream kinds of machine learning techniques have also been used to enhance sampling. For example, the *Reweighted Autoencoded Variational Bayes for Enhanced sampling* (RAVE) (Ribeiro et al., 2018) utilises a state-of-the-art deep learning approach called variational autoencoder (VAE) to estimate a low-dimensional representation of the protein-ligand dynamics and run successive iterations of new simulations to refine the model. RAVE has been applied to the estimation of the binding free energy of benzene from the L99A variant of lysozyme, with notable success (Lamim Ribeiro and Tiwary, 2019). In short, the neural network creates a low-dimensional representation to guide the dynamics, thus avoiding the pitfalls of exploring the full energy landscape. This reduced space, referred in the literature as “latent space”, has been introduced previously as a helpful trick to accelerate sampling in the *Intrinsic Map Dynamics* (iMapD) technique (Chiavazzo et al., 2017), and is also exploited in a recent promising method termed *Boltzmann generators* (Noé et al., 2019). In this technique, a neural network encodes the

system coordinate into a latent space, which is explored via Monte Carlo moves and then can reverse the transformation, obtaining conformation sampling at an accelerated rate, as well as estimates of thermodynamic properties.

Finally, in this group one can find the family of works belonging to the category of *Weighted Ensemble* (WE) techniques (Zuckerman and Chong, 2017; Dickson and Brooks, 2014; Zwier et al., 2015). Similarly to Adaptive Sampling, in WE an ensemble of short unbiased simulations is run. In the beginning, the multiple copies in the ensemble have an equal weight. As the simulations explore more states the statistical weights are updated and the replicas are merged or split according to the weights prioritising those states with less sampling. WE methods have been applied to the study of kinetics of protein-ligand systems (Dickson and Lotz, 2016; Dickson and Lotz, 2017; Dickson, 2018), successfully studying the unbinding of 11 min time scale with only 6 microseconds of molecular dynamics (Lotz and Dickson, 2018).

1.6 Markov State Models

The increase in parallel computers mentioned at the beginning of the previous section has facilitated the surge in popularity of *Markov State Models* (MSM). While the theoretical groundwork was first established in 1983 (Zwanzig, 1983), it was not fully developed until almost 20 years later (Swope et al., 2004). Nonetheless, the framework experienced, shortly after, a sudden increase of popularity and attention, receiving great advances in capabilities and robustness (Pande et al., 2010; Husic and Pande, 2018; Chodera and Noé, 2014).

A Markov State Model represents a model of the dynamics of a system that is assumed to be *Markovian* at a lag time τ . The Markovian property establishes that the transitions between different states are memory-less, that is, the transition probabilities do not depend on the past history, but only on the current state. The construction process of a MSM is outlined in Figure 1.10. First, a simulation or ensemble of simulations is run. The trajectories are discretised into a set of M clusters. Here, we will assume that the state decomposition corresponds to a “crisp” assignment, where each conformation is assigned to the closest state, creating a Voronoi diagram of the simulation on the coordinates used for the clustering, commonly known as *features*. This assumption, however, is not necessary and the estimation of MSMs has been done with a state discretisation that does not create a full partition of the state space using *milestoning* (Schütte et al., 2011). After the clustering, we use a sliding window approach to generate a count matrix from the discrete trajectories (see panel (d) of Figure 1.10). That is, with trajectories indexed from 0 to $L - 1$, where L is the length of the trajectory, we label as i the state at index 0 and j the state at index τ and add one to the matrix entry in row i and column j , and we repeat the procedure for indexes $(1, \tau + 1)$, $(2, \tau + 2)$ until $(L - \tau, L)$. From the count matrix, a transition matrix is estimated, which gives the probability to jump from one state to another, following the Markovian assumption.

The transition matrix is the core of an MSM, as it permits the extraction of plenty of information. If the matrix is diagonalised we obtain a set of eigenvectors and eigenvalues.




Figure 1.10: Schematic representation of a MSM. Panel (a): Four-well energy potential that is sampled. Panel (b): The MSM eigenvectors approximate to the real dynamics’ eigenfunctions. Panel (c): The relaxation time scales estimated from the MSM converge to the exact time scales, with a rate depending on many factors, such as the discretisation. Panel (d): Outline of the construction of a MSM. Trajectories are discretised, from which a transition matrix is estimated. The largest eigenvector approximate the stationary distribution of the dynamics. Extracted with permission from: Chodera, J. D. and Noé, F. (2014). Markov state models of biomolecular conformational dynamics. *Current Opinion in Structural Biology*, 25:135–144.

The eigenvectors describe “probability fluxes”, that is, how the probability distribution that describes the dynamics is affected by a given dynamical process (Prinz et al., 2011), while the eigenvalues establish the time scale at which the processes decay (Swope et al., 2004), according to the formula

$$t_i = \frac{-\tau}{\ln |\lambda_i(\tau)|} \quad (1.14)$$

where t_i is the associated time scale, also known as implied time scale, λ_i is the eigenvalue and τ is the lag time used. Note that the largest eigenvalue will be 1, and its associated eigenvector is called the *stationary distribution*, π . This eigenvector represents the equilibrium distribution of the Markov chain, meaning that when the transition matrix is applied to a system initialised from π , this distribution is left invariant. The rest of eigenvalues are smaller than one, representing dynamical processes that will eventually decay, leaving only the stationary distribution.

Markov State Models have two main advantages: i) they facilitate the integration of an ensemble of independent, usually short, parallel simulations into a single model, instead

of a long simulation, and ii) the obtained model allows the extraction of kinetic and equilibrium information. There are, however, some weak points to the method: i) the spatial discretisation greatly affects the model quality, one must choose a set of coordinates that represent well the dynamics of the system. The choice of a clustering algorithm and other parameters such as the number of states can be difficult. And ii) the lag time chosen should be as large as the time scale of the slowest dynamical process of the system. Since this is not known beforehand, a method for assessing whether the lag time is appropriate exists. Proposed by Swope et. al. (Swope et al., 2004) this technique requires estimating several MSM at increasing lag times, and checking at which values the implied time scales converged to a constant value, since if a model is Markovian at a lag time τ it will also be Markovian at any lag time larger than that (see panel (c) of Figure 1.10 for an example of this technique).

The process of construction of an MSM like the one outlined above is the most common one, however recent developments in the field have open the door to alternatives modes of construction that can alleviate some of the weak points described. The selection of features for the space discretisation is difficult and system-dependent. In an attempt to simplify and automate the process, dimensionality reduction techniques were introduced. With these methods, one could use a large number of coordinates from the system and run a *Principal Component Analysis* (PCA) (Mu et al., 2005) or *Time-lagged Independent Component Analysis* (TICA) (Naritomi and Fuchigami, 2011; Pérez-Hernández et al., 2013; Schwantes and Pande, 2015) to obtain a simplified set of features on which the discretisation is carried out. Building on the groundwork of the TICA method, Nüske et. al. developed the *Variational Approach to Conformational dynamics* (VAC) (Nüske et al., 2014), which presented a variational approach to building MSMs. According to the VAC, models could be objectively scored by the magnitude of the eigenvalues. Since in an MSM the eigenvalues are systematically underestimated, the best models would maximise the magnitude of their eigenvalues. Using this principle, McGibbon and Pande introduced the *Generalized Matrix Rayleigh Quotient* (GMRQ) score that allowed the comparison between different discretisation schemes and introduced the concept of cross-validation, ubiquitous in machine learning fields, for a more robust validation of the sampling quality of a simulation (McGibbon and Pande, 2015). In 2018 *VAMPNets* were introduced (Mardt et al., 2018), this technique uses a neural network trained to maximise the variational score and estimate an MSM, without the need of user interaction. With the introduction of ideas from artificial intelligence fields, the construction of MSMs has become more user-friendly and automatable which will surely increase the popularity of the technique.

1.7 PMF, probability distribution and ΔG

The construction of potentials of mean force is ubiquitous in free energy calculations, particularly in pathway-based methods. A PMF, denoted as $W(\chi)$, describes the free energy changes as a function of a coordinate or a small subset of coordinates of the system, χ . These coordinates can be a subset of the natural coordinates of the system, for example the coordinates of some atom, or more complex combination, such as the centre of mass

of a molecule or the distance between the protein and ligand. Numerically, the PMF can be understood as the projection of the free energy onto the subset of coordinates χ . With this in mind, the PMF can be defined as

$$e^{-\beta W(\chi)} = \frac{\int dx \delta(\chi'(x) - \chi) e^{-\beta E(x)}}{\int dx e^{-\beta E(x)}} \quad (1.15)$$

where $\delta(\chi'(x) - \chi)$ is the Dirac delta function for the coordinate χ , x represents the $3N$ coordinates of the atoms of the system and the integrals perform Boltzmann averaging over all the coordinates, leaving as a result the projection of the free energy over χ . For an illustrative discussion of the PMF derivation, one can review the introductory-level discussion by Zuckerman (Zuckerman, 2013), and for a more detailed view one can consult the reviews by Gallicchio and Levy (Gallicchio and Levy, 2011) and Zhou and Gilson (Zhou and Gilson, 2009). The right-hand side of equation 1.15 describes a probability distribution over the coordinates χ , $\pi(\chi)$, allowing to establish a connection between the PMF and probability

$$W(\chi) = -k_B T \ln(\pi(\chi)) \quad (1.16)$$

Examples of the kind of coordinates used for PMF building include protein-ligand distances (Gumbart et al., 2013; Deng et al., 2018), the z coordinate after aligning the system to the z axis (Doudou et al., 2009) or the RMSD to a reference conformation (Woo and Roux, 2005; Wang et al., 2006a; Gumbart et al., 2013). These works have several common points: i) the PMFs were one-dimensional and ii) they used restraints on the ligand and the protein to restrict the exploration to conformations relevant for binding. The restraints were important to reduce the computational cost, but introduce a higher degree of complexity in the setup. The two points were addressed in a work by Buch et al. (Buch et al., 2011), where they run an ensemble of molecular dynamics simulations of the trypsin-benzamidine system, with an aggregated simulation time of $50 \mu s$. The ligand was allowed to diffuse freely within a restricted volume delimited by a box around the binding site. The simulations were analysed using a Markov State Model, clustering using the coordinates of the C7 atom of benzamidine. Using equation 1.16, a three-dimensional PMF was estimated from the MSM stationary distribution. The standard binding free energy was then calculated as

$$\Delta G_b^o = -k_B T \ln \frac{V_b}{V_0} - \Delta W \quad (1.17)$$

with standard volume $V_0 = 1661 \text{ \AA}^3$, $V_b = \int_b e^{-\beta W(r)} dr$ is the average sampled bound volume and ΔW is the depth of the PMF (Zhou and Gilson, 2009), which is defined as the value of the PMF in the “bulk” region, that is, in the region where the ligand is not interacting with the protein and the PMF is flat. The value is determined after setting the minimum of the PMF to zero and shifting all the values by the original minimum value. This work was, to the best of our knowledge, the first one to combine MSM methodologies with a PMF to calculate binding free energies. The simpler setup with respect to previous

PMF calculations offers the advantage of being automatable, if enough computational resources are available.

Despite the versatility of the PMF technique, studies with large benchmarks including multiple proteins and ligands are difficult to find in the literature, especially so for those using MSM methodologies. Some examples include a study employing funnel metadynamics to estimate the binding affinity of 2 ligands to the human peroxiredoxin 5 protein (Troussicot et al., 2015). Another example demonstrated the estimation of the binding free energy of six ligand fragments to the FKBP protein using extensive equilibrium simulations (Pan et al., 2017). Despite these few enlightening examples, to the best of our knowledge, there is no publication describing the prediction of absolute binding free energies for multiple receptors with numerous ligands per receptor using unbiased simulations. The absence of further studies in this direction illustrates the important limitations faced by unbiased simulations, and reinforces the need for smarter enhanced sampling strategies that can make accurate free energy prediction techniques efficient enough for their routine use in real-case scenarios.

1.8 Protein-ligand systems used throughout this thesis

Often, when developing a new method, be it for free energy calculations or for molecular simulations, tests are run with a small set of systems that are simple and well studied, such as alanine tripeptide for protein dynamics, villin for protein folding or host-guest systems for protein-ligand binding. This is a sound strategy for initial tests, since it is easy to derive testable assertions for such systems. However, one should be careful to not over-trust such tests, because these simple systems are not entirely representative of the complex nature of protein systems, and there will be many characteristics that will not be properly captured by them.

Focusing on binding free energies, Mobley and Gilson reviewed common challenges that should be tested against (Mobley and Gilson, 2017). For example, sampling varying degrees of protein conformation changes, different ligand binding modes or different protonation states. It would be useful to create standardised benchmarks of systems that raise some of the mentioned issues, so that when a new method is designed it can be tested against them. Of course, this cannot be the work of a single group, but it must be a community effort. In this thesis, we have strongly considered the issues mentioned and, while we have not built or proposed a full set of benchmarks, we have selected widely different systems to test the methods developed. In this section we will overview, in no particular order, the different proteins chosen, and discuss what makes them relevant for our use cases.

1.8.1 Trypsin

Trypsin (TRP) is a serine protease, found mostly in the small intestine (Rawlings and Barrett, 1994). Its function is to catalyse the proteolysis of ingested protein, that is,




Figure 1.11: Panel (a) Trypsin-benzamidine complex. Panel (b) Plasminogen kringle domain 1 with L01 as a ligand. Panel (c) Urokinase-type plasminogen activator with 39L as a ligand. The protein is shown in cartoon representation, coloured in white, while the ligand is shown in licorice representation, with carbon coloured cyan, oxygens red, nitrogens blue and hydrogens white. The cofactor for trypsin, a calcium ion is shown as a cyan sphere.

the breakdown of proteins into smaller peptide chains, in the digestive system. Due to its small size and considerable rigidity, its binding to benzamidine has been extensively studied computationally (Doudou et al., 2009; Buch et al., 2011; Plattner and Noé, 2015; Takahashi et al., 2014). In Figure 1.11a, the trypsin-benzamidine complex is shown. The structure, taken from PDB entry 3PTB, shows a wide, solvent-exposed binding site, easily accessible for the small benzamidine ligand. This fact made trypsin a perfect initial test for both enhanced sampling exploration and free energy calculations.

1.8.2 Plasminogen kringle domain 1

The plasminogen kringle domain 1 (Plasmin) is one of the kringle domains of plasminogen. In the kringle domains lysine binding sites can be found, which are closely linked to blood clotting action. The plasminogen protein is involved in blood protein degradation, which is why it is an interesting drug target in situations with heavy blood loss, such as during surgery (Cheng et al., 2014). Plasmin is the simplest protein used throughout this thesis, with only 80 residues, as can be seen in Figure 1.11b. The structure, with PDB id 4CIK, shows a wide-open active site, where fragment-like ligands can bind. This system has served as a simple test for the binding free energy studies, with 16 ligands (see Table 1.1 for ligand structures) with binding free energies between -2.1 and -3.9 kcal/mol (Schmidt et al., 2017). Despite its simplicity, the reduced range of affinities makes it a quite challenging system to model computationally.

1.8.3 Urokinase-type plasminogen activator

The urokinase-type plasminogen activator, also known as urokinase (URO), is a protein of the family of the serine proteases. Its main function is the catalysis of the activation of plasmin through cleavage of plasminogen (Jankun and Skrzypczak-Jankun, 1999; Matthews et al., 2011). Urokinase is an interesting pharmaceutical target due to its rel-




Figure 1.12: Panel (a) Complex of progesterone with its endogenous ligand, progesterone. Panel (b) Complex of Mitogen-activated protein kinase 1 with EK2 as ligand. The protein is shown in cartoon representation, coloured in white, while the ligand is shown in licorice representation, with carbon coloured cyan, oxygens red, nitrogens blue and hydrogens white.

evance for blood clotting conditions as well as its possible relation to cancer migration (Tang and Han, 2013). Structurally, URO is similar to trypsin (compare Figure 1.11a to Figure 1.11c) with a solvent exposed cavity. However, the ligands used in this thesis in complex with URO are generally strong binders, with affinities between -5.16 and -12.70 kcal/mol (Dunbar et al., 2013) (see Table 1.2 for the ligand structures), making URO a more complex test case than trypsin. The protein structure is extracted from PDB entry 4FU9.

1.8.4 Progesterone receptor

The progesterone receptor (PR) is a nuclear hormone receptor capable of sensing the present of the hormone progesterone. Upon binding to the endogenous ligand, the receptor undergoes a conformational change that triggers the transcription of mRNA. The progesterone receptor has been studied as an important target for diseases such as diabetes, cancer or heart conditions (Beato and Klug, 2000; Smith and Muscat, 2005; Conzen, 2008; Wang et al., 2010). PR represents probably the most challenging system of the ones used in this work, due to its buried and narrow binding site (see Figure 1.12a). Sampling binding events requires the occurrence of low frequency concerted protein motions, alongside with local loop rearrangements. The five ligands used are shown in Table 1.3 with a range of experimental binding free energies between -5.3 and -9.1 kcal/mol (Grebner et al., 2017). The protein structure is extracted from PDB entry 1A28.

1.8.5 Mitogen-activated protein kinase 1

The mitogen-activated protein kinase 1 (ERK2) is a member of the MAP kinase family. These proteins mediate in multiple signalling pathways by phosphorylating nuclear targets. This receptor is involved in many processes such as cellular proliferation or transcription regulation and mutations of this receptor have been found in various cancer types (Owaki et al., 1992). The structure and binding site of ERK2 differs from all remaining proteins used in this work, as can be seen in Figure 1.12b. The binding site is partially solvent exposed and broad, fitting several different binding modes. The experimental affinities for the five ligands are in the range of -7.0 to -9.7 kcal/mol (Dunbar et al., 2013), see Table 1.4. The receptor was extracted from PDB entry 4FV3.

Table 1.1: Two dimensional structures of the ligands of the plasmin receptor, alongside with their experimental affinities

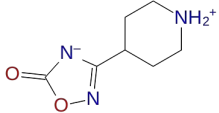
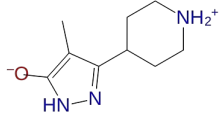
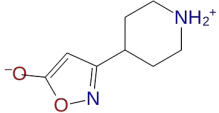
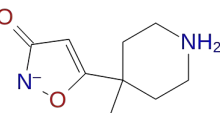
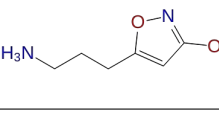
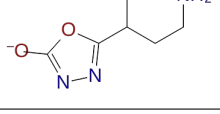
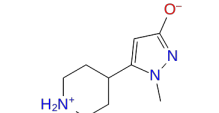
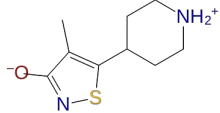
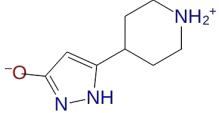
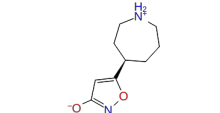
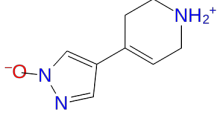
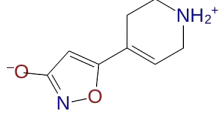
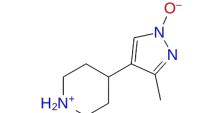
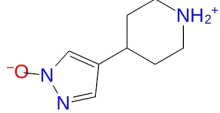
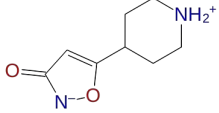
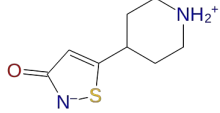
	
L05, $\Delta G_{exp} = -2.16$ kcal/mol	L12, $\Delta G_{exp} = -2.43$ kcal/mol
	
L06, $\Delta G_{exp} = -2.70$ kcal/mol	L02, $\Delta G_{exp} = -2.70$ kcal/mol
	
L13, $\Delta G_{exp} = -2.92$ kcal/mol	L09, $\Delta G_{exp} = -3.02$ kcal/mol
	
L10, $\Delta G_{exp} = -3.24$ kcal/mol	L08, $\Delta G_{exp} = -3.24$ kcal/mol
	
L11, $\Delta G_{exp} = -3.40$ kcal/mol	L04, $\Delta G_{exp} = -3.40$ kcal/mol
	
L14, $\Delta G_{exp} = -3.46$ kcal/mol	L03, $\Delta G_{exp} = -3.62$ kcal/mol
	
L16, $\Delta G_{exp} = -3.67$ kcal/mol	L15, $\Delta G_{exp} = -3.73$ kcal/mol
	
L01, $\Delta G_{exp} = -3.78$ kcal/mol	L07, $\Delta G_{exp} = -3.94$ kcal/mol

Table 1.2: Two dimensional structures of the ligands of the URO receptor, alongside with their experimental affinities

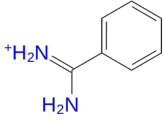
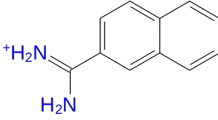
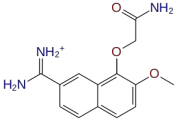
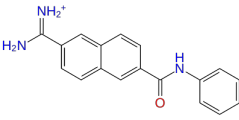
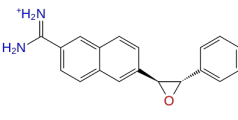
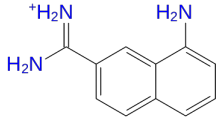
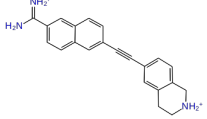
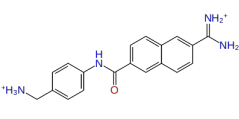
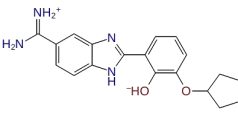
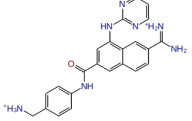
	
BAM, $\Delta G_{exp} = -5.16$ kcal/mol	2UP, $\Delta G_{exp} = -6.50$ kcal/mol
	
1UP, $\Delta G_{exp} = -7.70$ kcal/mol	675, $\Delta G_{exp} = -7.71$ kcal/mol
	
4UP, $\Delta G_{exp} = -7.73$ kcal/mol	6UP, $\Delta G_{exp} = -7.89$ kcal/mol
	
7UP, $\Delta G_{exp} = -8.99$ kcal/mol	239, $\Delta G_{exp} = -9.20$ kcal/mol
	
655, $\Delta G_{exp} = -9.20$ kcal/mol	UI1, $\Delta G_{exp} = -12.70$ kcal/mol

Table 1.3: Two dimensional structures of the ligands of the PR receptor, alongside with their experimental affinities

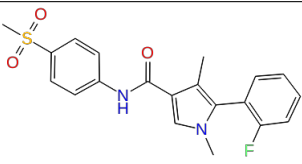
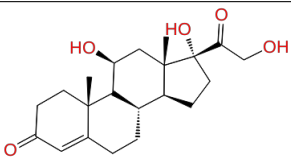
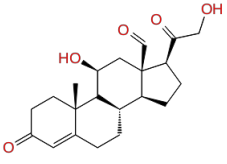
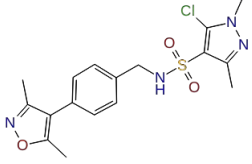
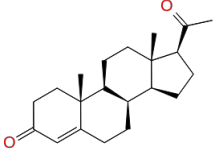
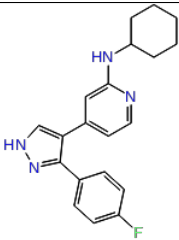
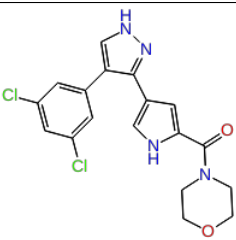
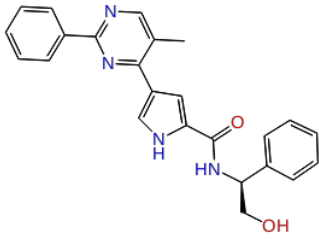
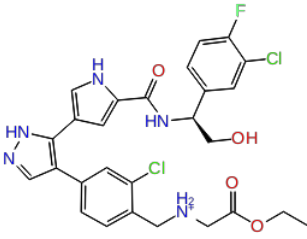
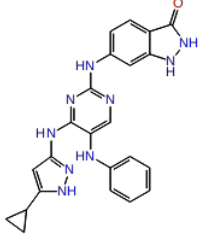
	
356, $\Delta G_{exp} = -5.35$ kcal/mol	HCY, $\Delta G_{exp} = -6.65$ kcal/mol
	
AS4, $\Delta G_{exp} = -8.21$ kcal/mol	400, $\Delta G_{exp} = -8.91$ kcal/mol
	
STR, $\Delta G_{exp} = -9.08$ kcal/mol	

Table 1.4: Two dimensional structures of the ligands of the ERK2 receptor, alongside with their experimental affinities

	
EK3, $\Delta G_{exp} = -6.95$ kcal/mol	EK2, $\Delta G_{exp} = -7.57$ kcal/mol
	
EK9, $\Delta G_{exp} = -9.57$ kcal/mol	EK6, $\Delta G_{exp} = -9.62$ kcal/mol
	
E63, $\Delta G_{exp} = -9.72$ kcal/mol	

Chapter 2

Results

In this chapter we report the main results obtained from this thesis. They contain a description of the methodology developed for the estimation of binding free energies and the application and accuracy measurements of said work. Section 2.1 includes a description of AdaptivePELE, as well as the first attempts at binding free energy estimation. Section 2.2 then explains the multi-stage protocol developed to accelerate the affinity calculations. Sections 2.3 and 2.4 cover the application of the mentioned protocol using different simulation techniques (PELE and MD, respectively). Finally, section 2.5 contains a preliminary attempt to reintroduce AdaptivePELE as the simulation propagator. During some of these sections there will be brief discussions about some relevant aspects of the results. These discussions are included for easier reading, and a more detailed one will be done in the next chapter.

2.1 AdaptivePELE

PELE has been successfully employed in exploring protein-ligand binding (Santiago et al., 2016; Edman et al., 2015; Grebner et al., 2016), as well as for free energy calculations (Takahashi et al., 2014; Grebner et al., 2017). However, all these examples involved testing few ligands and had a significant computational cost. AdaptivePELE was designed to overcome this limitation and permit the study of larger sets of protein-ligand systems with PELE.

2.1.1 AdaptivePELE algorithm

AdaptivePELE is an enhanced sampling technique that belongs in the same group as Adaptive Sampling or FAST. These methods, reviewed in section 1.5.4, are capable of accelerating molecular sampling without the need of adding any biasing potential. When compared with regular PELE simulations, AdaptivePELE shows an order of magnitude speed-up in the mapping of binding events (Lecina et al., 2017), consistently for widely different protein systems, from a simple trypsin protein to much more complex G-Protein

Coupled Receptors (GPCR).

The idea behind AdaptivePELE is based in reinforcement learning, a subfield of artificial intelligence that studies how agents can learn to solve a problem by maximising the rewards obtained from their actions (Sutton and Barto, 1998). More precisely, the theoretical framework discussed here is based on the *Multi Armed Bandit* (MAB) problem (Robbins, 1952). The most common description of the MAB problem requires maximising the money obtained from playing with k slot machines. The player can only pick one machine at a time and the rewards and winning probabilities of the machines are different and unknown to the player. If k is sufficiently big, it is easy to see that discovering the optimal strategy is difficult. The player is faced with the exploration-exploitation dilemma, where if he always plays the machine with the biggest known reward, the player might miss out on another machine with an even bigger reward, but if he keeps trying more machines to discover the maximum one the total obtained prize will be low.

The lessons learned from the MAB can be easily mapped to protein ligand systems by use of the free energy landscape. Ideally, the basins of the FEL would take the role of machines, and their reward would be the basin depth. However, this information is not available a priori, thus in AdaptivePELE the agent chooses between a set of different states, with a reward that is proportional to some metric that can approximate a reaction coordinate. The full AdaptivePELE algorithm, shown in Figure 2.1, is divided in three phases: propagation, clustering, and spawning. These steps are run for a number of iterations, also called epochs, or until an exit condition is met. The different steps will now be carefully explained.

Propagation. A number of PELE simulations, n , is run in parallel, either from the initial structures or from the conformations selected in a previous iteration. This is depicted in the top right image of Figure 2.1. In this example, all simulations start from a single initial state (white circle) and explore slightly different regions of the energy landscape.

Clustering. After the propagation phase, the simulation trajectories are clustered into a set of different states. Clustering uses the leader algorithm (Hartigan, 1975), a technique that is well suited for the incremental nature of the method. In this algorithm, one starts with an empty list of clusters and the first snapshot of the trajectories is added as the first cluster's representative structure. The subsequent snapshots are then compared to the existing clusters, and stops if the structure is deemed similar enough to the representative structure of a cluster, and the population of said cluster is incremented. If the snapshot is different to all existing states, it is assigned to a new cluster. The similarity criteria can be either the ligand root-mean-square deviation (RMSD) between the representative structure and the new conformation or the difference between the protein-ligand contact maps of the two structures.

Spawning. The most important phase of the AdaptivePELE algorithm is the spawning step. At this stage, the list of clusters is evaluated and a reward score is assigned to them. There are many possible schemes for this calculation, but the most commonly used ones available in AdaptivePELE are called *inverselyProportional* and *epsilon*. The first assigns a score that is inversely proportional to the population of the clusters, thus favouring the exploration of under-sampled regions of the energy landscape. The epsilon




Figure 2.1: Schematic representation of the AdaptivePELE algorithm. The algorithm starts from a set of conformations and runs short simulations (propagation phase, top right). Then the trajectories are clustered according to the ligand structure (middle right) and finally some promising states are selected for a new round (bottom right), until an exit condition is met. Image author: Daniel Lecina.

strategy combines the same exploration component of the inverselyProportional with a secondary approach to exploit known information. The method assigns a fraction ϵ of the n selected structures, to those clusters that maximise (or minimise) a certain metric, such as the ligand solvent-accessible surface area (SASA) or the distance between a pair of atoms. Based on the commonly known ϵ -greedy (Sutton and Barto, 1998) methodology of RL, the combination of a exploration oriented score with an exploitation one results in a more efficient sampling, particularly for values of ϵ of approximately up to 0.25 (Lecina et al., 2017). In addition to the score calculated by the spawning methods mentioned, AdaptivePELE optionally incorporates an extra term called *density* designed to compensate the difficulty of generating acceptable poses when the ligand is in close contact with

the protein, both due to steric limitations and the increased ruggedness of the energy landscape. This density is commonly calculated as

$$\rho = \begin{cases} \frac{64}{(-4r+6)^3} & \text{if } r \leq 1 \\ 8 & \text{if } r > 1 \end{cases} \quad (2.1)$$

where the density, ρ , is a function of a parameter $r = \frac{c}{l}$, where l is the number of the ligand heavy atoms and c is the number of alpha carbons of the protein in contact with the ligand. After the overall score is calculated, the n top states are filtered and a number of new trajectories is assigned to each state according to the score. Since in a typical simulation the number of clusters can get relatively large (numbers between 500 and 1000 are common), filtering the states ensures that resources are focused in regions of interest.

Exit condition. Finally, the structures selected in the spawning step are prepared to be used in the next iteration. However, before that the codes check for an exit condition. This is typically a number of epochs or iterations, but AdaptivePELE also offers the possibility of stopping the simulation if a maximum or minimum value is reached for a metric. For example, one might be interested in simulating the unbinding of a protein-ligand system, so a suitable exit condition would consist of stopping once a value for the SASA close to one is obtained, or once a certain number of trajectories have encountered such a value. If no exit condition is triggered, then the algorithm proceeds by preparing the structures for the next round of simulations, and goes back to the propagation phase.

2.1.2 Binding free energies from AdaptivePELE simulations

As mentioned in the previous section, PELE has already been used in conjunction with Markov State Models to estimate a potential of mean force and calculate binding free energies (Takahashi et al., 2014). In that paper, the authors accurately predicted the binding free energy of trypsin and three structurally similar ligands, which required approximately 1 week of simulation time per ligand with 64 processors. Although the time could be substantially reduced increasing the number of processors, the computational cost prevented the technique from being scalable to a higher number of ligands. For this reason, we set out to study whether the simulation time necessary could be reduced by use of AdaptivePELE.

The first system used was the trypsin-benzamidine complex shown in Figure 1.11a, which we will refer to as TRP from here on, with an experimentally determined binding free energy of -6.2 kcal/mol (Mares-Guia and Shaw, 1965). The process of the free energy estimation consists of three steps: i) AdaptivePELE simulation, ii) MSM generation, and iii) PMF and binding free energy estimation.

AdaptivePELE simulation. A simulation is started from an initial structure where benzamidine is fully solvated, with approximately 23 Å RMSD from the native structure (bound), as shown in Figure 2.2. The simulation consists of a series of short PELE simulations, followed by clustering using the leader algorithm on the ligand RMSD and spawning trajectories with a score inversely proportional to the cluster population. The

PELE simulations were run for a number of steps l , which we will refer here as the *epoch length*. Simulations used the OBC implicit solvent and ran ANM every step with modes changing every six. The ligand perturbation consisted of small steps of either 1 or 2 Å, each with a 50% probability, and rigid body rotation between -36° and 36° or between -18° and 18° , again, with 50% probability. The most important parameter for an AdaptivePELE run is the epoch length, that is, the length of the PELE simulation. Shorter epochs imply more interruptions and more opportunities to distribute the computing capabilities more efficiently. However, if the next step is to analyse the simulations with MSM, simulations need to be long enough to allow for a sufficiently large lag time for the model to satisfy the Markovian property.

MSM generation. The first step in the estimation of an MSM is the discretisation of the input simulations. To do so, one must first decide on a set of features to use in the clustering. For all trajectories we extracted the coordinates of the amidine carbon and ran K-Means clustering to obtain 100 clusters. We used the python PyEMMA package (Scherer et al., 2015) both for clustering and MSM construction. A count matrix was created using a sliding window approach, with a lag time $\frac{l}{2}$, in order to maximise both the collected statistics and the Markovianity of the model. From the count matrix, a transition matrix is estimated using a maximum likelihood estimator (Prinz et al., 2011). The transition matrix can be diagonalised, offering plenty of information about the system dynamics. In this work, however, we focus mainly on the stationary distribution, π , that allowed us to estimate the binding free energy.

PMF and binding free energy estimation. Following the method presented by Buch et. al. (Buch et al., 2011), the stationary distribution is used to estimate a PMF using equation 1.15 and later the binding free energy is estimated using 1.17, see section 1.7 for more details. What remains to be described in the mentioned section is the process for volume estimation. The volume of the states is calculated by considering the maximum and minimum values of the discretised coordinates. A cube is constructed with those points as vertices, and is later cut into smaller cubes. Each small cube is assigned to the nearest cluster centre and the total volume of each cluster is then the sum of the populated small cubes belonging to the particular state. The whole process of MSM generation and free energy estimation is repeated several times (usually 10) to approximate the estimation error. This is done using a bootstrapping technique, where at each iteration a number of trajectories is sampled with replacement from the complete set of simulations, and the standard deviation of the predictions is taken as an error estimate.

We started by testing the influence of the epoch length in the free energy estimation by running simulations with 50 processors. We ran AdaptivePELE with 50, 100, 200 and 400 steps per epoch and calculated the estimated binding free energy as more epochs are introduced. The lag time used for the MSM should not be larger than half the simulation length or else the statistics collected would be insufficient, and a non-negligible fraction of the simulation would be wasted due to the sliding window approach (Doerr and De Fabritiis, 2014). Thus, for example for simulations of 50 steps the lag time would be 25 and for the simulations of 400 steps the lag time would be 200. Similarly, the number of clusters used was 100, since using more clusters creates numerical problems in the initial epoch due the small amount of data. Figure 2.3 shows that after about ten iterations the




Figure 2.2: Initial position of the ligand used in the AdaptivePELE simulation. The native pose is also shown for comparison, coloured in dark blue.




Figure 2.3: Evolution of the binding free energy prediction as more epochs are used. The horizontal line marks the experimental value, while the different coloured lines represent simulations with different epoch lengths. Error bars correspond to the standard deviation obtained from the bootstrapping procedure.

50 steps simulation seems to have converged around the experimental value (marked by the horizontal black line). The same cannot be said about the rest of the simulations. With a similar amount of aggregated sampling, the runs with longer epochs do not come closer than 1 kcal/mol to the experimental value. The sampling redistribution done at the end of each epoch appears to help in the prediction accuracy. In those simulations with longer epochs and less redistribution attempts, the prediction accuracy improves minimally.

To further test the effect of the AdaptivePELE spawning in the free energy calculations, we ran additional simulations with 200 and 400 steps per epoch and 144 and 528 processors. Lag time and number of clusters were selected in the same way as the previous simulations. The results are shown in Figure 2.4a for the 200 steps and Figure 2.4b for the 400 steps. It is clear from the plots that when the number of processors was increased the accuracy was significantly improved. With 528 processors both the simulations with 200 and 400 steps were within 1 kcal/mol of the experimental value. Also, with 400 steps and 144 processors similar accuracy was achieved. Shorter epochs seemed the most efficient choice for the estimation of the binding affinity. However, this imposes an important limitation on the MSM constructed, as mentioned before. The short simulations of 50 steps might give a sufficiently good approximation for a relatively fast binding like the trypsin-benzamidine one, but it is not likely to give good results for more complicated systems.

To further improve the simulation efficiency, we modified the initial setup. Instead of




Figure 2.4: Evolution of the binding free energy prediction as more epochs are used, and with a different number of processors. The horizontal line marks the experimental value, while the different coloured lines represent simulations with different numbers of processors. Error bars correspond to the standard deviation obtained from the bootstrapping procedure. Panel (a) Evolution of the prediction with different number of processors with simulations of 200 steps. Panel (b) Evolution of the prediction with different number of processors with simulations of 400 steps.

starting the simulation from a single starting conformation, we traced the binding path from a previous simulation and used those conformations as starting points. We assumed that by beginning the simulations from those diverse structures, sampling transitions in and out the binding site would be more likely, thus improving the free energy estimation. The path used is shown in Figure 2.5, with ligand structures coloured from red to blue as benzamidine enters the binding site. To test this approach, we ran simulations with 50 processors and 100 steps per epoch, using 100 clusters and 25 steps lag time for the MSM. The evolution of the prediction for each epoch is shown in Figure 2.6. Comparing it to the red line in Figure 2.3, it is clear that with the same computational resources and same model conditions the prediction was better from the very first epoch. At all epochs the estimated value is within 1 kcal/mol from the experimental value, with small fluctuations around it.

Considering the apparent improvement brought by the additional initial exploration, we decided to adopt this strategy and extend the testing system to incorporate the URO receptor (see section 1.8.3 for more details) with 3 ligands: *BAM*, *655* and *UI1* (2D structures and experimental affinities for these ligands are shown in Table 1.2). The new strategy inserted an extra step to the schema explained before, the new sequence looked like: i) short AdaptivePELE exploration, ii) “production” AdaptivePELE simulation, iii) MSM generation, and iv) PMF and binding free energy estimation.

The first step consisted of an AdaptivePELE simulation with short epochs of only four steps, with 144 processors, clustering using the ligand RMSD and spawning from those




Figure 2.5: Illustration of a binding pathway for trypsin-benzamidine. Protein is shown in white cartoon, while ligand is coloured from red to blue, with red marking the initial poses in the pathway and blue the final ones.




Figure 2.6: Evolution of the binding free energy prediction for the simulations starting from the pathway as more epochs are used. The horizontal line marks the experimental value, while the red line represents the predicted values. Error bars correspond to the standard deviation obtained from the bootstrapping procedure.




Figure 2.7: Evolution of some simulation metric as the simulation of URO with the BAM ligand progresses. Lines are coloured according to the epoch they belong. Panel (a) Evolution of the ligand RMSD to the native pose. Panel (b) Evolution of the ligand SASA.

states that have been explored the least. The goal of this simulation was to quickly explore the binding mechanism to extract a diverse set of structures that will serve as seeds for the longer, “production” AdaptivePELE simulation. In Figure 2.7 we plot the evolution of the ligand RMSD to the native pose, in panel (a), and the ligand SASA, in panel (b), along an example of such short simulation for the binding of the BAM ligand to URO. From the plots, it is clear that within the first four epochs RMSD values below 2 Å were achieved, which translated to about 5 minutes of wall clock time. All the output of this short simulation was then reclustered to extract a reduced set of structures to begin the next phase. We selected 40 such structures, this number was chosen because it is high enough to include sufficient diversity of ligand conformations but low enough to allow us to start several processors from each structure, which we assumed would also help improve the sampling efficiency. Next, we ran an AdaptivePELE simulation of 200 steps, with 240 processors, 100 steps as lag time and 100 clusters for the MSM discretisation. The MSM and free energy estimation steps were identical to those presented in the trypsin simulations.

Results from the new approach are shown in Figure 2.8. In the plot, the different colours are mapped to the three ligands used. The red lines represent values for the BAM molecule, blue for 655 and violet for UI1. The horizontal lines mark experimental values, while the lines with error bars represent the predicted values. It is clear from the plot that the method is not able to distinguish between the three ligands, as all seem to converge towards the binding free energy of BAM, which is around 4 kcal/mol higher than 655 and 7 kcal/mol higher than UI1. There are several possible explanations for the poor performance observed. One is the length of the simulations. In the case of trypsin a lag time of 25 steps seemed sufficient to predict the free energy of binding, but the length required for the URO receptor might be much longer, especially for ligands that are strong binders




Figure 2.8: Evolution of the binding free energy prediction for the URO simulations with the three selected ligands. The horizontal lines mark the experimental values, red for BAM, blue for 655 and violet for UI1, while the lines with error bars represent the predicted values, with the same colour code. Error bars correspond to the standard deviation obtained from the bootstrapping procedure.

like 655 and UI1. Secondly, at the beginning of each epoch simulations are started from a distribution that is surely different to the equilibrium distribution, since the states are selected from those that are less visited. The precise impact of such perturbation in the dynamics of the system is difficult to measure but it seems reasonable to assume that it will interfere with the accuracy of the free energy estimation.

For the reasons outlined above, we decided to swap the production AdaptivePELE simulations with longer, regular PELE simulations. The initial adaptive simulation was kept to attempt to preserve the sampling advantage, but running longer PELE simulations would: i) remove whatever noise, if any, was introduced by the spawning step, and ii) allow the use of a longer lag time. PELE simulations were run with 240 processors for 2000 steps. Data was collected and processed in the exact same way as previously explained. The number of clusters for the MSM was kept at 100 but several lag times were used to evaluate its impact (25, 50, 100, 200 and 400 PELE steps). To evaluate the prediction accuracy we used three statistics: the first, the coefficient of determination, R^2 , corresponds to the proportion of the variance in the dependent variable explained by the independent variable. This measure has a value between 0 and 1, with a value of 1 meaning that the independent variable (in our case the free energy estimation) perfectly predicts the experimental values. The second statistic is the Spearman correlation. This value, also between 0 and 1, compares the rankings of two variables. It allows us to determine whether the predicted affinity is able to correctly rank the different ligands in a set




Figure 2.9: Prediction accuracy for different lag times. Measures include R^2 and Spearman correlation in panel (a) and RMSE in panel (b)

even if the error with respect to the experimental values are large. Finally, we calculated the Root-Mean-Square Error of the predicted values, calculated as

$$RMSE = \sqrt{\frac{\sum_i (\Delta G_{exp}^i - \Delta G_{calc}^i)^2}{N}} \quad (2.2)$$

where the index i runs over all the ligands in the set and N is the size of said set. Figure 2.9a shows the value of the coefficient of determination and the Spearman correlation for different MSM lag times. Figure 2.9b also plots the RMSE value for different MSM lag times. The measures were split into two figures to ease the visualisation of their magnitudes and avoid visual artifacts due to the different ranges. The values for R^2 and Spearman correlation suggest that the predicted binding free energy is able to separate and correctly rank the three ligands at all lag times. However, the values of the RMSE indicate a dramatic decrease in the error for a lag time of 400 steps. For reference, the values of the R^2 , Spearman correlation and RMSE for the AdaptivePELE simulations shown before are 0.67, 0.5 and 5.15 respectively. Analysis of the statistics should be cautiously done due to the small number of ligands used (only three) which will inevitably lead to noisy estimations. Nonetheless, a qualitatively different performance is observed between the predictions done with regular PELE simulations with a previous AdaptivePELE exploration and those done with two subsequent AdaptivePELE simulations.

Overall, we learned from these preliminary tests that the efficiency gain in exploration offered by AdaptivePELE can help reduce the computational time needed for binding free energy estimations with regular PELE simulations. However, from the simulations with URO we observed that using AdaptivePELE with the same clustering and spawning

methods used for quick exploration simulations and short epochs does not result in accurate binding affinity predictions. Combining these lessons we envision two possible routes for more efficient free energy estimation: i) combining a very short AdaptivePELE simulation with a longer regular PELE simulation, and ii) modify the clustering and spawning methods used in AdaptivePELE for those longer, “production” simulations, to rule out the possible noise-inducing effects introduced. The first option will be thoroughly explored in sections 2.2 to 2.4, while the latter option will be examined in section 2.5.

2.2 Protocol for the estimation of binding free energies

To prove the efficiency gained by running a short AdaptivePELE to extract seeds for a longer simulation we decided to expand the test systems to multiple protein systems. This requires two tasks: i) put together a diverse set of proteins with multiple ligands each, and ii) create a reasonably automatic protocol. The test benchmark consisted of the following receptors: URO, plasmin, ERK2 and PR (see section 1.8 for more details on those proteins and the ligands used). The selected proteins present different degrees of difficulty. Plasmin is a simple and small protein, but its ligands are grouped into a narrow range of affinities, making their prediction more difficult. URO showed promising results in previous tests, so more ligands were incorporated to fully test the prediction accuracy. ERK2 is a much more complex case, mainly due to its wide and flexible binding site. Finally, PR is the most difficult protein to sample in the whole benchmark due to its buried binding site only accessible from a narrow entrance canal.

The protocol requires previous knowledge about the binding site, as well as, at least a good approximation of the binding mode. One could start from a crystal structure if available, although a docking pose should also work. Exact information about the binding mode is not needed since, like most pathway-based methods, the binding mode should be recovered from the production simulations. The protocol was implemented in python, in collaboration with Daniel Soler from Nostrum Biodiscovery, and chains the several steps explained below automatically, as depicted in Figure 2.10.

Initialisation. The initial structure needs to be processed before starting any simulations. Mostly, protein preparation is automatically performed, tasks like adding possibly missing hydrogen atoms, assign the force field atom types or generate the atomic parameters for the ligand.

Equilibration. After processing the input structure, an equilibration step takes place. This consists of a short simulations (~ 500 PELE steps) with the goal of refining possible small inaccuracies that might be present in the input structure. This procedure uses very weak ligand perturbation parameters, ensuring that it will remain in the binding site. The equilibration trajectories are clustered into 10 clusters, and for each group a representative structure is selected to use in the next step. The whole procedure requires about 1-2 minutes of wall clock time.

Ligand unbinding simulation. The equilibrated structures are used as input to a short AdaptivePELE simulation. The goal of this run is to map the unbinding process in order




Figure 2.10: Schematic summary of the binding free energy protocol. The top left image depicts a summary of the ligand exit simulation, with structures coloured from red to blue as the simulation advances. The top right image depicts the centre 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 coloured according to their equilibrium probability




Figure 2.11: Evolution of the SASA along the several epochs of the unbinding simulation. Values of SASA larger than 0.8 typically mean almost completely solvated ligand structures. Colours represent the epoch number. Panel (a) shows the evolution for an AdaptivePELE simulation with the 1UP ligand for URO. Panel (b) shows the evolution of the STR ligand with PR.

to generate representative poses to use as seeds in the production simulation. We map the unbinding process instead of the binding since it requires less information about the receptor, making it easier to automatise, and we assume that the resulting structures will be representative of both processes. Epochs of only two steps are used, with an spawning criterion distribution 75% of the processors to structures with the largest SASA and the remaining cores to the least visited clusters. The simulation is run until at least four independent trajectories reach a SASA value larger than 0.95. In Figure 2.11 two examples of ligand unbinding simulation are shown. Figure 2.11a shows the unbinding of the 1UP ligand to URO. The exit condition was met after six epochs. Similarly, Figure 2.11b shows the simulation output for ligand STR and PR. Here, the occluded binding site makes unbinding more difficult and about twice as many steps were required to meet the exit condition.

Clustering and “box” construction. To extract “seeds” from the unbinding simulation, that is, structures that will serve as inputs for running additional simulations, we cluster the output of the initial run. The clustering is done with the K-Means implementation of PyEMMA, and results in 40 conformations that we assume will approximate both the binding and unbinding process well. This number is arbitrarily set, but, based on prior experience detailed in the previous section, results in a useful compromise of having enough different states to sample most of the relevant conformation space for binding while being able to run several trajectories from each starting point. Additionally, the clusters are used to construct the simulation “box”, that is, the volume available for ligand exploration. Multiple shapes can be used for this box, such as a single sphere around the binding site, multiple spheres along the unbinding path (as shown in Figure 2.12) or a




Figure 2.12: Initial setup for the production simulation. Each dot represents the centre of mass of the ligand in an initial structure. The multiple spheres depict the simulation box, that is, the volume accessible to the ligand.

cylinder covering the same path. The ligand is restricted inside this closed volume, but is free to explore the allowed region, without any external force when inside. The objective of setting this box is to reduce the computational time required to obtain converged free energy estimations.

Production simulation. From the 40 selected structures in the previous step, we start an ensemble of longer simulations, that we will refer to as production simulations. As we saw in the previous section, AdaptivePELE simulations did not produce consistent binding free energy predictions, and for this reason the production runs are done using regular unbiased simulations. These can either use PELE (see section 2.3) or MD (see section 2.4). The resulting trajectories from these simulations are processed, extracting the centre of mass of the ligand, to construct the MSM. The centre of mass is used because it is a measure that is generally applicable to any protein-ligand system and thus aids in the automation of the protocol.

MSM and binding free energy estimation. The process of MSM estimation follows the same schema previously considered. The processed trajectories are clustered using K-Means and each snapshot is assigned to the closest cluster. A count matrix is constructed with a lag time, τ , which leads to a transition matrix and a stationary distribution describing the




Figure 2.13: Example of visualisation analysis useful for evaluating the modelling accuracy. In panel (a) the clusters are coloured according to their PMF value (red means higher probability, blue lower probability). Panel (b) shows a one-dimensional projection of the PMF with respect to the distance to the initial structure. Each point corresponds to a cluster, however, it should be noted that the figures are only used for illustrative purposes, and no relationship should be inferred from the two panels.

equilibrium probability of each state. From the distribution a PMF and binding free energy are estimated as detailed in section 1.7.

The main, new contributions to this step of the protocol are the analysis and visualisation tools shown in Figure 2.13. They give two complementary views of the PMF estimation that highlight possible shortcomings that can affect the prediction accuracy. Figure 2.13a shows a representation of the MSM clusters centres with the URO receptor. The clusters are coloured according to their PMF value, with the lowest values shown in red and the highest in blue. This visualisation allows for a quick check if the assumptions made in the PMF derivation are held, mainly whether the minimum corresponds to the binding pose and the fully solvated structures shows maximal, close to constant, PMF values. Additionally, this kind of image can also point to possible sampling artifacts caused by the simulation box, as exemplified by the few reddish clusters on the top right corner.

Figure 2.13b shows a related plot, a one-dimensional projection of the estimated PMF over the euclidean distance to the initial conformation, which points to the active site bound mode. This offers a simplified version of the information conveyed by Figure 2.13a, but it is easier to incorporate into the automatic protocol and provides an informative quick analysis of the simulation.

Once the entire protocol is run for the chosen set of ligands, the prediction accuracy is measured, like in the previous URO example, using three statistical metrics: i) the Spearman correlation, ii) the coefficient of determination, R^2 , and iii) the RMSE. In the following sections we report the results obtained from applying the protocol to some

test systems, in section 2.3 we use the version of the protocol that uses PELE as the propagator, we call this version of the protocol *PELE-MSM*. In section 2.4 we demonstrate the results of applying a second version of the protocol that uses MD as the propagator, which is why we refer to this variation as *MD-MSM*. Finally, in section 2.5, we introduce another version that again uses AdaptivePELE to run the production simulations but with different spawning and clustering methods than those shown in section 2.1.

2.3 PELE-MSM

The method termed PELE-MSM employs the protocol described in the previous section using PELE as the simulation propagator for the production runs. To thoroughly test its accuracy we ran three independent iterations of the whole protocol using 128 cores, with simulations of 1000 PELE steps. The three iterations were completely independent, starting from the same initial structure but with different random seeds. The multiple rounds were incrementally used in the free energy prediction. That is, we first estimated the binding affinity with the trajectories from the first iteration only. Then, we combined the simulations of the first two rounds and again estimated the binding free energy. Lastly, all trajectories were used to perform a final estimation. This process allowed us to investigate the effect of the amount of sampling in the prediction accuracy as well as providing a quick evaluation of the convergence of the calculations.

The benchmark used to test the method consisted of a total of four proteins and 23 ligands: plasmin with ligands L01, L02, L05, L06, L07 and L15 (see Table 1.1 for ligand structures and experimental affinities); URO with ligands 1UP, 2UP, 4UP, 6UP, 7UP, 675 and 39L (see Table 1.2 for ligand structures and experimental affinities); ERK2 with ligands EK2, EK3, EK6, EK9 and E63 (see Table 1.4 for ligand structures and experimental affinities); and PR with ligands STR, AS4, HCY, 356 and 400 (see Table 1.3 for ligand structures and experimental affinities). The work described in this section was recently published in the *Journal of Chemical Theory and Computation* (Gilabert et al., 2019).

2.3.1 Binding Free Energy for the Plasmin System

The initial structure used for the Plasmin runs was the top Glide SP docking pose, which closely matched the binding pose shown in the crystallographic structure (PDB id 4CIK). Due to the small size of the protein, this was the fastest system to test, requiring only 3 hours of wall-clock time per iteration. Figure 2.14 shows the correlation between the computational estimations of the binding free energy and the experimental values. Prediction has a R^2 values of 0.72, a Spearman correlation of 0.75 and RMSE of 1.11 kcal/mol. The graph clearly shows a separation of the ligands into two halves, one with the three strongest binders and one with the three weakest ones. This observation is supported by the statistical measures that show low error, with strong correlation and ranking. This positive result is obtained despite the fact that all six ligands are quite weak binders (best in the set has a binding free energy of -3.94 kcal/mol); as a reference, Glide SP and XP scores, although predicting the bound pose, give correlations of 0.02 and 0.04,




Figure 2.14: Correlation of the predicted binding free energy with respect to experimental values for the Plasmin system. Text labels identify the ligand that each point represents.

respectively.

To evaluate the accuracy as the amount of sampling increases, we calculated the coefficient of determination as more iterations are completed, which resulted in R^2 values of 0.55, 0.70 and 0.72 for 1, 2 and 3 iterations respectively. While the accuracy with only one iteration was already notable, it improved when more trajectories were added. This suggests that with only 127 simulations, one could get a quick approximation, but at least twice as many simulations are required to obtain accurate results. In Figure 2.15 the PMF obtained at the introduction of each iteration are shown, with Figure 2.15a showing the first round, Figure 2.15b the PMF for the first two, and Figure 2.15c shows the estimation for the combination of the three. There exist fewer differences in the PMF profiles than there are in the R^2 values. All three cases show a similar outline, with a well defined minimum at close distances to the native bound pose, and overall flat PMF values at distances larger than 10 Å.

The PMF profiles shown in Figure 2.15 present little fluctuations, with only a small group of about four states deviating from the main trend. We investigated the representative structures of said cluster, which is shown in Figure 2.16b. In Figure 2.16a the structure of the absolute minimum is shown, with the ligand and some important residues painted cyan and the protein illustrated using white ribbons. The ligand, which has a zwitterionic nature, is trapped between two positive and two negative ones. The positive charges come from two arginine residues, Arg 33 and 71, while the negative charges come from two aspartates, Asp 55 and 57. The small PMF deviation, which we will refer here as a secondary minimum, although it is of little importance, shows a similar but less stable binding mode to that of the active site. Both structures share a pair of arginine residues that, along with Arg 35, interact with the negative charge of the ligand. On the other end of the L01 molecule, its positive charge interacts with only one negative charge, Glu




Figure 2.15: Comparison of the evolution of the one-dimensional PMF with three iterations for the L01 ligand of the plasmin receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.

69, forming a similar but less balanced interaction than that of the binding pose.

For completeness, in Appendix A, we show the evolution of the PMF estimation for the three iterations for all the remaining ligands of the set. This is shown in Figure A.1 to A.5. It seems clear from the graphs that a similar analysis to that of L01 applies to the rest. All profiles are similar throughout the three iterations, with well-defined minima at small distances to the active conformations and approximately constant values at large distances. All in all, the PELE-MSM scheme showed accurate predictions of the binding free energies for the plasmin receptor with around 9 hours of wall-clock time and 128 processors. Importantly, these positive results are achieved despite the reduced range of binding affinities of the ligand set used, for a system that despite its simplicity has been unsuccessfully studied using FEP and MM/GBSA with R^2 values of 0.16 and 0.2 respectively (Schmidt et al., 2017). The good performance of the method with this small test system, in line with the previous results obtained with trypsin suggest that the method is well suited to predict the binding free energy of fragments.

2.3.2 Binding Free Energy for the URO System

The study of the URO receptor with PELE-MSM started from a Glide docking pose, which could reproduce the crystallographic binding pose (using PDB entry 4FU9 as reference). This protein, which is larger than plasmin, required around 8 hours of wall-clock time per iteration. The correlation between the calculated and experimental values, shown in Figure 2.17, was qualitatively similar to that observed in plasmin. The ligands were separated into two groups, the two with highest affinity (39L and 7UP) on one side and the remaining five ligands on the other. The difference between the two groups is less marked than the one observed in the plasmin predictions, and the ranking of the ligands inside each group is also less precise. In quantitative terms, the prediction statistics show




Figure 2.16: Representative structure of the minima of the L01 ligand binding to the plasmin receptor. Protein is shown in white ribbons, ligand and key residues are shown in licorice representation with cyan carbons. Panel (a): Active site and absolute minimum. Panel (b): Secondary minimum at about 10 Å distance from the active site, with the active site structure depicted in blue for reference.

an R^2 value of 0.59, Spearman correlation of 0.61 and RMSE of 5.86 kcal/mol. The free energy predictions were significantly off from experimental values, however the other measures showed some predictive power of the ligand ranking within the set.

The PMF profiles obtained for URO show signs of not fully converged simulations. This is particularly true for the 7UP ligand; in Figure 2.18 its PMF profiles are shown. With only one iteration significant deviations from an ideal profile are observed, with the PMF minimum located at around 15 Å. A representative structure of such states is shown in Figure 2.19, with the bound pose shown in dark blue for reference. In this structure the ligand forms an ionic bond with Asp 147 as well as hydrophobic interactions. The protein-ligand interaction does not seem particularly strong, as evidenced by the profile including the second and third iterations (Figures 2.18b and 2.18c) where the mentioned minimum vanishes and the PMF presents a shape more similar to that seen in the plasmin receptor. The rest of the ligands do not show such inaccuracies in their PMF profiles (Figure A.6 to A.11). They do, however, show overall less well-defined PMF outlines, which together with the moderate success reflected by the measured suggest that more sampling would be necessary to obtain better estimations, likely due to the bigger size of the protein and stronger interactions formed with the larger ligands, most of which have net charges, some even with two positive charges.




Figure 2.17: Correlation of the predicted binding free energy with respect to experimental values for the URO system. Text labels identify the ligand that each point represents.




Figure 2.18: Comparison of the evolution of the one-dimensional PMF with three iterations for the 7UP ligand of the URO receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure 2.19: Representative structure of the minimum in the first iteration of the 7UP ligand binding to the URO receptor. Protein is shown in white ribbons, ligand and key residues are shown in licorice representation with cyan carbons, with the active site structure depicted in blue for reference.




Figure 2.20: Correlation of the predicted binding free energy with respect to experimental values for the ERK2 system. Text labels identify the ligand that each point represents.

2.3.3 Binding Free Energy for the ERK2 System

The initial structures were extracted from a Glide docking. In order to reproduce the crystallographic pose (PDB entry 4FV3) we needed to use the hydrogen bond constraint options (to Met 106), otherwise, Glide produced several incorrect poses without the constraint option. For this system, we again performed three independent simulations, each taking approximately 8 h of wall-clock time.

Contrary to the examples shown before, PELE-MSM has no predictive power on the ERK2 set. In Figure 2.20, it is apparent that the method is not able to separate or rank the ligands. The prediction provides an R^2 value of 0.01, Spearman correlation of 0 and RMSE of 6.98 kcal/mol.

Figures 2.21, and A.12 to A.15, provide evidence to why the accuracy of the method is so low in this example. All those PMF profiles have a “V shape”, with a minimum around 10 Å from the initial pose, contrary to the well-behaved profiles of the plasmin or URO. The failure to identify the correct PMF minimum arises from modelling error. We examined the force field total energy internally used in the PELE simulation with respect to the distance, shown in Figure 2.22 for the EK2 ligand. There, it is clear that the interactions in the native pose are not properly captured by the energy function, having an energy 50 kcal/mol higher than those poses in the PMF absolute minimum.

The structures of the active site and the PMF minimum are shown in Figures 2.23a and 2.23b, respectively. The native pose, despite forming more protein-ligand interactions than those in the PMF minimum, had a less favourable energy. This could be explained by two factors, a poor force field parametrisation or the inability of the implicit solvent to capture important water network interactions in the bound pose.




Figure 2.21: Comparison of the evolution of the one-dimensional PMF with three iterations for the EK2 ligand of the ERK2 receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure 2.22: Plot of PELE's internal energy values with respect to the centre of mass distance, that is, the same distances used in the one-dimensional PMF plots




Figure 2.23: Representative structure of the minima of the EK2 ligand binding to the ERK2 receptor. Protein is shown in white ribbons, ligand and key residues are shown in licorice representation with cyan carbons. Panel (a): Active site structure. Panel (b): PMF minimum at about 10 Å distance from the active site, with the active site structure depicted in blue for reference.

To determine the influence of the implicit solvent in the energy evaluation, we ran a PELE simulation of the same length and number of processors but introducing four explicit waters, positioned as shown in Figure 2.24a. From the image, it is already apparent that the waters form bridging interactions between the ligand and the protein. This simulation was run using the newly developed water sampling algorithm part of Martí Municoy’s PhD thesis (manuscript in preparation). In such a simulation, waters are sampled using a parallel MC scheme at each PELE step and are later subjected to energy minimisation. Figure 2.24b shows the total energy obtained from the simulation, which, if compared to Figure 2.22, shows a remarkable improvement in energy calculation accuracy, while showing limited sampling at low distances. While the regular PELE-MSM is not able to predict binding free energies for a kinase such as ERK2, with the addition of the Water-PELE mode, once it is fully developed, the prediction accuracy might prove sufficient to make it a relevant method for the study of kinases in the context of drug discovery.

2.3.4 Binding Free Energy for the PR System

Finally, we applied the PELE-MSM protocol to the PR system. Initial structures were generated from Glide docking, which could reproduce the crystallographic pose of progesterone. We then placed the other ligands similarly in the binding site (we used the 1A28 crystal as the protein structure for all ligands). Each iteration of the protocol required approximately 8 hours of wall clock time. Similarly to the results obtained for the ERK2 receptor, the correlation plot of the binding free energy estimation versus the experimental value (shown in Figure 2.25) showed no correlation, with a R^2 value of 0.0, Spearman




Figure 2.24: Preliminary analysis of the impact of explicit waters in the PELE simulations of the ERK2 receptor with ligand EK2. Panel (a): Initial structure with water positions and interactions marked. Panel (b): Plot of PELE's internal energy values with respect to the centre of mass distance in the simulation with four explicit waters.

of 0.0, RMSE of 4.41 kcal/mol.

Contrary to the ERK2 case, the problem with the PR simulation does not appear to be with the force field or model assumptions, but with insufficient sampling. In Figure 2.26, the PMF profile obtained for the STR ligand is shown. The PMF minimum is indeed located at very low distances. There are, however, very few states located at less than 10 Å, and the PMF does not flatten as the distance increases. Instead, we observed two seemingly disconnected sets of states, with most of them clustered between 15 and 20 Å and similar PMF values. The maximum value of the profile is not at these large distances, but at what seems to be a sort of transition state.

The sampling difficulties are better illustrated by plotting the SASA of the ligand as the production advances, as shown in Figure 2.27. In the figure, a total of 381 trajectories (the combined results of the three independent iterations run are plotted, of which only ~10 show transitions involving states of low SASA values, within the active site and entrance channel. Moreover, these transitions occurred mostly in early simulation times. Thus, the simulation forms two almost fully disconnected sets of states, as observed in the PMF plot, which leads to a spurious free energy estimation. This is likely caused by the fully occluded nature of the PR binding site which prevents sufficient observation of both binding and unbinding events for the MSM to produce a correct estimate of the equilibrium probabilities.




Figure 2.25: Correlation of the predicted binding free energy with respect to experimental values for the PR system. Text labels identify the ligand that each point represents.




Figure 2.26: One-dimensional PMF of the first iteration for the STR ligand of the PR receptor for the three iterations.




Figure 2.27: SASA evolution during the production simulation for three iterations of 127 trajectories each for the PR system with the STR ligand.

2.3.5 Effect of the previous AdaptivePELE simulation


After testing the four receptors, we evaluated the impact that the initial steps of the protocol had in the free energy estimation. We ran a simulation matching the resources requirements of one protocol iteration with the L05 ligand and plasmin protein and the 39L ligand of the URO system. We ran a PELE simulation with 128 processors starting from the initial structure of the previous protocol run, that is, the docking pose matching the crystallographic structure.

Figure 2.28 compares the SASA evolution of the 250 initial steps with (Figure 2.28a) and without (Figure 2.28b) the previous exploratory AdaptivePELE simulation. It is clear from the plots that adding an extra enhanced sampling simulation at the beginning of the protocol introduces plenty of additional transitions between regions of low and high SASA. The simulations starting from the bound pose did not show any unbinding event in the first 250 steps, and needed around 1000 PELE steps to produce a conformation with a SASA value larger than 0.9, as shown in Figure 2.29a. On the other hand, the simulations starting from multiple seeds showed a uniform exploration of the whole range of SASA values, with multiple transitions between the different regions. If we compare the PMF profile for the simulation starting from the bound (Figure 2.29b) with the one obtained in first iteration of the protocol run starting from multiple structures (Figure A.2a), the profile is more noisy, with more dispersion of the clusters, probably as a consequence of the small range of states explored in the simulation.

A similar analysis was done for the 39L ligand binding to the URO receptor. In Figure 2.30a, we show the evolution of the ligand SASA for the production simulation. Similarly to the L05 case, only one trajectory showed an unbinding event, with SASA larger than 0.8. Moreover, the intermediate regions with SASA values between 0.3 and 0.6 were only populated at the very beginning of the simulation, but were barely sampled overall. In Figure 2.30b, the PMF estimated from the same simulations is shown. There, we again see a profile that strongly deviates from the assumptions made about the dynamics. The PMF presents a bifurcating shape with the maximum at distances around 5 Å and a roughly constant region at distances larger than 10 Å.

While it might be expected to not see many unbinding events for the 39L ligand, which has a free energy of binding of approximately -9 kcal/mol, it is more surprising that a similar phenomenon is observed in the weakly binding L05 ligand. Being able to sample several occurrences of binding and unbinding events is crucial for accurate binding free energy estimation, thus the introduction of a previous enhanced sampling exploration led to improved sampling.

In summary, PELE-MSM has shown success in ranking a set of several ligands for two proteins, plasmin and URO. In the former, the absolute free energy prediction was on par with state of the art methods, with RMSE values less than 2 kcal/mol. The accuracy for the URO receptor was more modest, but still could be useful in separating strongly binding ligands for weaker ones. For the other two receptors, it has highlighted two well-known issues in biomolecular simulations: limited accuracy of the models and insufficient sampling. An alternative to the modelling assumptions would be the use of MD simulations (as we will see in the next section), and the sampling limitations could be tackled



(a) SASA evolution from multiple structures (b) SASA evolution from single structures

Figure 2.28: Comparison of the exploration with and without a previous AdaptivePELE simulation for the L05 ligand for the plasmin system. Panel (a): Evolution of the SASA during the 250 PELE steps of the first iteration of the simulation starting from seeds extracted from and AdaptivePELE simulation. Panel (b): Evolution of the SASA during the 250 PELE steps of the first iteration of the simulation from a single bound conformation in the active site. Panel (c): Evolution of the SASA during the entire production simulation of the first iteration starting from a single bound conformation in the active site.




Figure 2.29: Comparison of the exploration without a previous AdaptivePELE simulation for the L05 ligand for the plasmin system. Panel (a): Evolution of the SASA during the first iteration of the simulation starting from a single bound conformation in the active site. Panel (b): PMF of the first iteration of the simulation starting from a single conformation in the active site.




Figure 2.30: Comparison of the exploration without a previous AdaptivePELE simulation for the 39L ligand for the URO system. Panel (a): Evolution of the SASA during the first iteration of the simulation starting from a single bound conformation in the active site. Panel (b): PMF of the first iteration of the simulation starting from a single conformation in the active site.

using enhanced sampling methods. One thing to note is the high RMSE values of all systems but plasmin. This considerable deviation from the experimental values could be in part explained by the limited solvent exploration caused by the simulation box.

2.4 MD-MSM

A second version of the protocol for the calculation of binding free energies was developed using MD simulations for the production. This schema, named MD-MSM, was an exact replica of the PELE-MSM protocol swapping the PELE simulations with MD ones. The MD simulations were run using an extension to AdaptivePELE which used the python OpenMM API (Eastman et al., 2017) to carry out the simulations. The newly developed MD mode started processing the input structure (or multiple structures in our case) to fix common problems such as identifying disulphide bonds. The parameters are generated using Amber force fields, by default Amber14 (Maier et al., 2015) for proteins, and the ligands are parameterised with GAFF (Wang et al., 2004) using the Ambertools suite (Wang et al., 2006b). Using the TLeap software from Ambertools, the system is solvated using a cubic box and ions are added to neutralise.

Input structures are minimised for 2000 steps of the L-BFGS optimisation algorithm, followed by 400 ps of NVT equilibration at 300 K with restraints to the protein and ligand heavy atoms of $5 \text{ kcal mol}^{-1} \text{ \AA}^{-2}$. The system is subjected to a second equilibration run under NPT conditions at 300 K and 1 atm with restraints to protein alpha carbons and ligand heavy atoms of $0.5 \text{ kcal mol}^{-1} \text{ \AA}^{-2}$. The time step used in the equilibration simulation, as well as the production ones, is 2 fs. Finally, the production simulations are run with no restraints but restricting the volume accessible to the ligand, as was done in the PELE-MSM protocol, and conformations were saved every 125 ps.

The benchmark used to test the method consisted of a total of three proteins and 25 ligands: plasmin with ligands L01 to L16 (see Table 1.1 for ligand structures and experimental affinities); URO with ligands 1UP, 2UP, 7UP and 39L (see Table 1.2 for ligand structures and experimental affinities), and ERK2 with ligands EK2, EK3, EK6, EK9 and E63 (see Table 1.4 for ligand structures and experimental affinities). The work described in this section was developed in collaboration with Oriol Gracia during his master thesis and was submitted for publication (under revision at the moment of writing this thesis) in the *Journal of Chemical Information and Modelling* (Gilabert et al., 2020).

2.4.1 Binding Free Energy for the Plasmin System

As was done in the PELE-MSM study, the first system used to test the new version of the protocol was the plasmin receptor. We did so by using a series of 16 ligands. For each molecule we performed 80 simulations of 200 ns, starting from 40 structures selected from the initial stages of the protocol as described in the PELE-MSM section. The simulations amount to a total aggregated time of $256 \mu\text{s}$ ($16 \mu\text{s}$ per ligand), and were processed in the same way as the PELE runs, building an MSM with a lag time of 400




Figure 2.31: Correlation of the predicted binding free energy with respect to experimental values for the plasmin system. Text labels identify the ligand that each point represents.

saved steps, corresponding to 50 ns, and 1000 clusters. The number of clusters used with MD simulations was notably larger than the ones used with PELE. This was due to the higher amount of conformations to analyse in MD simulations. Early tests varying the number of clusters showed that a number of clusters similar to the ones used with PELE yielded unreliable models, while increasing this number improved significantly the MSM quality.

The correlation of the predicted free energy values with the experimental values is shown in Figure 2.31. The plot shows a positive correlation, with a low RMSE of 0.64 kcal/mol, R^2 of 0.343 and Spearman correlation of 0.66. The general agreement between the predicted and experimental values is reflected in the Spearman correlation, even if the calculated values are generally overestimated. The coefficient of determination is low due to two groups of ligands that deviate from the general trend. On the one hand, ligands L02, L05, L06 and L12 are correctly ordered but their estimated affinities appear shifted down with respect to the general trend drawn by the other ligands. On the other hand, the binding free energy of ligand L03 is overestimated and lies well outside the general trend. There is no apparent reason that could explain the misclassification of this ligand, but the shift observed in the group of four ligands suggest that our method has a lower bound in the range of affinities that is able to predict at -3 kcal/mol. Overall, despite the modest values obtained for the prediction statistics, the method seems to perform quite well, ranking correctly most of the ligands in a set that is difficult to rank due to the small ranges of affinities covered.




Figure 2.32: Analysis of the evolution of the prediction statistics (Spearman 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 starting point is based on adding the first 75 ns segments of all trajectories (with subsequent increments of 25ns).

In Figure 2.32, we plot the evolution of the coefficient of determination and the Spearman correlation as different simulation lengths are considered. The x axis starts at a length of 75 ns due to the MSM lag time of 50 ns, and increases by 12.5 ns. For example, the dotted line marks the performance with simulation of 100 ns (which amounts to an aggregated simulation time of 8 μs given that there are 80 trajectories and 16 ligands). The figure shows a slight improvement when the simulation length increases until 100 ns, beyond that, the metrics fluctuate around the values reported in the previous paragraph.

Additionally, we ran some additional replicas to assess the convergence of the method results. First, we ran two extra ensembles of 80 simulations for ligand L04 varying the random seed and then we repeated the whole protocol, generating a new set of initial structures and running simulations from them, for the L06 ligand. The results from these additional runs are shown in Table 2.1. The agreement for both tests was excellent. The predictions with different seeds for ligand L04 were within the standard deviation, and the difference in the binding free energy estimation of the two independent protocol runs with ligand L06 was only of 0.1 kcal/mol. This, together with the analysis of Figure 2.32 suggests that the amount of sampling used in the simulation is sufficient to obtain reliable estimations.

Table 2.1: Comparison of estimated binding free energies obtained from runs starting from different seeds or different initial structures for ligands L04 and L06 and plasmin receptor.

Ligand	Seed	Initial structures	Estimated ΔG (kcal/mol)
L04	75686	Set 1	-3.43 ± 0.27
L04	42353	Set 1	-3.21 ± 0.17
L04	12351	Set 1	-3.70 ± 0.19
L06	75686	Set 1	-3.66 ± 0.19
L06	75686	Set 2	-3.77 ± 0.22

In section B.1 of the Appendix, the PMF profiles for all 16 ligands are shown. The presented landscapes strongly agree with the assumptions made in the free energy calculation. All of them show a well-defined minimum, with roughly constant values at distances larger than 10 Å. Some of the ligands show signs of noise and non-convergence, like ligand L13 (Figure B.4a) or ligand L16 (Figure B.4d), but overall, as seen with the PELE-MSM examples, obtaining clean PMF profiles from the simulations goes hand in hand with accurate binding free energy predictions.

2.4.2 Binding Free Energy for the URO System

As a second test, we executed the MD-MSM protocol with the URO receptor and four ligands. Since the protein is bigger than plasmin, we hypothesised that more parallel trajectories would lead to a better estimation and ran 120 simulations for each molecule of length 100 ns, starting from 40 structures selected from the initial stages of the protocol as described in the PELE-MSM section. The simulations amount to a total aggregated time of 48 μs (12 μs per ligand), and were processed in the same way as the PELE runs, building an MSM with a lag time of 400 saved steps, corresponding to 50 ns and 1000 clusters.

In Figure 2.33 we show the correlation between the predicted free energy values and the experimental ones. The plot shows a clear disagreement between the two, with an obvious anti-correlation and the ligand split in two pairs, 7UP and 39L which have the highest affinity in the set were predicted as weaker binders than 1UP and 2UP, contrary to the experimentally determined binding affinities. This was indicated by the negative value of the Spearman correlation, of -0.6. The coefficient of determination showed a misleadingly high value of 0.454. The RMSE obtained was 4.31 kcal/mol. While both the correlation plot and the Spearman correlation illustrate without a doubt that the protocol is not able to predict binding free energies for this receptor, the value of the statistics should be carefully considered due to the small number of data points. The reduced sample size could lead to high fluctuations that would difficult the interpretations, although this is not the case here, as they correctly identify the poor performance of the method.

To further investigate the causes of the failure, we plotted the PMF profiles of all four ligands in Figure 2.34. The profiles for ligands 1UP and 2UP (Figure 2.34a) show clean profiles, with a well defined minimum at short distances from the initial structure. The




Figure 2.33: Correlation of the predicted binding free energy with respect to experimental values for the URO system. Text labels identify the ligand that each point represents.




Figure 2.34: One dimensional projection of the PMF over the distance to the initial structure of the simulations with the URO system. Panel (a): PMF for ligands 1UP and 2UP (weaker binders). Panel (b): PMF for ligands 39L and 7UP (stronger binders).

same is not true for ligands 39L and, especially, 7UP. As shown in Figure 2.34b, the PMF profile for 39L has a well-defined minimum at short distances, but has several secondary minima at different distances. The absolute minimum for 7UP, however, is located at around 10 Å, with secondary minima at distances around 0 and 25 Å.

In Figure 2.35, representative structures for the minima of ligands 7UP and 39L are shown. The representative structures of the absolute minimum of the 7UP PMF are illustrated in Figure 2.35a. They correspond to a group of states located at around 10 Å of the binding site. The clusters in that region can be separated into two groups that feature very similar poses, with roughly the same centre of mass but flipped (shown in cyan and purple). Due to the almost symmetrical structures of the ligand the structures depict similar interactions with the protein, a hydrogen bond with Glu 151 and $\pi - \pi$ interactions between the fused rings of the ligand and Tyr149 and Tyr28 (shown coloured in green). The secondary minimum observed in the 39L PMF at around 11 Å from the binding site forms a sort of “pre-binding” conformation, with a very similar orientation to the bound pose but out of the binding site. There ligand forms several interactions with the protein: hydrogen bonds with Ser197, Val39 and Asp55, a salt-bridge interaction with Asp55 and a cation- π interaction with Tyr50.

The analysis of the relevant structures extracted from the PMF profiles suggests the poor predictive performance could be caused by metastable trapping of the largest ligands in the protein surface. There, those molecules can form several interactions with the protein due to their large size and electrical charge which would translate into longer transition times and thus require longer simulation times to observe binding or unbinding events. Therefore, we decided to run another set of longer simulations with the same set of ligands to test whether the effects observed would vanish with more sampling. We ran 40




Figure 2.35: Representative structures for relevant PMF minima for the 7UP and 39L ligands of the URO receptor. The representative structures are shown with carbons colored in cyan. Interacting amino acids and the native bound pose are shown with carbons coloured in green. Non-carbon atoms are coloured according to the atom type (blue for nitrogen and red for oxygen). Panel (a): Representative structures for the absolute PMF minimum of the 7UP ligand. The mentioned structures are shown with carbons coloured in cyan and purple. Panel (b): Representative structure of the secondary minimum of the PMF profile of the 39L ligand.

trajectories of 1 μs each, starting from the same 40 initial structures used in the previous runs. The new set of simulations amounted to a total aggregated simulation time of 160 μs (40 μs per ligand). The MSM parameters were kept the same as in the previous runs.

Figure 2.36 shows the correlation between the predicted and experimental values obtained from the new set of simulations. The anti-correlation observed with the simulations of 100 ns disappeared with more sampling. However, the new set of simulations did not improve dramatically the performance, predicting as equal the three ligands 1UP, 2UP and 7UP, which, according to their experimental values, are separated by at least 1 kcal/mol. The statistics obtained are a R^2 value of 0.477, Spearman correlation of 0.8 and RMSE of 4.94 kcal/mol. The error has increased, but the other two metrics show an important improvement, even if misleading after the analysis of the correlation plot.

The PMF profiles, shown in Figure 2.37, are different from those obtained for the shorter simulations. For the ligands 39L and 7UP (Figure 2.37b), the profiles are clean, with the absolute minimum at a distance close to zero, and some secondary minima of much less importance than those shown in Figure 2.34. The same is not true for the 1UP and 2UP ligands (Figure 2.37a). There, we see a bimodal PMF distribution with two equal minima, one at 0 Å, and one at around 20 Å. The minimum at the largest distance is more sparsely populated than the one at short distance, suggesting that it might be caused by one trajectory sampling some infrequent but strong interaction of the ligands with the protein surface.




Figure 2.36: Correlation of the predicted binding free energy with respect to experimental values for the URO system with the longer simulations. Text labels identify the ligand that each point represents.




Figure 2.37: One dimensional projection of the PMF over the distance to the initial structure of the longer simulations with the URO system. Panel (a): PMF for ligands 1UP and 2UP (weaker binders). Panel (b): PMF for ligands 39L and 7UP (stronger binders).




Figure 2.38: Depiction of the new cylindrical region available for the ligand. The cylinder is delimited by the small dots, the protein is shown in white cartoon and the ligands are shown in cyan, the structures shown correspond to the secondary minima found in the short simulations.

The analysis of the PMF profiles, in agreement with the correlation plot, indicates that the use of long simulations can improve some of the problems identified with the first test. Nonetheless, the accuracy obtained is far from sufficient to provide a reliable estimation for the reduced set of ligands picked. As a final attempt to improve the prediction accuracy, we repeated the long simulations, but modifying the volume available for the ligand.

In the new set of trajectories, the ligand was allowed to explore a region defined by a cylinder, as shown in Figure 2.38. This new encompassing volume allows the ligand to freely explore the binding site, as well a fully solvated state. Through the use of a flat bottom potential restraint, once the atom closest to the centre of mass of the ligand exits the defined cylinder, an attractive force is applied, restricting the possibility of interaction with other sites in the protein surface. The purpose of this restriction, similar in nature to the one applied in funnel metadynamics, is to, in addition of avoiding metastable states, to reduce the number of fully solvated states accessible for the ligand, which is assumed to improve the convergence of the calculations and require less computational power.

The correlation of the estimated values for the new simulation ensemble is shown in Figure 2.39. The new results showed considerable improvements with respect to previous ones. The four ligands are separated into two groups, and the two strongest binders 39L and 7UP were grouped together and correctly predicted as the two ligands with highest affinity, although the order between them was reversed. The other two molecules, 1UP and 2UP were predicted to be almost identical, despite having around 1 kcal/mol experimental difference. However, they were predicted as clearly weaker binders than the other two, something that was not achieved in previous attempts. The measured statistics were an R^2 value of 0.761, Spearman correlation of 0.6 and RMSE of 5.64 kcal/mol, which means a significant improvement over the previous attempts, despite the increase in the error of the absolute value.




Figure 2.39: Correlation of the predicted binding free energy with respect to experimental values for the URO system with the longer simulations and a cylinder shaped simulation “box”. Text labels identify the ligand that each point represents.

The PMF profiles for these new simulations are shown in Figure 2.40. In all four cases the issues observed in the previous simulation were not present. All profiles show a clear trend, with a minimum at close distances and PMF values that are roughly constant at distances larger than 15 Å. The PMF profiles indicate that in these new simulations there were no significant protein-ligand interactions outside of the binding, as evidenced by the improved statistics, which has contributed to a more accurate estimation.




Figure 2.40: One dimensional projection of the PMF over the distance to the initial structure of the longer simulations with a cylinder shaped simulation “box” with the URO system. Panel (a): PMF for ligands 1UP and 2UP (weaker binders). Panel (b): PMF for ligands 39L and 7UP (stronger binders).

2.4.3 Binding Free Energy for the ERK2 System

As a third and final test, we executed the MD-MSM protocol with the ERK2 receptor and five ligands. For this system we used the protocol that worked best with URO, 40 trajectories of 1 μs each, starting from 40 structures selected from the initial stages of the protocol as described in the PELE-MSM section. The simulations amount to a total aggregated time of 200 μs (40 μs per ligand), and were processed in the same way as the PELE runs, building an MSM with a lag time of 400 saved steps, corresponding to 50 ns and 1000 clusters. As we did with the URO protein, we used two different simulation boxes, spherical and cylindrical.

We first ran the simulations with a spherical box, and analysed them similarly to the previous cases. First, we looked at the correlation with respect to the experimental measurements, shown in Figure 2.41. The plot shows good correlations, with almost perfect ranking, except for the EK9 ligands which is underestimated. The deviation of the EK9 ligands has an important impact in the R^2 value of 0.526. On the other hand, the Spearman correlation is high, with a value of 0.9, and the RMSE is quite high too, at 4.77 kcal/mol.




Figure 2.41: Correlation of the predicted binding free energy with respect to experimental values for the ERK2 system with the sphere shaped simulation box. Text labels identify the ligand that each point represents.

The correlation plot and the statistics indicate an accurate estimation of the binding free energy, contrary to the simulations done with PELE. However, the PMF profiles present

strong similarities to the PELE case presented in section 2.3.3, where the simulations did not properly capture the active conformation as the absolute minimum. The new profiles, shown in Figure 2.42, show that only for ligand EK9 the absolute minimum was correctly identified. For the four remaining ligands, the minima was placed at different distances. For ligand E63, the minimum is located at around 5 Å from the active pose, deeper in the binding site. For the other three ligands, the minima are located in the same region, binding inside a small cavity in the protein surface at between 15 and 20 Å from the binding site. The mentioned structures are shown in Figure C.1 and C.2 in section C.1 of the Appendix.

The simulations with a cylinder-shaped simulation box for ERK2 produced, contrary to what we saw with URO, worse binding free energy prediction. The measured statistics were: an R^2 of 0.143, Spearman correlation of 0.5 and RMSE of 6.44 kcal/mol. In Figure 2.43 the correlation plot between the predicted and experimental free energies is shown. The graphs confirm what the statistics suggest, the accuracy of the MD-MSM simulation with a cylindrical box are significantly worse than those obtained with a spherical one. Four of the five ligands, with a difference in experimental affinities of around 2 kcal/mol, were predicted to have the same binding free energy, within a 0.25 kcal/mol range.

The analysis of the PMF profiles, shown in Figure 2.44, points to inaccurate models, just like was observed in the simulation with the spherical box. The profiles show that in only two ligands, EK2 and EK9, the minimum was located at a small distance of the active site. In the rest, minima were situated at different places, for EK6, the ligand was placed just outside of the binding site, the minimum for E63 was inside a small cavity close to the active site, and finally, the EK3 minimum was found in another cavity in the protein surface at almost 20 Å from the active site. The structures of the minima described here can be seen in Figures C.3 and C.4 in section C.2 of the appendix.

In light of the results obtained with MSM-PELE and the ERK2 receptors we hypothesised that one of the causes for the poor predictive performance was the implicit solvent of PELE. After redoing the same test with MD simulations, which included explicit waters, we see that the solvent model was not the main problem, as in the MD simulation the active conformation was often misidentified. However, the explicit waters did show a small improvement with respect to PELE, as the minimum was correctly identified in some of the cases (even if the model was still noisy), while none of the ligands showed any success with the implicit solvent in PELE.




Figure 2.42: One-dimensional projection PMF for all the ligands used in the ERK2 set in the simulations with a spherical simulation box.




Figure 2.43: Correlation of the predicted binding free energy with respect to experimental values for the ERK2 system with the longer simulations and a cylinder shaped simulation box. Text labels identify the ligand that each point represents.




Figure 2.44: One-dimensional projection PMF for all the ligands used in the ERK2 set in the simulations with a cylindrical simulation box.

2.4.4 Effect of the previous AdaptivePELE simulation

As done in the PELE-MSM section, we attempted to evaluate the acceleration provided by the use of a previous exploratory AdaptivePELE simulation. We did so by running an additional set of simulations with the plasmin system starting from a fully solvated structure. Simulation and MSM parameters were kept the same, thus one can obtain a measure of the acceleration by comparing the prediction statistics.

Figure 2.45 shows the correlation of the estimated binding free energies from the new simulations with the experimental values. If compared with Figure 2.31, it is clear that the correlation observed has decreased significantly. All but two ligands were estimated at around -3.5 kcal/mol, forming no clear trend. The remaining two ligands, L01 and L03 were grossly overestimated, with over 1.5 kcal/mol error. One should note that, again, the L03 ligand results in an inaccurate estimation, suggesting that this molecule in particular is problematic for our setup. The bad predictive performance was confirmed by the reduction in the predictive metrics, with R^2 going from 0.343 to 0.178, Spearman correlation from 0.66 to 0.41 and RMSE slightly increasing from 0.64 kcal/mol to 0.79 kcal/mol.




Figure 2.45: Correlation of the predicted binding free energy with respect to experimental values for the plasmin system starting the simulation from fully solvated structures. Text labels identify the ligand that each point represents.

In section B.2 of the Appendix, the PMF profiles for all 16 ligands are shown. As was seen in the simulations following the entire protocol, the obtained landscapes strongly agree

with the assumptions made in the free energy calculation. All of them show a well-defined minimum, with roughly constant values at distances larger than 10 Å.

In summary, MD-MSM has not shown an improvement in accuracy or efficiency with respect to PELE-MSM. For plasmin, results were overall similar to those achieved with PELE, albeit with a larger set of ligands. For URO, the obtained results with reduced set of ligands were initially significantly worse than the ones obtained with PELE, requiring a change in the region accessible to the ligand and remarkably long simulations, that required days of compute time to produce a positive correlation. Finally, for the ERK2 system, the statistics measured were improved with respect to the ones obtained with PELE-MSM. Nonetheless, further analysis showed models that predict binding poses and sites wildly different than those observed crystallographically. Despite seeing an improvement in accuracy by the use of a previous enhanced sampling simulation, using unbiased MD simulations does not appear to be a viable strategy for the estimation of absolute binding free energies.

2.5 AdaptivePELE-MSM

The final version of the protocol for the calculation of binding free energies developed consisted solely of enhanced sampling simulations, both for exploration and production runs. This schema, named AdaptivePELE-MSM, was structured similarly to the previous protocols but using a modified AdaptivePELE schema for the production simulations.

In the initial tests using AdaptivePELE for production runs shown in section 2.1, we observed that the results were less accurate than those obtained with regular PELE simulations. Our main hypothesis to explain the lack of accuracy was based on the perturbation introduced by the spawning method, which redirected the sampling efforts to the less populated states. This strategy, that we have named here *inverselyProportional*, is often referred to in the literature as *counts* adaptive sampling and has been shown to efficiently explore the high dimensional energy landscape in protein systems, but to spend significant simulation time in high energy state that are not too relevant for the dynamics (Zimmerman et al., 2018; Shamsi et al., 2018; Betz and Dror, 2019).

We hypothesised that the cited shortcoming of the spawning method was the main cause of the poor predictive performance observed in the AdaptivePELE simulations. In sections 2.3 and 2.4, we evidenced an improvement in the accuracy of the binding free energy prediction with the introduction of a previous short, exploratory AdaptivePELE simulation. However, the tests with more complex systems, like ERK2 and PR, demonstrated that the limited sampling capabilities of unbiased molecular simulations were a limiting factor for the estimation of binding free energies.

Thus, in need of methods that provide more sampling of the relevant events (ligand binding and unbinding) we turned again to AdaptivePELE. In the new production simulations, we used the same clustering used in the MSM estimation, the K-Means algorithm instead of the leader one. In the spawning step, we built an MSM that would serve as a guide for the sampling in the next iteration. To choose the most relevant states, we followed

the adaptive sampling approach presented by Hinrichs and Pande (Hinrichs and Pande, 2007). There, the authors developed an score function based on the uncertainty of the eigenvalue estimation in an MSM, which can be calculated as

$$r_i = \left(\frac{\bar{q}_i}{w_i + 1} - \frac{\bar{q}_i}{w_i + m + 1} \right) \quad (2.3)$$

where r_i is the score for each cluster, w_i is the sum of the transitions from state i , m is the number of new trajectories that will be spawned (the number of processors), and \bar{q}_i is a factor that measures the sensitivity of the eigenvalues with respect to the transition probabilities, more detail on the derivation of equation 2.3 are available in the original publication (Hinrichs and Pande, 2007).

The resampling strategy is based on focusing the processing efforts on those states that contribute the most to said uncertainty to improve the overall quality of the model. After the less converged states are selected, a new round of simulations begins, in the exact same fashion of the original AdaptivePELE algorithm. The post-processing, MSM and binding free energy estimation steps were the same used in the previous versions of the protocol.

To test the new version of the protocol we used a fraction of the benchmark used to test the PELE-MSM method, consisting of two proteins and twelve ligands: URO with ligands 1UP, 2UP, 4UP, 6UP, 7UP, 675 and 39L (see Table 1.2 for ligand structures and experimental affinities); and PR with ligands STR, AS4, HCY, 356 and 400 (see Table 1.3 for ligand structures and experimental affinities).

2.5.1 Binding Free Energy for the URO System

To test the efficiency of the new clustering and spawning strategies, we ran the new protocol, AdaptivePELE-MSM, with the URO receptor, with the same set of ligands that we used in the PELE-MSM test. We ran three different versions of the protocol, i) with production simulations of 200 steps per epoch, ii) with production simulations of 400 steps per epoch and iii) with production simulations of 600 steps per epoch. All versions used 128 cores. In this manner, we can not only evaluate the effectiveness of the new strategy by comparing the prediction accuracy with the PELE-MSM runs from section 2.3.2, but we can also analyse the impact of the epoch length parameter. The MSM lag time was set at half the epoch length, and the number of clusters was chosen to be 100, like in the PELE-MSM protocol.

For all three versions, we ran three epochs, and calculated the binding free energy incrementally at the end of each iteration. That is, after the second epoch the trajectories of epochs 1 and 2 were used, and after the third one all trajectories were used. In Figure 2.46 we show the correlation with respect to the experimental binding free energies for the three protocols using the three computed epochs. All three plots show an overall good ranking of the ligands, with considerable separation between ligands 39L and 7UP and the rest. Also, for all three cases, the 2UP ligand is overestimated with respect to

ligand like 1UP, something that was also observed in the PELE-MSM runs with the URO receptor.




Figure 2.46: Correlation of the predicted binding free energy with respect to experimental values for the URO system in the simulations with AdaptivePELE-MSM, in the final epoch. Text labels identify the ligand that each point represents.

To evaluate the impact of the spawning strategy we calculated the same three statistics used throughout this thesis (R^2 , Spearman correlation and RMSE) at the end of each epoch. The results are shown in Figure 2.47. We can observe two general behaviours. For simulations with 400 and 600 steps per epoch, we see small improvements as the epochs accumulate, but the same is not true for the runs with 200 steps per epoch.

If we look at the evolution of the coefficient of determination (Figure 2.47a), we see that for the simulations with 400 and 600 steps, there is no improvement as the epochs accumulate, with a value around 0.6. The adaptive procedure did not appear to yield great improvement in those cases. On the other hand, the run with 200 steps per epoch showed a remarkable improvement from the second epoch to the first, and stayed roughly constant in the third epoch. The final value for R^2 in this version was lower than those observed in the simulations with 400 and 600 steps.




Figure 2.47: Evolution of the measured statistics for the URO system in the simulations with AdaptivePELE-MSM, as more epochs are used.

The analysis of the Spearman correlation (Figure 2.47b) is similar to that of the coefficient of determination. The runs with 400 and 600 steps show similar values for all three epochs. The simulations with 200 steps per epoch, however, showed a clear improvement for all epochs, and finished with roughly the same accuracy for this metric as the other two versions of the AdaptivePELE-MSM protocol, with a fraction of the simulation time.

Finally, we look at the evolution of the RMSE (Figure 2.47c). Here, all three variants of the AdaptivePELE-MSM protocol show a similar behaviour, with an increase in error as the epochs accumulate. In section D.1 of the Appendix, we show the PMF profiles for all seven ligands after the third epochs for the runs with 200 steps (Figures D.1 and D.2), 400 steps (Figures D.3 and D.4) and 600 steps (Figures D.5 and D.6). In all cases, the profiles were clean, with only a single minimum at close distances to the bound conformation, no secondary minima and roughly constant PMF values at distances larger than 10 Å.

2.5.2 Binding Free Energy for the PR System

The first test with AdaptivePELE-MSM showed a predictive performance similar, or slightly better, for the URO receptor. Thus, we extended the testing benchmark including the PR protein, which posed a significant challenge for PELE-MSM in terms of sampling. We ran four different versions of the protocol with different epoch lengths: i) 200 steps, ii) 400 steps, iii) 600 steps, and iv) 1000 steps. All versions, like the simulations in the previous section used 128 processors. The MSM parameters were chosen in a similar fashion to URO, with the number of clusters fixed at 100 and the lag time set at half the epoch length, to maximise the gathered statistics used.

Like we did in the URO case, we estimated the binding free energy incrementally, along the 14 calculated epochs. This number was higher than for the URO system due to the higher sampling requirements of PR. In Figure 2.48 we show the correlation with respect to the experimental binding free energies for all four epoch lengths tested. It is clear from the plots that the protocol was not able to accurately predict the binding free energy. For all four versions tested there is no clear correlation between the estimated and experimental values. With the long epochs, however, there is some apparent correlation, albeit low. On the other hand, for simulations with 600 and 200 steps the points are roughly uniformly distributed, and for the case of 400 steps there is a small anti-correlation. The different distributions suggest that the predictions were noise-driven and the sampling obtained was not sufficient to obtain an accurate model of the ligand binding to PR.

The evolution of the prediction statistics agrees with the assessment above. The plot of the coefficient of determination (Figure 2.49a) shows mostly no correlation, with some fluctuations in the 400, 600, and 1000 steps simulations. Similarly, the evolution of the Spearman correlation (Figure 2.49b) shows values fluctuating between 0 and -0.75 for all simulations but the one with 1000 steps per epoch. For this last case, starting from the fifth epoch the Spearman correlation value above 0.5, which accounts for the small correlation seen in Figure 2.48d. The RMSE (Figure 2.49c) shows a clear decreasing trend as the epochs accumulate, with a lower error for the longer simulations.

In section D.2 of the Appendix, we show the PMF profiles for the five ligands, estimated using all simulation epochs. The profiles, shown in Figures D.7 to D.14, present two differentiable patterns. In one, exemplified in Figures D.7a and D.7c the minimum is correctly predicted at low distances, but the PMF does not show a flat, constant value at larger distances, but instead all points above 12 Å form a secondary minimum, with an abrupt barrier at 10 Å. The second pattern, illustrated in Figures D.7b and D.8b, clearly distinguishes two minima, as the first pattern, but with their weights switched. That is, the absolute minimum, in these cases, is the one above 12 Å, while the minimum around 0 Å has a higher PMF value. Like in the first pattern, the two minima are separated by a spiked barrier at around 5 Å.

Overall, the reduced benchmark tests for AdaptivePELE-MSM have shown promising but partially unsatisfactory results. For URO, the ordering of the ligands was mostly correct, but the RMSE was high. On the other hand, the results for the PR system were disappointing. Despite using a considerable amount of computational power, the correlation with the experimental values was poor, and the obtained models pointed to




Figure 2.48: Correlation of the predicted binding free energy with respect to experimental values for the PR system in the simulations with AdaptivePELE-MSM, in the final epoch. Text labels identify the ligand that each point represents.

significant sampling issues. The adaptive sampling reward function used here has not provided sufficient acceleration to tackle a highly complex system like PR.




Figure 2.49: Evolution of the measured statistics for the PR system in the simulations with AdaptivePELE-MSM, as more epochs are used.

Chapter 3

Discussion

3.1 Protocol development and the importance of scientific software

Moving beyond the traditional paradigm of theory and experimentation, computational science has transformed many aspects of research. Whether simulations are regarded as the “third pillar” of science or just a set of methods for theoretical work, it is virtually impossible to find a scientific field that does not benefit from simulations.

In works where the development of new methods takes a central role, such as this one, an important amount of work hours are devoted to software development tasks, which are often “invisible” to other scientists. This is in part to be expected, as the main goal of any research endeavour, even the computational ones, is to better comprehend some natural phenomenon. This development work, however, is of crucial importance, and if done with best practices in mind it can allow for investigations that would not be possible without.

Good software is written to be understood by humans, so it is maintainable and welcomes contribution and collaboration. Maintainability is crucial if the program is expected to be used by other scientists, or by the authors in future times. Programs that are not catered to a single use should be written in a clear manner, both for ease of use and for reproducibility. This last point is essential to ensure the quality of the scientific results. Further, good code should be modular, to facilitate the incorporation of new techniques or ideas. In important, open problems like the prediction of binding free energies failing to stay up to date with newly discovered ideas could prevent altogether the production of relevant research. Finally, good software should be reliable, which can be achieved with thorough testing procedures, and scalable, which is particularly important if the software has to run in supercomputing facilities. The development work done in this thesis, focused in AdaptivePELE and the automatic protocol, was carried out with such good practices in mind.

An important challenge we encountered was to make the software robust enough to run in different machines. As an example, during the development of the extension of Adap-

tivePELE that allowed us to run MD simulations, we had to adapt the software to run on a GPU cluster, which raised particular issues that needed to be resolved. Since our approach relies on ensembles of simulations, we wanted to run them in parallel on multiple nodes. We implemented a basic communication protocol in AdaptivePELE, which allowed us to run the ensemble of simulations in many nodes in parallel (typically we used 10 nodes) and using the 4 GPUs of each node. This had the advantage that we only needed to run one command to launch the simulations of all ligands, which in turn increased our throughput, so we could study sets of four or more ligands for multiple proteins, and even test several configurations, using MD simulations.

The development of a fully automatic protocol for the estimation of binding proved crucial to the study of binding free energies in a benchmark composed of four proteins. The main characteristics of the protocol are user-friendliness, modularity and a high degree of automation. The modularity of the implementation allowed us to test three different simulation techniques, with little tuning, so their strengths and weaknesses could be easily compared. Its user-friendliness and automation level were crucial to study multiple proteins with several ligands each. Automating cumbersome tasks like protein preparation steps, or ligand parametrisation eliminated points that require manual attention, greatly increased our efficiency. The protocol implementation not only facilitated running the simulations, but incorporated also the subsequent analysis, including the construction of the MSM and the estimation of the binding free energy. Even if further analysis needed to be customised on a per-system basis, this automatic, preliminary analysis provided us with the opportunity to quickly assess the performance and plan, when needed, further testing or parameter variation.

The software development work done during this thesis has made possible the computational estimation of protein-ligand binding free energies using pathway-based methods on a number of proteins and ligands that, to the best of our knowledge, has not been published before. Despite the advancements, the problem of binding free energies is far from solved, as will be clear in the next section, and surely new insights will be uncovered in the future that will bring us closer to the solution. Making it possible to incorporate these still unknown ideas shows that time spent in hidden tasks like testing and designing are good time investments for a computational scientist.

3.2 Successes and limitations of presented protocol for binding free energies

As we mentioned in the previous section, the automatic protocol we developed allowed us to estimate the binding free energy for five proteins. One of them, trypsin, was only used for the initial tests of the software, but the remaining four were studied using three different simulation methods, with varying degrees of success.

Firstly, the system here referred as plasmin is a kringle domain of plasminogen. This is a small protein, of only 80 residues, which allowed for extensive sampling. We predicted the binding to a set of six ligands using PELE simulations and 16 ligands using MD. The

prediction was successful in both cases. With PELE-MSM, we obtained an R^2 of 0.72, a Spearman correlation of 0.75 and an RMSE of 1.11 kcal/mol. With MD-MSM, the larger amount of ligands used caused the R^2 value to decrease to 0.343, but the Spearman correlation was similarly high at 0.66 and error was lower, with an RMSE of 0.64 kcal/mol. In both cases the ligands were mostly correctly ranked, with a low error in the absolute binding free energy values.

Secondly, we used the urokinase-type plasminogen activator protein. URO was a test system for all simulation methods selected with the protocol, PELE, MD and AdaptivePELE, and was also for the runs using PELE and AdaptivePELE before developing the protocol. The accuracy for all three versions was similar. With PELE-MSM, we obtained an R^2 of 0.59, a Spearman correlation of 0.61 and an RMSE of 5.86 kcal/mol. With AdaptivePELE-MSM, the collected statistics with an epoch length of 600 steps, which seemed to give the best performance, were an R^2 of 0.632, a Spearman correlation of 0.75 and an RMSE of 5.35 kcal/mol. Finally, for the MD-MSM case we observed an R^2 of 0.761, a Spearman correlation of 0.6 and an RMSE of 5.64 kcal/mol. The results for PELE and AdaptivePELE were similar, although as noted in section 2.5.1 the AdaptivePELE runs took about half the computing resources that PELE required. The simulations with MD-MSM used only four ligands and the accuracy with that reduced set was on par, for the runs using a cylindrical box and long simulations, but was significantly worse for the other parameters tested, as we will discuss in detail in section 3.3.2. Overall the ranking in the ligand was good, but slightly worse than for plasmin. The absolute error, however, was much higher.

Third, we ran the PELE and MD versions of the protocol with the mitogen-activated protein kinase 1. This protein is similar in size to URO, but its binding site is larger (and more exposed), which complicates the binding free energy predictions. For both versions of the protocol we used the same five ligands. With PELE-MSM we obtained an R^2 of 0.01, a Spearman correlation of 0.0 and an RMSE of 6.98 kcal/mol, while in the MD case the values were an R^2 of 0.526, a Spearman correlation of 0.9 and an RMSE of 4.7 kcal/mol. The measured statistics seem completely different, with MD performing significantly better than PELE. In both cases, however, analysing the PMF profiles made clear that the estimated models could not properly capture the binding process. We believe that this limitation is mostly related to modelling parameters, mainly force field inaccuracies and solvent models, as we will discuss further in section 3.2.2.

The last protein studied was the progesterone receptor, which is a protein that presents significant sampling challenges due to its buried binding site. We ran PELE-MSM and AdaptivePELE-MSM with this protein, using five ligands in both cases. In this case, PELE-MSM showed no predictive power, with an R^2 value and a Spearman correlation of 0 and an RMSE of 4.41 kcal/mol. With AdaptivePELE-MSM we obtained similar results, but slightly better statistics in the best case with 1000 steps per epoch. After running 13 epochs, the measured statistics were an R^2 of 0.355, a Spearman correlation of 0.8 and an RMSE of 4.38 kcal/mol. Despite the seemingly improved results with AdaptivePELE, the analysis of the PMF profiles shows that the obtained models are not much more precise than the PELE ones. The PR receptor constitutes a prime example of sampling limitations for binding free energy estimation, as will see in section 3.2.1.

3.2.1 Sampling limitations

It is well-known that the usability of molecular simulation methods is hampered by the long time scales at which many processes of interest happen. A clear example of such processes is protein-ligand binding. Non-covalent complexes might occur on a time scale of milliseconds, or even longer, which makes simulations that sample multiple occurrences of binding and unbinding prohibitively expensive.

This well-known limitation motivated the central idea of this thesis, the combination of short exploratory enhanced simulation with a longer simulation for data collection, that takes advantage of information acquired during the first one. Despite this, we observed clear signs of insufficient sampling in some of the tests, particularly in the ones run with the PR system.

In the PELE-MSM runs for PR we observed a clear lack of transitions in and out of the binding site, as measured by the SASA value. In Figure 3.1 we show the evolution of the SASA during the production simulations of the three iterations for the STR ligand with PR (one should note that this plot is the same as in Figure 2.27, we reproduced it again here for easier reading). There is a clear gap in sampled SASA values around 0.2, which marks a transition state between conformations that have entered the binding site. This lack of sampling prevents the estimated model from properly reproducing the protein-ligand binding process, and ultimately explains the poor predictive performance observed with PELE-MSM and the PR protein.

The AdaptivePELE-MSM version of the protocol was designed to improve the sampling issues encountered with PELE-MSM and the PR system. In Figure 3.2 we show the plots of the evolution of the SASA for the AdaptivePELE-MSM production simulations for the four epoch length tested (200, 400, 600 and 1000 steps). If we compare those plots with the PELE one, we can see that the sampling is increased in the intermediate region, with better results as the epoch length is increased. The increase in sampling is also reflected in the measured statistics, particularly for the simulations of 1000 steps per epoch. However, the predictive performance is still poor and the resulting models, as reflected by the PMF profiles, show the same inaccuracies that the ones obtained with PELE-MSM (see section 3.3.1 for more details on the accuracy comparison of AdaptivePELE-MSM and PELE-MSM).

We also approximately measured the amount of transition events in and out of the binding site for each epoch length and ligand. We did so by counting, in each trajectory, how many times the ligand goes from a distance less than 2.5 Å to a distance larger than 10 Å to the crystallographic pose, for unbinding events, and the opposite for binding events. The results are shown in Figure 3.3, with number of bindings in Figure 3.3a and number of unbinding events in Figure 3.3b. If we look at the number of binding events, the only set of simulations that reliably produces such transitions is the one with 1000 steps per epoch, which still failed to find any binding for the ligand labelled 400. The number of events decreased with the epoch length, the simulations with 600 steps found binding events for only three of the ligands, the runs with 400 steps only for two and the shortest simulations found no such event for any of the ligands. The situation is similar but slightly better when looking at the number of unbinding events. Again the longest simulations




Figure 3.1: SASA evolution during the production simulation for three iterations of 127 trajectories each for the PR system with the STR ligand.

were the most reliable finding unbinding for all five ligands, even if the amount is low. The same was not true for the other runs, with unbinding events detected in four ligands for the simulations with 600 steps, but only in one ligand for the runs with 400 and 200 steps.

Despite the additional sampling that AdaptivePELE-MSM introduced, the number of transitions in and out of the binding site is still limited, which prevented an accurate binding free energy estimation. In a previous study with PELE (Grebner et al., 2017) the accurate binding affinity prediction to PR was achieved using three ligands. In that paper, the authors reportedly used 600 trajectories for 48 hours of wall clock time for each ligand. With PELE-MSM, we used 3 iterations of 1000 steps with 128 trajectories, which required around 24 hours of wall clock time, approximately 9 times less if we consider the number of CPU hours. The longest AdaptivePELE-MSM simulations used 128 processors for 14 epochs of 1000 steps, with each epoch requiring approximately 6 hours of wall clock time, amounting to around 5 times more than PELE-MSM and 2.7 times less than the original publication by Grebner et. al. Comparing the resources used by each paper and the resulting accuracy it is clear that the amount of computational resources required to obtain accurate binding free energy predictions for the PR protein is exceedingly high, and the methods presented here do not allow for a reliable estimation of multiple ligands in a fast manner.




Figure 3.2: SASA evolution during the production simulation with AdaptivePELE-MSM for the STR ligand and PR. The colour represents the simulation epoch. Panel (a): Simulations with 200 steps per epoch. Panel (b): Simulations with 400 steps per epoch. Panel (c): Simulations with 600 steps per epoch. Panel (d): Simulations with 1000 steps per epoch.

3.2.2 Modelling limitations

Aside from insufficient sampling, the other most common cause of failures in binding free energy predictions is modelling inaccuracies, mainly attributed to limited force field accuracy. Inadequate sampling should be identified almost always as the cause of prediction failures as current simulation methods are incapable of exploring the time scales to reach true equilibrium, as clearly exposed by David Mobley (Mobley, 2012). True equilibrium sampling is impossible to assess, and methods that measure simulation convergence are limited to attempt to identify cases of insufficient sampling, but can never ensure that sampling was sufficient in simulations of systems as complex as proteins.




Figure 3.3: Plot of the number of transitions obtained for the simulations of the PR receptor using different epoch lengths with AdaptivePELE-MSM. Panel (a): Plot of binding events. Panel (b): Plot of unbinding events.

Here we will discuss the poor results obtained with the ERK2 systems, that are likely attributable to a combination of sampling and modelling issues, with emphasis on the latter. With PELE-MSM the prediction showed no correlation, and the estimated models were not even able to properly identify the (experimentally known) binding pose. While the relatively short simulations likely produced insufficient sampling, the model PMF profiles and plots of internal energy point to severe modelling limitations that are most likely the key factor. If we compare the potential energy calculated by PELE in the PELE-MSM runs that use and implicit solvent with the one obtained with a combined implicit solvent and four added explicit solvent molecules (Figure 3.4, note that the two plots were already shown in the Results chapter, Figures 2.22 and 2.24b, and are only shown here for the sake of discussion) we observe significant differences in the energy values at distances of less than 5 Å, and a displacement of the energy minimum towards the low distance region. This suggests that the interactions that the introduced water molecules form are important to stabilise the binding pose, and without them the implicit solvent introduces an error that prevents the accurate estimation of binding free energies. The water molecules could also prevent the ligand from adopting alternative binding poses by filling spaces that would be free in pure implicit solvent models. It should be noted, however, that the simulation with explicit waters is not exactly comparable to the PELE-MSM runs, since it was only a quick test using a functionality of PELE that is still under development.

Since PELE did not fully support an explicit solvent model, we introduced the ERK2 as a test system for the MD-MSM version of the protocol. In those runs, the results, measured by the three statistics that we have used throughout this thesis, were much better than those of PELE, but the PMF profiles still showed issues identifying the binding pose. We analysed the evolution of the R^2 and the Spearman correlation for the two test runs, one with a spherical simulation box and one with a cylindrical one (Figure 3.5),




Figure 3.4: Plot of PELE’s internal energy values with respect to the centre of mass distance in the PELE-MSM simulations of the ERK2 receptor with ligand EK2. Panel (a): Simulation with implicit solvent. Panel (b): Simulation with implicit solvent and four explicit waters.

as the simulation length increased. The evolution for the simulations with a spherical box (Figure 3.5a) showed a consistent improvement of both metrics as the simulation length increased, up to around 400 ns. From there on, the value of the metric does not show great variations, only small fluctuations. Similarly, for the simulations with a cylindrical box, there is a clear increasing trend until the simulations reach 250 ns. From there, the Spearman correlation stays roughly constant, but the R^2 value shows a clear decreasing trend. These plots support the idea that inadequate sampling is not the main factor to explain the inaccurate models with ERK2, as extending the simulation length renders no improvement or even a reduction in the prediction accuracy. On the other hand, the results with MD-MSM indicate that while the use of explicit solvent appears to have a measurable impact in the model accuracy, it can not be the only cause of poor performance, given the lack of accuracy of some of the models obtained with MD simulations with explicit waters.

In this work, we have illustrated, through two protein systems, the difficulty in establishing the shortcomings in simulations for free energy calculations. The two most common issues, inadequate sampling and inaccurate models, are often intertwined and it is extremely difficult to precisely establish to what extent is each to blame in the simulations failures. Lack of sampling is regarded as the first limitation that should be ruled out, as it is easier to measure than the modelling inaccuracies, that can come from multiple sources. In the PR system, that seemed to be the case, as the insufficient sampling was easily measurable, and a previously published work with a similar set-up but more sampling had significantly better results. In the ERK2 system, there was no evident sign of insufficient sampling, although that does not guarantee that sampling was adequate. It was clear, however, that the model used did not properly capture the binding process as evidenced by the incapacity of identifying the binding pose. This could be attributed to a combination of issues with the solvent model (particularly in PELE simulations) and the force field




Figure 3.5: Analysis of the evolution of the prediction statistics (Spearman and R^2) as the simulation length increases for the ERK2 system. The starting point is based on adding the first 75 ns segments of all trajectories (with subsequent increments of 25ns). Panel (a): Statistics evolution for the simulations with a spherical simulation box. Panel (b): Statistics evolution for the simulations with a cylindrical simulation box.

parameters (both in PELE and MD simulations).

3.3 Effectiveness of unbiased simulations for free energy calculations

To overcome the limited sampling offered by unbiased molecular simulations for long time scale events like protein-ligand binding, enhanced sampling methods have been the object of extensive research. Many of these techniques require a complex and system-dependent setup, such as the discovery of efficient collective variables for metadynamics. Since one of the goals of this thesis was to obtain a method that could predict the binding free energies for different systems, using techniques with a system-specific setup would not be feasible.

As a compromise between sampling efficiency and user-friendliness we decided to combine both enhanced and unbiased simulations. We first used a quick exploratory simulation with AdaptivePELE, which required no specific setup and could be easily automated to work with different proteins, and combined it with regular unbiased simulations. This approach worked well for plasmin and URO, two systems that possess a well-defined, tight and accessible binding site, but failed with systems with more flexible binding sites and deeply buried (less accessible) ones.

3.3.1 Role of enhanced sampling methods

In sections 2.3.5 and 2.4.4 we tested whether the initial AdaptivePELE simulation provided a noticeable sampling improvement to the subsequent production simulations using unbiased methods. There, running simulations after this initial step resulted in less noisy models and a better estimation, as measured by the statistical metrics used (R^2 , Spearman correlation and RMSE). Nonetheless, the tests with the PR and ERK2 receptors showed that even this accelerated sampling was not sufficient for the accurate prediction of binding free energies.

A final version of the protocol was developed that removed altogether unbiased simulation methods from the protocol. With an adaptive sampling based approach (Hinrichs and Pande, 2007), this new version, AdaptivePELE-MSM, showed an acceptable binding free energy estimation with reduced computational resources for the URO system, but it did not prove successful for PR.

Before delving further into the analysis of the PELE and AdaptivePELE-MSM results for URO, we should note the differences in the amount of sampling used in each version of the protocol. Three factors need to be taken into account: the number of processors, the number of steps per simulation and the number of runs. The number of runs and the number of processors is the same in all four versions, 3 and 128. Finally, the number of steps per simulation was 1000 for the PELE-MSM simulation, 200, 400 or 600 for the AdaptivePELE simulations. The total amount of sampling used for the PELE-MSM results was 5 times that used for the 200 steps per epoch run, 2.5 times the time used for the 400 steps per epoch run and about 1.6 times the amount of sampling used for the simulation with 600 steps per epoch.

In Figure 3.6, we plot the evolution of the measured statistics after each AdaptivePELE-MSM epoch (lines with markers) and PELE-iterations (solid, horizontal lines). In some graphs only two horizontal lines are visible, that is caused by two PELE-MSM values that are so close that their lines are overlapped. For all three metrics, there is little difference and the fact that increasing the PELE-MSM iterations does not generally improve the estimations points to random fluctuations as the most likely explanation. Judging from the comparisons, it would seem that with only one simulation of 400 PELE steps is sufficient to obtain an acceptable ranking, with R^2 and Spearman correlation around 0.6. The RMSE value seems initially better for the AdaptivePELE-MSM runs, but as the epochs accumulate the value increases to only 0.5 kcal/mol lower than the PELE-MSM runs.

The positive results of AdaptivePELE-MSM with URO contrast with the negative results obtained with AdaptivePELE simulations before establishing the protocol (see section 2.1.2 for more details). The main difference between the two runs is the spawning strategy used, in the early simulations the reward function considered as more relevant those state that were less visited, while the AdaptivePELE-MSM reward function gave priority to those states that contributed more to the uncertainty of the eigenvalue estimation. Likely, the former gave worse results due to tendency, observed multiple times, to sample high energy states that are not too relevant for the dynamics (Zimmerman et al., 2018; Shamsi et al., 2018; Betz and Dror, 2019).




Figure 3.6: Evolution of the measured statistics for the URO system in the simulations with AdaptivePELE-MSM, as more epochs are used. The horizontal line marks the statistics measured using PELE-MSM protocol after each iteration.

Contrary to the URO case, and aiming at evaluating sampling deficiencies, in the PR runs with AdaptivePELE-MSM we used more computational power than with PELE-MSM. The runs used the same number of processors, but while in PELE-MSM we used 3 iterations, with AdaptivePELE-MSM we used 14. Considering the different epoch lengths, with the 200 steps per epoch we roughly matched the sampling effort of PELE-MSM, with 400 we double it, 600 was thrice as expensive and for the longest epochs (1000 steps), the resources used were almost five times as much.

In figure 3.7 we show the evolution of the statistics for all four epochs lengths, and compare them to the value obtained with all three PELE-MSM iterations (horizontal line). For the coefficient of determination we see that the shortest run stays in line with the PELE-MSM result, while the longer simulations fluctuate above the PELE mark, but end with low correlation values. The Spearman correlation paints a different picture. There is a clear difference between the 1000 steps per epoch simulation and the rest. Again, the simulation of 200 steps shows no correlation, with small fluctuations, in line with the

PELE results. The runs with 400 and 600, show fluctuations between 0 and -0.75. The runs with 1000 steps, however, shows a clear improvement with respect to PELE starting from the fifth epoch, showing a partially correct ranking of the ligands, with a Spearman correlation finishing at 0.5. The error, measured through the RMSE, despite decreasing with more sampling stays always higher than the one observed with PELE.




Figure 3.7: Evolution of the measured statistics for the PR system in the simulations with AdaptivePELE-MSM, as more epochs are used. The horizontal line marks the statistics measured using PELE-MSM protocol.

The insufficient sampling recorded by the PR runs, together with the lack of improvement observed with URO using AdaptivePELE-MSM indicate that better reward functions are needed. In the field of MSM theory, there have been important advances that allow the evaluation of the quality of a constructed MSM, in terms of the variational approach to conformational mechanics (Nüske et al., 2014) or the GMRQ score (McGibbon and Pande, 2015). However, these schemes generate a score for the entire model and are useful for tuning different parameters, such as the number of clusters. But a reward function inside an adaptive sampling method needs a way to score each individual state,

in order to select the most promising one for further rounds. Therefore, the search for better spawning criteria to improve binding free energies remains an open question, and an interesting avenue for future research.

3.3.2 Sampling efficiency of longer vs multiple shorter simulations

Throughout this thesis, we have relied on MSMs for the analysis of the simulations, and to obtain an estimation of the equilibrium probability. These kinds of models are very popular due to the vast amount of information that they offer and their capacity to integrate ensembles of short parallel simulations into a single model. There is, however, a potential issue that might affect the quality of the Markov model if the simulations are not long enough. As pointed out in early groundwork (Swope et al., 2004), the lag time used to construct the model must be at least as long as the longest equilibration time among the states used to discretise the simulations. That in turn pushes the simulation length to several times this lag time, in order to be able to collect sufficient statistics.

We suspected that this effect could be at play in the early runs of the MD-MSM protocol with the URO receptor. For this system, we ran three different sets of simulations, one with 120 trajectories of length 100 ns, one with 40 trajectories of length 1 μ s with a spherical box and one with 40 trajectories of length 1 μ s with a cylindrical box. The first set provided the worst predictive performance, with an anti-correlation of the predicted values versus the experimental ones. This anti-correlation was not found in any of the two other sets, with longer simulations. The performance of the simulations with the spherical box was not great, grouping three of the four ligands at the same predictive value, while the difference in the experimentally measured affinity was greater than 1 kcal/mol for each. This was not true for the final set of simulations, with the cylindrical box. There, there is a weak correlation between the predictions and experimental values, but at least the predictions correctly separate the two ligands with higher affinity from the two with lower affinity.

To characterise the sampling obtained by each set of simulations we approximated the number of binding and unbinding events by measuring the ligand distance to the bound pose and counting how many times each simulation went from a bound pose to free one, and vice versa. We considered as bound all conformation with a distance less than 2.5 Å, and as unbound those with a distance larger than 6 Å. This analysis was not intended as a rigorous characterisation of the protein-ligand complex, but merely as a lower bound on the true number of binding and unbind events. The results are shown in Figure 3.8. The simulations of 100 ns are labelled as “Short, sphere”, the longer ones with a spherical box as “Long, sphere” and the ones with the cylindrical box are labelled “Long, cylinder”.

The number of binding events (Figure 3.8a) is bigger for the shorter simulations with ligand 1UP and 2UP, which are the one with lowest affinity. However, for the other two ligands, the simulations that produced more transitions are the long ones with the cylinder box. Both the short and long simulations with the spherical box show less binding events, which could be explained by metastable trapping of the largest ligands in small cavities in the protein surface, as suggested by the PMF profile (see section 2.4.2 for more details).

This is particularly true for the long, sphere simulations which showed only two bindings for the 39L ligand and no binding events for the 7UP. For reference, the long, cylinder simulations produced 19 and 9 transitions, respectively.




Figure 3.8: Plot of the number of transitions obtained for the simulations of the URO receptor using different simulation boxes with MD-MSM. Panel (a): Plot of binding events. Panel (b): Plot of unbinding events.

The results were quite different for the number of unbinding events (Figure 3.8b). Of the four ligands, only for 1UP we observed transitions, and merely one of the long, sphere simulations and 2 for the long, cylinder ones. The short simulations showed no unbinding events at all. Even in the version of the protocol that showed better results, the sampling is limited. To further assess this, we analysed the evolution of the statistical metrics as the simulation length increased for the long simulations, shown in Figure 3.9. In the plots, we also marked in two vertical lines the results after only considering the first 100 ns of simulation (dotted line) and the results at the point in which the long simulations match the aggregated simulation time of the ensemble of short simulations (solid vertical line). With only 100 ns simulations the results are, unsurprisingly, bad since there are three times less trajectories than in the short simulations ensemble. Comparing the point in which all three sets of runs have the same amount of sampling, however, we clearly see that the long simulations are better for this system, since while the short simulations predictions were anti-correlated, the long simulations give positive results. At this point the results of the sphere box are better than the cylindrical one, but if we observe the whole evolution it is clear that the latter produces more reliable sampling. The long, sphere runs show wide fluctuations, between no correlation and negative one, at length greater than 500 ns. On the other hand, the plot for the cylindrical box shows a much more stable evolution, with only small fluctuations and always positive values for the Spearman correlation.

Despite the apparent advantage of the cylindrical box, MD simulations have proven incapable of sampling the URO system in an efficient enough manner for the estimation of binding free energies for even moderately sized sets of ligands. Even with 1 μ s long




Figure 3.9: Analysis of the evolution of the prediction statistics (Spearman and R^2) as the simulation length increases for the URO simulations. The solid vertical line marks where the long simulations of the URO system have the same aggregated simulation time as the short ones, 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). Panel (a): Statistics evolution for the simulations with a spherical simulation box. Panel (b): Statistics evolution for the simulations with a cylindrical simulation box.

simulations, the sampling was poor and no unbinding event was observed. Similarly, in a recent paper (Pramanik et al., 2019), while comparing the efficiency and accuracy of estimations from metadynamics versus unbiased MD simulations, it was reported that 7 of the unbiased MD replicas run up to 2 μs without showing any unbinding event. It seems clear that the future of binding free energy estimation will require more sophisticated enhanced sampling methods, since unbiased techniques would need orders of magnitude more powerful computing capabilities than the ones currently available, which does not seem likely in the short-term future.

Chapter 4

Conclusions

The conclusions that can be extracted from this thesis are:

- The development of an automatic, modular and user-friendly protocol for the estimation of binding free energies has allowed us to predict binding affinities for a large amount of protein-ligand systems, with respect to comparable pathway-based methods. The protocol appeared to work best with fragment systems and protein with well defined and not very flexible binding sites.
- Unbiased simulation techniques are not efficient enough for the prediction of binding free energies for realistic protein-ligand systems in a constrained time frame. Even with optimisations designed to improve sampling, they fail to collect sufficient statistics about protein-ligand binding and unbinding processes.
- Enhanced sampling methods are required to predict binding affinities in a way that can impact real-world drug discovery campaigns. However, further work is required in development of such techniques that can be generally applied to many protein system, as well as easing their applicability.
- The design of realistic and diverse testing benchmarks is crucial to properly identify strengths and weaknesses of newly developed methods. Further work on the standardisation of such benchmarks would help in the comparison of the multiple existing methods for the prediction of binding free energies.
- Establishing software development best practises during the implementation of new techniques increases scientific productivity, even if the development time is longer. Further, establishing automatic and reproducible protocols results in less error-prone procedures.

Appendices

Appendix A

PMF plots for PELE-MSM benchmark

A.1 PMF plots for plasmin system




Figure A.1: Comparison of the evolution of the one-dimensional PMF with three iterations for the L02 ligand of the plasmin receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure A.2: Comparison of the evolution of the one-dimensional PMF with three iterations for the L05 ligand of the plasmin receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure A.3: Comparison of the evolution of the one-dimensional PMF with three iterations for the L06 ligand of the plasmin receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure A.4: Comparison of the evolution of the one-dimensional PMF with three iterations for the L07 ligand of the plasmin receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure A.5: Comparison of the evolution of the one-dimensional PMF with three iterations for the L15 ligand of the plasmin receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.

A.2 PMF plots for URO system




Figure A.6: Comparison of the evolution of the one-dimensional PMF with three iterations for the 1UP ligand of the URO receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure A.7: Comparison of the evolution of the one-dimensional PMF with three iterations for the 2UP ligand of the URO receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure A.8: Comparison of the evolution of the one-dimensional PMF with three iterations for the 4UP ligand of the URO receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure A.9: Comparison of the evolution of the one-dimensional PMF with three iterations for the 6UP ligand of the URO receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure A.10: Comparison of the evolution of the one-dimensional PMF with three iterations for the 239 ligand of the URO receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.





Figure A.11: Comparison of the evolution of the one-dimensional PMF with three iterations for the 675 ligand of the URO receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.

A.3 PMF plots for ERK2 system



(a) PMF using one iteration (b) PMF using two iterations (c) PMF using three iterations

Figure A.12: Comparison of the evolution of the one-dimensional PMF with three iterations for the EK3 ligand of the ERK2 receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure A.13: Comparison of the evolution of the one-dimensional PMF with three iterations for the EK6 ligand of the ERK2 receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.




Figure A.14: Comparison of the evolution of the one-dimensional PMF with three iterations for the EK9 ligand of the ERK2 receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.





Figure A.15: Comparison of the evolution of the one-dimensional PMF with three iterations for the E63 ligand of the ERK2 receptor for the three iterations. Panel (a) shows the PMF resulting from the first iteration only. Panel (b) shows the PMF resulting from combining the first and second iterations. Panel (c) shows the PMF resulting from combining the first three iterations.


Appendix B

PMF plots for MD-MSM benchmark


B.1 PMF Plots for plasmin system




(a) L01



(b) L02




(c) L03




(d) L04


Figure B.1: One-dimensional projection PMF for all the ligands used in the plasmin set.




(a) L05



(b) L06




(c) L07




(d) L08


Figure B.2: One-dimensional projection PMF for all the ligands used in the plasmin set.




(a) L09



(b) L10




(c) L11




(d) L12


Figure B.3: One-dimensional projection PMF for all the ligands used in the plasmin set.




(a) L13



(b) L14



(c) L15



(d) L16

Figure B.4: One-dimensional projection PMF for all the ligands used in the plasmin set.

B.2 Plots for plasmin system without previous exploration




Figure B.5: One-dimensional projection PMF for all the ligands used in the plasmin set in the simulations without a previous AdaptivePELE simulation.




Figure B.6: One-dimensional projection PMF for all the ligands used in the plasmin set in the simulations without a previous AdaptivePELE simulation.




Figure B.7: One-dimensional projection PMF for all the ligands used in the plasmin set in the simulations without a previous AdaptivePELE simulation.







Figure B.8: One-dimensional projection PMF for all the ligands used in the plasmin set in the simulations without a previous AdaptivePELE simulation.

Appendix C

Structures for MD-MSM simulations with the ERK2 benchmark

C.1 Structures of the PMF minima for the ERK2 ligands with spherical box




Figure C.1: Structure of the absolute minimum of the PMF for all the ligands used in the ERK2 set in the simulations with a spherical simulation box. Protein is shown in white cartoon, the minimum structure is coloured cyan, with the active conformation shown in dark blue for reference. Waters at less than 4 Å of the ligand are also shown.




Figure C.2: Structure of the absolute minimum of the PMF for all the ligands used in the ERK2 set in the simulations with a spherical simulation box. Protein is shown in white cartoon, the minimum structure is coloured cyan, with the active conformation shown in dark blue for reference. Waters at less than 4 Å of the ligand are also shown.

C.2 Structures of the PMF minima for the ERK2 ligands with cylindrical box




Figure C.3: Structure of the absolute minimum of the PMF for all the ligands used in the ERK2 set in the simulations with a cylindrical simulation box. Protein is shown in white cartoon, the minimum structure is coloured cyan, with the active conformation shown in dark blue for reference. Waters at less than 4 Å of the ligand are also shown.





Figure C.4: Structure of the absolute minimum of the PMF for all the ligands used in the ERK2 set in the simulations with a cylindrical simulation box. Protein is shown in white cartoon, the minimum structure is coloured cyan, with the active conformation shown in dark blue for reference. Waters at less than 4 Å of the ligand are also shown.


Appendix D

PMF plots for AdaptivePELE-MSM benchmark

D.1 PMF plots for URO system



(a) 1UP



(b) 2UP

Figure D.1: One-dimensional projection PMF for all the ligands used in the URO set in the simulations with AdaptivePELE-MSM with 200 steps per epoch.




Figure D.2: One-dimensional projection PMF for all the ligands used in the URO set in the simulations with AdaptivePELE-MSM with 200 steps per epoch.




Figure D.3: One-dimensional projection PMF for all the ligands used in the URO set in the simulations with AdaptivePELE-MSM with 400 steps per epoch.




Figure D.4: One-dimensional projection PMF for all the ligands used in the URO set in the simulations with AdaptivePELE-MSM with 400 steps per epoch.




Figure D.5: One-dimensional projection PMF for all the ligands used in the URO set in the simulations with AdaptivePELE-MSM with 600 steps per epoch.




Figure D.6: One-dimensional projection PMF for all the ligands used in the URO set in the simulations with AdaptivePELE-MSM with 600 steps per epoch.

D.2 PMF plots for PR system






Figure D.7: One-dimensional projection PMF for all the ligands used in the PR set in the simulations with AdaptivePELE-MSM with 200 steps per epoch.



(a) 356



(b) 400

Figure D.8: One-dimensional projection PMF for all the ligands used in the PR set in the simulations with AdaptivePELE-MSM with 200 steps per epoch.






Figure D.9: One-dimensional projection PMF for all the ligands used in the PR set in the simulations with AdaptivePELE-MSM with 400 steps per epoch.



(a) 356



(b) 400

Figure D.10: One-dimensional projection PMF for all the ligands used in the PR set in the simulations with AdaptivePELE-MSM with 400 steps per epoch.






Figure D.11: One-dimensional projection PMF for all the ligands used in the PR set in the simulations with AdaptivePELE-MSM with 600 steps per epoch.



(a) 356



(b) 400

Figure D.12: One-dimensional projection PMF for all the ligands used in the PR set in the simulations with AdaptivePELE-MSM with 600 steps per epoch.




Figure D.13: One-dimensional projection PMF for all the ligands used in the PR set in the simulations with AdaptivePELE-MSM with 1000 steps per epoch.




Figure D.14: One-dimensional projection PMF for all the ligands used in the PR set in the simulations with AdaptivePELE-MSM with 1000 steps per epoch.

Appendix E

Thesis summary

E.1 Thesis summary

The advances in computing power have motivated the hope that computational methods can accelerate the pace of drug discovery pipelines. For this, fast, reliable and user-friendly tools are required. One of the fields that has gotten more attention is the prediction of binding affinities. Two main problems have been identified for such methods: insufficient sampling and inaccurate models.

This thesis is focused on tackling the first problem. To this end, we present the development of efficient methods for the estimation of protein-ligand binding free energies. We have developed a protocol that combines enhanced sampling with more standard simulations methods to achieve higher efficiency. First, we run an exploratory enhanced sampling simulation, starting from the bound conformation and partially biased towards unbound poses. Then, we leverage the information gained from this short simulation to run longer unbiased simulations to collect statistics.

Thanks to the modularity and automation that the protocol offers we were able to test three different methods for the long simulations: PELE, molecular dynamics and AdaptivePELE. PELE and molecular dynamics showed similar results, although PELE used less computational resources. Both seemed to work well with small protein-fragment systems or proteins with not very flexible binding sites. Both failed to accurately reproduce the binding of a kinase, the Mitogen-activated protein kinase 1 (ERK2). On the other hand, AdaptivePELE did not show a great improvement over PELE, with positive results for the Urokinase-type plasminogen activator (URO) and a clear lack of sampling for the Progesterone receptor (PR).

We demonstrated the importance of a well-designed suite of test systems for the development of new methods. Through the use of a diverse benchmark of protein systems we have established the cases in which the protocol is expected to give accurate results, and which areas require further development. This benchmark consisted of four proteins, and over 30 ligands, much larger than the test systems typically used in the development of pathway-based free energy methods.

In summary, the methodology developed in this work can contribute to the drug discovery process for a limited range of protein systems. For many others, we have observed that regular unbiased simulations are not efficient enough and more sophisticated, enhanced sampling methods are required.

E.2 Resum de la thesis

Els grans avenços en la capacitat de computació han motivat l'esperança que els mètodes de simulacions per ordinador puguin accelerar el ritme de descobriment de nous fàrmacs. Per a què això sigui possible, es necessiten eines ràpides, acurades i fàcils d'utilitzar. Un dels problemes que han rebut més atenció és el de la predicció d'energies lliures d'unió entre proteïna i lligand. Dos grans problemes han estat identificats per a aquests mètodes: la falta de mostreig i les aproximacions dels models.

Aquesta tesi està enfocada a resoldre el primer problema. Per a això, presentem el desenvolupament de mètodes eficients per a l'estimació de d'energies lliures d'unió entre proteïna i lligand. Hem desenvolupat un protocol que combina mètodes anomenats *enhanced sampling* amb simulacion clàssiques per a obtenir una major eficiència. Els mètodes d'*enhanced sampling* són una classe d'eines que apliquen algun tipus de pertorbació externa al sistema que s'està estudiant per tal d'accelerar-ne el mostreig.

En el nostre protocol, primer correm una simulació exploratòria d'*enhanced sampling*, començant per una mostra de la unió de la proteïna i el lligand. Aquesta simulació esta parcialment esbiaixada cap a aquells estats del sistema on els dos components es troben més separats. Després utilitzem la informació obtinguda d'aquesta primera simulació més curta per a córrer una segona simulació més llarga, amb mètodes sense biaix per obtenir una estadística fidedigna del sistema.

Gràcies a la modularitat i el grau d'automatització que la implementació del protocol ofereix, hem pogut provar tres mètodes diferents per les simulacions llargues: PELE, dinàmica molecular i AdaptivePELE. PELE i dinàmica molecular han mostrat resultats similars, tot i que PELE utilitza menys recursos. Els dos han mostrat bons resultats en l'estudi de sistemes de fragments o amb proteïnes amb llocs d'unió poc flexibles. Però, els dos han fallat a l'hora de reproduir els resultats experimentals per a una quinasa, la *Mitogen-activated protein kinase 1* (ERK2). D'altra banda, AdaptivePELE no ha mostrat una gran millora respecte a PELE, amb resultats positius per a la proteïna *Urokinase-type plasminogen activator* (URO) i una clara falta de mostreig per al receptor de progesterona (PR).

En aquest treball hem demostrat la importància d'establir un banc de proves equilibrat durant el desenvolupament de nous mètodes. Mitjançant l'ús d'un banc de proves divers hem pogut establir en quins casos es pot esperar que el protocol obtingui resultats acurats, i quines àrees necessiten més desenvolupament. El banc de proves ha consistit de quatre proteïnes i més de trenta lligands, molt més dels que comunament s'utilitzen en el desenvolupament de mètodes per a la predicció d'energies d'unió mitjançant mètodes basats en camins (*pathway-based*).

En resum, la metodologia desenvolupada durant aquesta tesi pot contribuir al procés de recerca de nous fàrmacs per a certs tipus de sistemes de proteïnes. Per a la resta, hem observat que els mètodes de simulació no esbiaixats no són prou eficients i tècniques més sofisticades són necessàries.

Bibliography

- Atilgan, A. R., Durell, S. R., Jernigan, R. L., Demirel, M. C., Keskin, O., and Bahar, I. (2001). Anisotropy of fluctuation dynamics of proteins with an elastic network model. *Biophysical Journal*, 80(1):505–515.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (1995). Gambling in a rigged casino: the adversarial multi-armed bandit problem. *Annual Symposium on Foundations of Computer Science - Proceedings*, 68(68):322–331.
- Banks, J. L., Beard, H. S., Cao, Y., Cho, A. E., Damm, W., Farid, R., Felts, A. K., Halgren, T. A., Mainz, D. T., Maple, J. R., Murphy, R., Philipp, D. M., Repasky, M. P., Zhang, L. Y., Berne, B. J., Friesner, R. A., Gallicchio, E., and Levy, R. M. (2005). Integrated Modeling Program, Applied Chemical Theory (IMPACT). *Journal of Computational Chemistry*, 26(16):1752–1780.
- Barducci, A., Bussi, G., and Parrinello, M. (2008). Well-tempered metadynamics: A smoothly converging and tunable free-energy method. *Physical Review Letters*, 100(2):1–4.
- Beato, M. and Klug, J. (2000). Steroid hormone receptors: An update. *Human Reproduction Update*, 6(3):225–236.
- Berendsen, H. J. C., Postma, J. P. M., van Gunsteren, W. F., and Hermans, J. (1981). Interaction Models for Water in Relation to Protein Hydration. pages 331–342.
- Betz, R. M. and Dror, R. O. (2019). How Effectively Can Adaptive Sampling Methods Capture Spontaneous Ligand Binding? *Journal of Chemical Theory and Computation*, 15(3):2053–2063.
- Borrelli, K. W., Vitalis, A., Alcantara, R., and Guallar, V. (2005). PELE: Protein energy landscape exploration. A novel Monte Carlo based technique. *Journal of Chemical Theory and Computation*, 1(6):1304–1311.
- Bowman, G. R., Ensign, D. L., and Pande, V. S. (2010). Enhanced modeling via network theory: Adaptive sampling of markov state models. *Journal of Chemical Theory and Computation*, 6(3):787–794.
- Buch, I., Giorgino, T., and De Fabritiis, G. (2011). Complete reconstruction of an enzyme-inhibitor binding process by molecular dynamics simulations. *Proceedings of the National Academy of Sciences of the United States of America*, 108(25):10184–10189.

- Buch, I., Harvey, M. J., Giorgino, T., Anderson, D. P., and De Fabritiis, G. (2010). High-throughput all-atom molecular dynamics simulations using distributed computing. *Journal of Chemical Information and Modeling*, 50(3):397–403.
- Chen, W., Deng, Y., Russell, E., Wu, Y., Abel, R., and Wang, L. (2018). Accurate Calculation of Relative Binding Free Energies between Ligands with Different Net Charges. *Journal of Chemical Theory and Computation*, 14(12):6346–6358.
- 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., and Gustafsson, D. (2014). Discovery of the fibrinolysis inhibitor AZD6564, acting via interference of a protein-protein interaction. *ACS Medicinal Chemistry Letters*, 5(5):538–543.
- Chiavazzo, E., Covino, R., Coifman, R. R., Gear, C. W., Georgiou, A. S., Hummer, G., and Kevrekidis, I. G. (2017). Intrinsic map dynamics exploration for uncharted effective free-energy landscapes. *Proceedings of the National Academy of Sciences of the United States of America*, 114(28):E5494–E5503.
- Chodera, J. D. and Noé, F. (2014). Markov state models of biomolecular conformational dynamics. *Current Opinion in Structural Biology*, 25:135–144.
- Ciociola, A. A., Cohen, L. B., and Kulkarni, P. (2014). How drugs are developed and approved by the FDA: Current process and future directions. *American Journal of Gastroenterology*, 109(5):620–623.
- Conzen, S. D. (2008). Nuclear receptors and breast cancer. *Molecular Endocrinology*, 22(10):2215–2228.
- Csermely, P., Palotai, R., and Nussinov, R. (2010). Induced fit, conformational selection and independent dynamic segments: An extended view of binding events. *Trends in Biochemical Sciences*, 35(10):539–546.
- Deng, N., Cui, D., Zhang, B. W., Xia, J., Cruz, J., and Levy, R. (2018). Comparing alchemical and physical pathway methods for computing the absolute binding free energy of charged ligands. *Physical Chemistry Chemical Physics*, 20(25):17081–17092.
- Dickson, A. (2018). Mapping the Ligand Binding Landscape. *Biophysical Journal*, 115(9):1707–1719.
- Dickson, A. and Brooks, C. L. (2014). WExplore: Hierarchical exploration of high-dimensional spaces using the weighted ensemble algorithm. *Journal of Physical Chemistry B*, 118(13):3532–3542.
- Dickson, A. and Lotz, S. D. (2016). Ligand Release Pathways Obtained with WExplore: Residence Times and Mechanisms. *Journal of Physical Chemistry B*, 120(24):5377–5385.
- Dickson, A. and Lotz, S. D. (2017). Multiple Ligand Unbinding Pathways and Ligand-Induced Destabilization Revealed by WExplore. *Biophysical Journal*, 112(4):620–629.

- Dobson, C. M. (2004). Chemical space and biology. *Nature*, 432(7019):824–828.
- Doerr, S. and De Fabritiis, G. (2014). On-the-fly learning and sampling of ligand binding by high-throughput molecular simulations. *Journal of Chemical Theory and Computation*, 10(5):2064–2069.
- Doudou, S., Burton, N. A., and Henchman, R. H. (2009). Standard free energy of binding from a one-dimensional potential of mean force. *Journal of Chemical Theory and Computation*, 5(4):909–918.
- Dougherty, T. J. and Pucci, M. J. (2014). *Antibiotic discovery and development*.
- 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., and Carlson, H. A. (2013). CSAR data set release 2012: Ligands, affinities, complexes, and docking decoys. *Journal of Chemical Information and Modeling*, 53(8):1842–1852.
- 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., and Pande, V. S. (2017). OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS Computational Biology*, 13(7):1–17.
- Eckhardt, R. (1987). Stan Ulam, John von Neumann, and the Monte Carlo method. *Los Alamos Science*, 15(131-136):30.
- Edman, K., Hosseini, A., Bjursell, M. K., Aagaard, A., Wissler, L., Gunnarsson, A., Kaminski, T., Köhler, C., Bäckström, S., Jensen, T. J., Cavallin, A., Karlsson, U., Nilsson, E., Lecina, D., Takahashi, R., Grebner, C., Geschwindner, S., Lepistö, M., Hogner, A. C., and Guallar, V. (2015). Ligand Binding Mechanism in Steroid Receptors: From Conserved Plasticity to Differential Evolutionary Constraints. *Structure*, 23(12):2280–2290.
- Fischer, E. (1894). Einfluss der Configuration auf die Wirkung der Enzyme. *Berichte der deutschen chemischen Gesellschaft*, 27(3):2985–2993.
- Frauenfelder, H., Sligar, S. G., and Wolynes, P. G. (1991). The energy landscapes and motions of proteins. *Science*, 254(5038):1598–1603.
- Gallicchio, E. and Levy, R. M. (2011). Recent theoretical and computational advances for modeling protein-ligand binding affinities. *Advances in Protein Chemistry and Structural Biology*, 85(24):27–80.
- Gapsys, V., Pérez-Benito, L., Aldeghi, M., Seeliger, D., Van Vlijmen, H., Tresadern, G., and De Groot, B. L. (2020). Large scale relative protein ligand binding affinities using non-equilibrium alchemy. *Chemical Science*, 11(4):1140–1152.
- General, I. J. (2010). A note on the standard state’s binding free energy. *Journal of Chemical Theory and Computation*, 6(8):2520–2524.

- Genheden, S. and Ryde, U. (2015). The MM/PBSA and MM/GBSA methods to estimate ligand-binding affinities. *Expert Opinion on Drug Discovery*, 10(5):449–461.
- Gilabert, J. F., Gracia Carmona, O., Hogner, A., and Guallar, V. (2020). Combining Monte Carlo and molecular dynamics simulations for enhanced binding free energy estimation through Markov State Models. *Journal of Chemical Information and Modeling*, Under revision.
- Gilabert, J. F., Grebner, C., Soler, D., Lecina, D., Municoy, M., Gracia Carmona, O., Soliva, R., Packer, M. J., Hughes, S. J., Tyrchan, C., Hogner, A., and Guallar, V. (2019). PELE-MSM: A monte carlo based protocol for the estimation of absolute binding free energies. *Journal of Chemical Theory and Computation*, (15):6243–6253.
- Gilabert, J. F., Lecina, D., Estrada, J., and Guallar, V. (2018). Monte Carlo Techniques for Drug Design: The Success Case of PELE. In *Biomolecular Simulations in Structure-Based Drug Discovery*, pages 87–103.
- Gilson, M. K. and Zhou, H. X. (2007). Calculation of protein-ligand binding affinities. *Annual Review of Biophysics and Biomolecular Structure*, 36(1):21–42.
- Gleeson, M. P. (2008). Generation of a set of simple, interpretable ADMET rules of thumb. *Journal of Medicinal Chemistry*, 51(4):817–834.
- Grebner, C., Iegre, J., Ulander, J., Edman, K., Hogner, A., and Tyrchan, C. (2016). Binding Mode and Induced Fit Predictions for Prospective Computational Drug Design. *Journal of Chemical Information and Modeling*, 56(4):774–787.
- Grebner, C., Lecina, D., Gil, V., Ulander, J., Hansson, P., Dellsen, A., Tyrchan, C., Edman, K., Hogner, A., and Guallar, V. (2017). Exploring Binding Mechanisms in Nuclear Hormone Receptors by Monte Carlo and X-ray-derived Motions. *Biophysical Journal*, 112(6):1147–1156.
- Grünberg, R., Leckner, J., and Nilges, M. (2004). Complementarity of structure ensembles in protein-protein binding. *Structure*, 12(12):2125–2136.
- Gumbart, J. C., Roux, B., and Chipot, C. (2013). Standard binding free energies from computer simulations: What is the best strategy? *Journal of Chemical Theory and Computation*, 9(1):794–802.
- Halgren, T. A., Murphy, R. B., Friesner, R. A., Beard, H. S., Frye, L. L., Pollard, W. T., and Banks, J. L. (2004). Glide: A New Approach for Rapid, Accurate Docking and Scoring. 2. Enrichment Factors in Database Screening. *Journal of Medicinal Chemistry*, 47(7):1750–1759.
- Hartigan, J. (1975). *Clustering algorithms*. Wiley, New York.
- Harvey, M. J., Giupponi, G., and De Fabritiis, G. (2009). ACEMD: Accelerating biomolecular dynamics in the microsecond time scale. *Journal of Chemical Theory and Computation*, 5(6):1632–1639.

- Hinrichs, N. S. and Pande, V. S. (2007). Calculation of the distribution of eigenvalues and eigenvectors in Markovian state models for molecular dynamics. *Journal of Chemical Physics*, 126(24).
- Huang, S. Y., Grinter, S. Z., and Zou, X. (2010). Scoring functions and their evaluation methods for protein-ligand docking: Recent advances and future directions. *Physical Chemistry Chemical Physics*, 12(40):12899–12908.
- Huang, S. Y. and Zou, X. (2007). Ensemble docking of multiple protein structures: Considering protein structural variations in molecular docking. *Proteins: Structure, Function and Genetics*, 66(2):399–421.
- Husic, B. E. and Pande, V. S. (2018). Markov State Models: From an Art to a Science. *Journal of the American Chemical Society*, 140(7):2386–2396.
- Izrailev, S., Stepaniants, S., Balsera, M., Oono, Y., and Schulten, K. (1997). Molecular dynamics study of unbinding of the avidin-biotin complex. *Biophysical Journal*, 72(4):1568–1581.
- Jacobson, M. P., Kaminski, G. A., Friesner, R. A., and Rapp, C. S. (2002). Force field validation using protein side chain prediction. *Journal of Physical Chemistry B*, 106(44):11673–11680.
- Jankun, J. and Skrzypczak-Jankun, E. (1999). Molecular basis of specific inhibition of urokinase plasminogen activator by amiloride. *Cancer Biochemistry Biophysics*, 17(1-2):109–123.
- Jarzynski, C. (1997). Nonequilibrium equality for free energy differences. *Physical Review Letters*, 78(14):2690–2693.
- Jorgensen, W. L. (2004). The Many Roles of Computation in Drug Discovery. *Science*, 303(5665):1813–1818.
- Jorgensen, W. L., Chandrasekhar, J., Madura, J. D., Impey, R. W., and Klein, M. L. (1983). Comparison of simple potential functions for simulating liquid water. *The Journal of Chemical Physics*, 79(2):926–935.
- Jorgensen, W. L. and Tirado-Rives, J. (1988). The OPLS Potential Functions for Proteins. Energy Minimizations for Crystals of Cyclic Peptides and Crambin. *Journal of the American Chemical Society*, 110(6):1657–1666.
- Jorgensen, W. L. and Tirado-Rives, J. (2005). Molecular modeling of organic and biomolecular systems using BOSS and MCPRO. *Journal of Computational Chemistry*, 26(16):1689–1700.
- Kaminski, G. A., Friesner, R. A., Tirado-Rives, J., and Jorgensen, W. L. (2001). Evaluation and reparametrization of the OPLS-AA force field for proteins via comparison with accurate quantum chemical calculations on peptides. *Journal of Physical Chemistry B*, 105(28):6474–6487.

- Kirkwood, J. G. (1935). Statistical Mechanics of Fluid Mixtures. *Journal of Chemical Physics*, 3(1933):300–313.
- Koshland, D. E. (1958). Application of a Theory of Enzyme Specificity to Protein Synthesis. *Proceedings of the National Academy of Sciences*, 44(2):98–104.
- Kutzner, C., Páll, S., Fechner, M., Esztermann, A., De Groot, B. L., and Grubmüller, H. (2015). Best bang for your buck: GPU nodes for GROMACS biomolecular simulations. *Journal of Computational Chemistry*, 36(26):1990–2008.
- Laio, A. and Parrinello, M. (2002). Escaping Free-Energy Minima. *Proc. Natl. Acad. Sci. U. S. A.*, 99:12562.
- Lamim Ribeiro, J. M. and Tiwary, P. (2019). Toward Achieving Efficient and Accurate Ligand-Protein Unbinding with Deep Learning and Molecular Dynamics through RAVE. *Journal of Chemical Theory and Computation*, 15(1):708–719.
- Lang, P. T., Brozell, S. R., Mukherjee, S., Pettersen, E. F., Meng, E. C., Thomas, V., Rizzo, R. C., Case, D. A., James, T. L., and Kuntz, I. D. (2009). DOCK 6: Combining techniques to model RNA-small molecule complexes. *Rna*, 15(6):1219–1230.
- Lecina, D. (2017). Studying protein-ligand interactions using a Monte Carlo procedure. *Universitat de Barcelona*, (Doctoral thesis).
- Lecina, D., Gilabert, J. F., and Guallar, V. (2017). Adaptive simulations, towards interactive protein-ligand modeling. *Scientific Reports*, 7(1):8466.
- Lee, T. S., Cerutti, D. S., Mermelstein, D., Lin, C., Legrand, S., Giese, T. J., Roitberg, A., Case, D. A., Walker, R. C., and York, D. M. (2018). GPU-Accelerated Molecular Dynamics and Free Energy Methods in Amber18: Performance Enhancements and New Features. *Journal of Chemical Information and Modeling*, 58(10):2043–2050.
- Li, Z. and Scheraga, H. A. (1987). Monte Carlo-minimization approach to the multiple-minima problem in protein folding. *Chemistry*, 84(October):6611–6615.
- Limongelli, V. (2020). Ligand binding free energy and kinetics calculation in 2020. *WIREs Comput Mol Sci*, (August 2019):1–32.
- Limongelli, V., Bonomi, M., and Parrinello, M. (2013). Funnel metadynamics as accurate binding free-energy method. *Proceedings of the National Academy of Sciences*, 110(16):6358–6363.
- Lotz, S. D. and Dickson, A. (2018). Unbiased Molecular Dynamics of 11 min Timescale Drug Unbinding Reveals Transition State Stabilizing Interactions. *Journal of the American Chemical Society*, 140(2):618–628.
- Maier, J. A., Martinez, C., Kasavajhala, K., Wickstrom, L., Hauser, K. E., and Simmerling, C. (2015). ff14SB: Improving the Accuracy of Protein Side Chain and Backbone Parameters from ff99SB. *Journal of Chemical Theory and Computation*, 11(8):3696–3713.

- Malek, R. and Mousseau, N. (2000). Dynamics of Lennard-Jones clusters: A characterization of the activation-relaxation technique. *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 62(6):7723–7728.
- Manousiouthakis, V. I. and Deem, M. W. (1999). Strict detailed balance is unnecessary in Monte Carlo simulation. *The Journal of Chemical Physics*, 110(6):2753–2756.
- Mardt, A., Pasquali, L., Wu, H., and Noé, F. (2018). VAMPnets for deep learning of molecular kinetics. *Nature Communications*, 9(1):1–14.
- Mares-Guia, M. and Shaw, E. (1965). Studies on the Active Center of Trypsin. the Binding of Amidines and. *The Journal of biological chemistry*, 240:1579–1585.
- Matthews, H., Ranson, M., and Kelso, M. J. (2011). Anti-tumour/metastasis effects of the potassium-sparing diuretic amiloride: An orally active anti-cancer drug waiting for its call-of-duty? *International Journal of Cancer*, 129(9):2051–2061.
- McGibbon, R. T. and Pande, V. S. (2015). Variational cross-validation of slow dynamical modes in molecular kinetics. *Journal of Chemical Physics*, 142(12).
- Metropolis, N. and Ulam, S. (1949). The Monte Carlo method. *Journal of the American Statistical Association*, 44(247):335–341.
- Mobley, D. L. (2012). Let’s get honest about sampling. *Journal of Computer-Aided Molecular Design*, 26(1):93–95.
- Mobley, D. L. and Gilson, M. K. (2017). Predicting Binding Free Energies: Frontiers and Benchmarks. *Annual Review of Biophysics*, 46(1):531–558.
- Mu, Y., Nguyen, P. H., and Stock, G. (2005). Energy landscape of a small peptide revealed by dihedral angle principal component analysis. *Proteins: Structure, Function and Genetics*, 58(1):45–52.
- Naritomi, Y. and Fuchigami, S. (2011). Slow dynamics in protein fluctuations revealed by time-structure based independent component analysis: The case of domain motions. *Journal of Chemical Physics*, 134(6).
- Noé, F., Olsson, S., Köhler, J., and Wu, H. (2019). Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457).
- Nüske, F., Keller, B. G., Pe, G., Mey, A. S. J. S., and Noé, F. (2014). Variational Approach to Molecular Kinetics. *J. Chem. Theory Comput.*, 10:1739–1752.
- Onufriev, A., Bashford, D., and Case, D. A. (2004). Exploring Protein Native States and Large-Scale Conformational Changes with a Modified Generalized Born Model. *Proteins: Structure, Function and Genetics*, 55(2):383–394.
- Owaki, H., Makar, R., Boulton, T. G., Cobb, M. H., and Geppert, T. D. (1992). Extracellular signal-regulated kinases in T cells: Characterization of human ERK1 and ERK2 cDNAs. *Biochemical and Biophysical Research Communications*, 182(3):1416–1422.

- Pan, A. C., Xu, H., Palpant, T., and Shaw, D. E. (2017). Quantitative Characterization of the Binding and Unbinding of Millimolar Drug Fragments with Molecular Dynamics Simulations. *Journal of Chemical Theory and Computation*, 13(7):3372–3377.
- Pande, V. S., Beauchamp, K., and Bowman, G. R. (2010). Everything you wanted to know about Markov State Models but were afraid to ask. *Methods*, 52(1):99–105.
- Patrick Walters, W., Stahl, M. T., and Murcko, M. A. (1998). Virtual screening - An overview. *Drug Discovery Today*, 3(4):160–178.
- Pérez, A., Marchán, I., Svozil, D., Sponer, J., Cheatham, T. E., Laughton, C. A., and Orozco, M. (2007). Refinement of the AMBER force field for nucleic acids: Improving the description of α/γ conformers. *Biophysical Journal*, 92(11):3817–3829.
- Pérez-Hernández, G., Paul, F., Giorgino, T., De Fabritiis, G., and Noé, F. (2013). Identification of slow molecular order parameters for Markov model construction. *Journal of Chemical Physics*, 139(1).
- Plattner, N. and Noé, F. (2015). Protein conformational plasticity and complex ligand-binding kinetics explored by atomistic simulations and Markov models. *Nature Communications*, 6(May).
- Ponomarenko, E. A., Poverennaya, E. V., Ilgisonis, E. V., Pyatnitskiy, M. A., Kopylov, A. T., Zgoda, V. G., Lisitsa, A. V., and Archakov, A. I. (2016). The Size of the Human Proteome: The Width and Depth. *International Journal of Analytical Chemistry*, 2016.
- Pramanik, D., Smith, Z., Kells, A., and Tiwary, P. (2019). Can One Trust Kinetic and Thermodynamic Observables from Biased Metadynamics Simulations?: Detailed Quantitative Benchmarks on Millimolar Drug Fragment Dissociation. *Journal of Physical Chemistry B*, 123(17):3672–3678.
- Prinz, J. H., Wu, H., Sarich, M., Keller, B., Senne, M., Held, M., Chodera, J. D., Schütte, C., and Noé, F. (2011). Markov models of molecular kinetics: Generation and validation. *Journal of Chemical Physics*, 134(17):1–23.
- Rawlings, N. D. and Barrett, A. J. (1994). Families of serine peptidases. *Methods in Enzymology*, 244(C):19–61.
- Reif, M. M., Hünenberger, P. H., and Oostenbrink, C. (2012). New interaction parameters for charged amino acid side chains in the GROMOS force field. *Journal of Chemical Theory and Computation*, 8(10):3705–3723.
- Ribeiro, J. M. L., Bravo, P., Wang, Y., and Tiwary, P. (2018). Reweighted autoencoded variational Bayes for enhanced sampling (RAVE). *Journal of Chemical Physics*, 149(7).
- Robbins, H. (1952). Some Aspects of the Sequential Design of Experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535.

- Robin Ganellin, C., Jefferis, R., and Roberts, S. M. (2013). *Introduction to biological and small molecule drug research and development: Theory and case studies*.
- Ruiz-Carmona, S., Alvarez-Garcia, D., Foloppe, N., Garmendia-Doval, A. B., Juhos, S., Schmidtke, P., Barril, X., Hubbard, R. E., and Morley, S. D. (2014). rDock: A Fast, Versatile and Open Source Program for Docking Ligands to Proteins and Nucleic Acids. *PLoS Computational Biology*, 10(4).
- Ryckaert, J. P., Ciccotti, G., and Berendsen, H. J. (1977). Numerical integration of the cartesian equations of motion of a system with constraints: molecular dynamics of n-alkanes. *Journal of Computational Physics*, 23(3):327–341.
- Santiago, G., De Salas, F., Lucas, M. F., Monza, E., Acebes, S., Martinez, Á. T., Camarero, S., and Guallar, V. (2016). Computer-Aided Laccase Engineering: Toward Biological Oxidation of Arylamines. *ACS Catalysis*, 6(8):5415–5423.
- Scherer, M. K., Trendelkamp-Schroer, B., Paul, F., Pérez-Hernández, G., Hoffmann, M., Plattner, N., Wehmeyer, C., Prinz, J. H., and Noé, F. (2015). PyEMMA 2: A Software Package for Estimation, Validation, and Analysis of Markov Models. *Journal of Chemical Theory and Computation*, 11(11):5525–5542.
- Schmidt, T. C., Eriksson, P. O., Gustafsson, D., Cosgrove, D., Frølund, B., and Boström, J. (2017). Discovery and Evaluation of Anti-Fibrinolytic Plasmin Inhibitors Derived from 5-(4-Piperidyl)isoxazol-3-ol (4-PIOL). *Journal of Chemical Information and Modeling*, 57(7):1703–1714.
- Schütte, C., Noé, F., Lu, J., Sarich, M., and Vanden-Eijnden, E. (2011). Markov state models based on milestoning. *Journal of Chemical Physics*, 134(20).
- Schwantes, C. R. and Pande, V. S. (2015). Modeling molecular kinetics with tICA and the kernel trick. *Journal of Chemical Theory and Computation*, 11(2):600–608.
- Sertkaya, A., Wong, H. H., Jessup, A., and Beleche, T. (2016). Key cost drivers of pharmaceutical clinical trials in the United States. *Clinical Trials*, 13(2):117–126.
- Shamsi, Z., Cheng, K. J., and Shukla, D. (2018). Reinforcement Learning Based Adaptive Sampling: REAPing Rewards by Exploring Protein Conformational Landscapes. *Journal of Physical Chemistry B*, 122(35):8386–8395.
- Shaw, B. D. E., Deneroff, M. M., Dror, R. O., Kuskin, J. S., Larson, R. H., Salmon, J. K., Young, C., Batson, B., Bowers, K. J., Chao, J. C., Eastwood, M. P., Gagliardo, J., Grossman, J. P., Ho, C. R., Ierardi, D. J., Kolossváry, I., Klepeis, J. L., Layman, T., Mcleavey, C., Moraes, M. A., Mueller, R., Priest, E. C., Shan, Y., Spengler, J., Theobald, M., Towles, B., and Wang, S. C. (2008). Anton, a Special-Purpose Machine for Molecular Dynamics Simulation. *Communications of the ACM*, pages 91–97.
- 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., and Wriggers, W.

- (2010). Atomic-level characterization of the structural dynamics of proteins. *Science*, 330(6002):341–346.
- Shirts, M. and Pande, V. S. (2000). Screen savers of the world unite. *Science*, 290(5498):1903–1904.
- Singhal, N. and Pande, V. S. (2005). Error analysis and efficient sampling in Markovian state models for molecular dynamics. *Journal of Chemical Physics*, 123(20).
- Smith, A. G. and Muscat, G. E. (2005). Skeletal muscle and nuclear hormone receptors: Implications for cardiovascular and metabolic disease. *International Journal of Biochemistry and Cell Biology*, 37(10 SPEC. ISS.):2047–2063.
- Susukita, R., Ebisuzaki, T., Elmgreen, B. G., Furusawa, H., Kato, K., Kawai, A., Kobayashi, Y., Koishi, T., McNiven, G. D., Narumi, T., and Yasuoka, K. (2003). Hardware accelerator for molecular dynamics: MDGRAPE-2. *Computer Physics Communications*, 155(2):115–131.
- Sutton, R. and Barto, A. (1998). Reinforcement Learning: An Introduction. *IEEE Transactions on Neural Networks*, 9(5):1054–1054.
- Swanson, J. M., Henchman, R. H., and McCammon, J. A. (2004). Revisiting Free Energy Calculations: A Theoretical Connection to MM/PBSA and Direct Calculation of the Association Free Energy. *Biophysical Journal*, 86(1 I):67–74.
- Swope, W. C., Pitera, J. W., and Suits, F. (2004). Describing Protein Folding Kinetics by Molecular Dynamics Simulations. 1. Theory [†]. *The Journal of Physical Chemistry B*, 108(21):6571–6581.
- Takahashi, R., Gil, V. A., and Guallar, V. (2014). Monte carlo free ligand diffusion with markov state model analysis and absolute binding free energy calculations. *Journal of Chemical Theory and Computation*, 10(1):282–288.
- Tang, L. and Han, X. (2013). The urokinase plasminogen activator system in breast cancer invasion and metastasis. *Biomedicine and Pharmacotherapy*, 67(2):179–182.
- Tiwary, P. and Parrinello, M. (2013). From metadynamics to dynamics. *Physical Review Letters*, 111(23):1–5.
- Torrie, G. M. and Valleau, J. P. (1977). Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling. *Journal of Computational Physics*, 23(2):187–199.
- Trott, O. and Olson, A. J. (2010). Software news and update AutoDock Vina: Improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. *Journal of Computational Chemistry*, 31(2):455–461.
- Troussicot, L., Guillièrè, F., Limongelli, V., Walker, O., and Lancelin, J. M. (2015). Funnel-metadynamics and solution NMR to estimate protein-ligand affinities. *Journal of the American Chemical Society*, 137(3):1273–1281.

- Tsai, C.-J., Kumar, S., Ma, B., and Nussinov, R. (1999). Folding funnels, binding funnels, and protein function. *Protein Science*, 8(6):1181–1190.
- Vanommeslaeghe, K., Hatcher, E., Acharya, C., Kundu, S., Zhong, S., Shim, J., Darian, E., Guvench, O., Lopes, P., Vorobyov, I., and Mackerell, A. D. (2010). CHARMM general force field: A force field for drug-like molecules compatible with the CHARMM all-atom additive biological force fields. *Journal of Computational Chemistry*, 31(4):671–690.
- Wales, D. J. (2013). Surveying a complex potential energy landscape: Overcoming broken ergodicity using basin-sampling. *Chemical Physics Letters*, 584:1–9.
- Wales, D. J. and Doye, J. P. (1997). Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms. *Journal of Physical Chemistry A*, 101(28):5111–5116.
- Wang, J., Deng, Y., and Roux, B. (2006a). Absolute binding free energy calculations using molecular dynamics simulations with restraining potentials. *Biophysical Journal*, 91(8):2798–2814.
- Wang, J., Wang, W., Kollman, P. A., and Case, D. A. (2006b). Automatic atom type and bond type perception in molecular mechanical calculations. *Journal of Molecular Graphics and Modelling*.
- Wang, J., Wolf, R. M., Caldwell, J. W., Kollman, P. A., and Case, D. A. (2004). Development and testing of a general Amber force field. *Journal of Computational Chemistry*.
- 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., and Abel, R. (2015). 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. *Journal of the American Chemical Society*, 137(7):2695–2703.
- Wang, W., Wang, J., and Kollman, P. A. (1999). What determines the van der Waals coefficient β in the LIE (linear interaction energy) method to estimate binding free energies using molecular dynamics simulations? *Proteins: Structure, Function and Genetics*, 34(3):395–402.
- Wang, X. X., Jiang, T., and Levi, M. (2010). Nuclear hormone receptors in diabetic nephropathy. *Nature Reviews Nephrology*, 6(6):342–351.
- Wlodarski, T. and Zagrovic, B. (2009). Conformational selection and induced fit mechanism underlie specificity in noncovalent interactions with ubiquitin. *Proceedings of the National Academy of Sciences of the United States of America*, 106(46):19346–19351.

- Woo, H. J. and Roux, B. (2005). Calculation of absolute protein-ligand binding free energy from computer simulations. *Proceedings of the National Academy of Sciences of the United States of America*, 102(19):6825–6830.
- Woods, C. J., Michel, J., Bodnarchuk, M., Genheden, S., Bradshaw, R., Ross, G. A., Cave-Ayland, C., Bruce-Macdonald, H., Cabedo Martinez, A. I., and Graham, J. (2017). ProtoMS 3.3.
- Zhou, H.-X. and Gilson, M. K. (2009). SI-Theory of free energy and entropy in noncovalent binding. *Chem. Rev.*, 109(9):4092–4107.
- Zhu, K., Shirts, M. R., and Friesner, R. A. (2007). 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. *Journal of Chemical Theory and Computation*, 3(6):2108–2119.
- Zimmerman, M. I., Porter, J. R., Sun, X., Silva, R. R., and Bowman, G. R. (2018). Choice of Adaptive Sampling Strategy Impacts State Discovery, Transition Probabilities, and the Apparent Mechanism of Conformational Changes. *Journal of Chemical Theory and Computation*, 14(11):5459–5475.
- Zuckerman, D. M. (2013). *Statistical Physics of Biomolecules: An Introduction*, volume 2012.
- Zuckerman, D. M. and Chong, L. T. (2017). Weighted Ensemble Simulation: Review of Methodology, Applications, and Software. *Annual Review of Biophysics*, 46(1):43–57.
- Zwanzig, R. (1983). From Classical Dynamics to Continuous Time Random Walks. *J. Stat. Phys.*, 30(2):255–262.
- Zwanzig, R. W. (1954). High-Temperature Equation of State by a Perturbation Method. I. Nonpolar Gases. *Journal of Chemical Physics*, 22(8):1420–1426.
- Zwier, M. C., Adelman, J. L., Kaus, J. W., Pratt, A. J., Wong, K. F., Rego, N. B., Suárez, E., Lettieri, S., Wang, D. W., Grabe, M., Zuckerman, D. M., and Chong, L. T. (2015). WESTPA: An interoperable, highly scalable software package for weighted ensemble simulation and analysis. *Journal of Chemical Theory and Computation*, 11(2):800–809.