Prediction of Allosteric Binding Modes in Homologous NMDA-Receptors

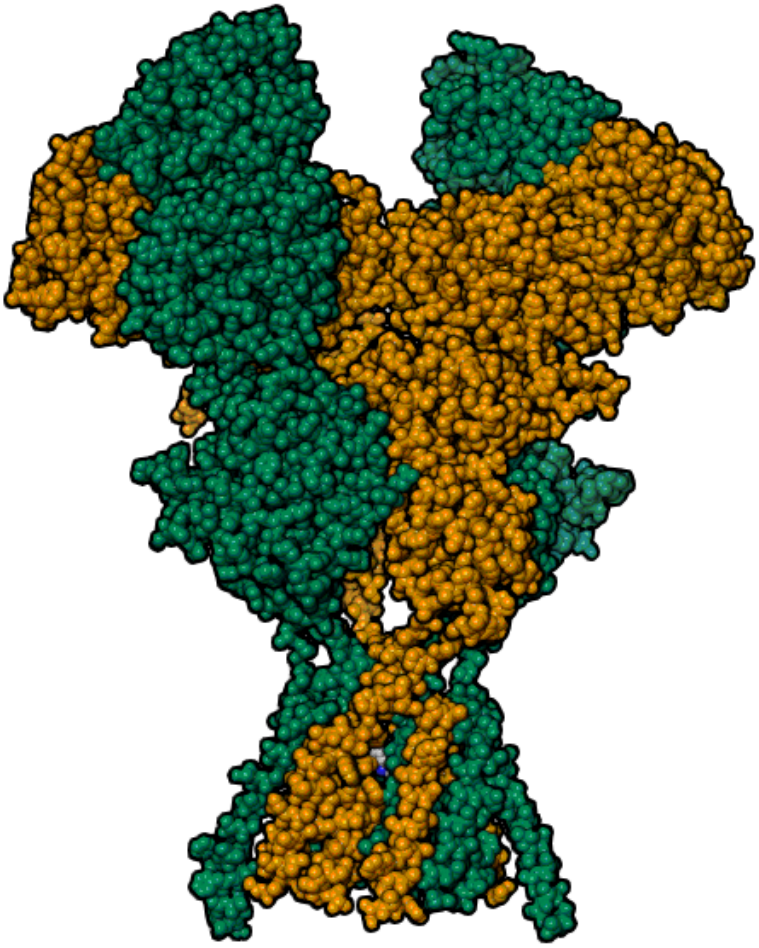

By: *Anton Sellerberg*
Supervisor*: Pär Söderhjelm*
Examiner: *Mikael Akke*

Division of biophysical chemistry
Department of Chemistry
Lund University

September 2nd, 2024

Abstract

The NMDA receptors are natural glutamate binders that play critical roles in the function of the central nervous system, with their dysfunction implicated in serious psychiatric disorders and neurodegenerative conditions. Understanding their roles in the brain and developing treatments will benefit from the existence of small molecules specifically targeting different NMDA receptor subtypes. This study has investigated allosteric binding sites of two homologous NMDA receptors. Agonist binding domains in GluN1/N2A (crystal structure) and GluN1/N2B (homology model) were initially subjected to Induced Fit Docking of a group of 15 analogous ligands. Ten of these are known binders of at least one receptor. Five ligands are newly synthesised analogues of known binders. Following Induced Fit Docking 23 ligand-receptor complexes of type GluN1/N2A and 20 ligand-receptor complexes of GluN1/N2B were investigated with parallel molecular dynamics (100ns) and triplicate metadynamics (20ns). Results suggest stable binding modes for 11 ligands in GluN1/N2A and 8 ligands in GluN1/N2B.

Table of Contents

1. Introduction

Some vital aspects of biological systems such as ourselves cannot be observed directly. Many relevant phenomena occur on scales that are either so small, or so large that they cannot be assessed through experimentation. In trying to understand the intricacies of molecular interactions or the stability of an entire ecosystem, one often needs to turn to simulations. These controlled environments are free from the restraints of physical experimentation. Instead they force us to grapple with the limitations of experimental models, and the limited capabilities of our computers.

Drawing meaningful conclusions from simulated biological systems remains challenging. Difficulty often arises because our understanding of many biological processes is incomplete, and computer hardware lacks the power to model biological complexity in complete detail. There are nevertheless ways of bridging biology and simulation in a way that advances both.

One way of advancing our understanding of a biological system, and approaching medical interventions in neurodegenerative disease is by making the biochemical distinction between related subgroups of human glutamate receptors. To functionally tell these structures apart is a first step towards studying their roles in the brain. The primary excitatory neurotransmitter of the mammalian brain is found in the form of L-glutamate (Glu). Its functional diversity in the brain is due to the existence of two families of receptors: the *Ionotropic Glutamate Receptors* (iGluRs) and their counterpart, the *Metabotropic Glutamate Receptors* (mGluRs). These function in tandem to impact the central nervous system over complex timescales and locations (Reiner and Levitz, 2018). The normal activity of glutamate receptors in the brain has a central role in synaptic plasticity and the encoding of memory (Morris, 2013), while dysfunction of Glu-signalling has been linked to neurodegenerative conditions like Alzheimer's and Huntington's disease (Lewerenz and Maher, 2015). It is known to play a role in the development of psychiatric disorders like schizophrenia, bipolar- and major depressive disorder (Li, Yang and Lin, 2019).

This work approaches the problem of drug discovery within two subtypes of iGluRs, via computational methods designed to evaluate and specify their interactions with a set of small molecules, most of which are known to interact with one or both of these receptors.

On the way towards describing, and medically targeting the differences between closely related protein complexes, a high quality crystal structure paired with molecular simulations is a helpful starting point. The study of *molecular dynamics* (MD) centres around computational models of events at microscopic scales. These simulate interactions between atoms and calculate the motion that follows. By the now long held insight of Newtonian mechanics, that *Force equals mass times acceleration,* it is possible to relate the forces acting on every atomic particle to their

future motion. Given a close enough approximation of these forces, an evolution of a theoretical system of particles offers insight into the properties of a real one. Forces between particles in the system are governed by *force fields*, vital ingredients to these dynamic simulations. These mathematical models describe things like Van der Waals-forces, electrostatic interactions and the flexibility of intramolecular bonds, where the force field in question must be sufficiently accurate to give biologically useful results (Lindorff-Larsen et al., 2012).

This offers a detailed model of molecular behaviour, as the *dynamics*, or time-dependent behaviour of the system, evolves under the influence of classical mechanics. As Newton's equations of motion are resolved one step at a time, it becomes possible to track the transitions between molecular states. Changes in positions, velocities and orientations lend insight into real-time processes like chemical reactions and molecular binding events (Hollingsworth and Dror, 2018).

There is a great potential in advanced computational tools like molecular dynamics and its related methods. As available computing power continues to grow and be more efficiently utilised (Götz et al., 2012), (Stone et al., 2016), one could expect to see methods like MD spread into a wider range of disciplines. Aside from the growing power of simulation, another reason for new opportunities is the increasing availability of crystal structures in neuroscience. Such protein structures, integral to the processes taking place in the brain, present possible targets for drug discovery (Hilger, Masureel and Kobilka, 2018), (G. Brent Dawe et al., 2016).

The thesis presented here is organised as such: Section 2 presents theoretical background information on a number of relevant topics, beginning with the subgroup of iGluRs known as NMDA receptors. This biological background is followed by descriptions of computational methods and models that have been relevant to this work, with the section ending in an argument for the chosen method. Section 3 covers the specifics of how this work has been carried out. In section 4 the results are first presented and discussed, followed by a general discussion of the work as a whole, and some suggestions for future research.

2. Background

This section presents theoretical backgrounds on the topics and methods most relevant to this thesis.

2.1 The NMDA receptor complex

The family of cation-permeable, ligand-gated, ionotropic glutamate receptors has three functional classes. These are paired with a fourth, mysterious group known as GluD1-2, which doesn't behave as an ion channel under normal circumstances, in spite of its structural similarity to the AMPA and Kainate receptors. The three classes of ionotropic glutamate receptors are: AMPA-, Kainate-, and NMDA receptors (Hansen et al., 2018).

These were at first identified (and named) pharmacologically, by their respective agonist molecules. The NMDA receptors are distinguished by a voltage dependent inhibition from extracellular Mg^{2+}, and a high permeability of Ca^{2+}. They require the binding of *two co-agonists*: *glutamate* and *glycine* (or D-serine) for the activation of the ion channel (Hansen et al., 2018).

One big reason for paying attention to the NMDA receptor is that it is dependent on both pre- and postsynaptic activity. This *coincidence function* has been linked to the encoding of information in the brain. The requirement of two simultaneous events for receptor activation allows them to act as sensors of *synchronous neural activity*, allowing for neurons that fire together, to wire together (Morris, 2013).

There are seven genes that encode the seven subunits of the NMDA receptors. All the receptors we know of are *heterotetrameric* assemblies of those, meaning they consist of four non-identical building blocks (see figure 1). All building blocks belong to three types: *GluN1*, *GluN2*(A to D), and *GluN3*(A & B). GluN1 is the glycine binder, obligatory to all NMDARs. GluN2, with its subtypes A through D are binders of glutamate. The functional differences of its subtypes give the central nervous system a means of controlling properties of NMDARs, depending on developmental period and brain region. It is the developmental aspect of the GluN2-subunits that has driven much of the research towards finding small molecules that can distinguish between them. Assemblies and functions of GluN3-types are less certain, and beyond the scope of this work. (Hansen et al., 2018).

In terms of structure and function, all subunits share a four-domain architecture. Three of these domains can be seen in figure 1. The subunit architecture consists of an extracellular amino-terminal domain (ATD), a bilobed agonist-binding domain (ABD), a pore-forming transmembrane domain (TBD) and an intracellular carboxy-terminal domain (Hansen et al., 2018).

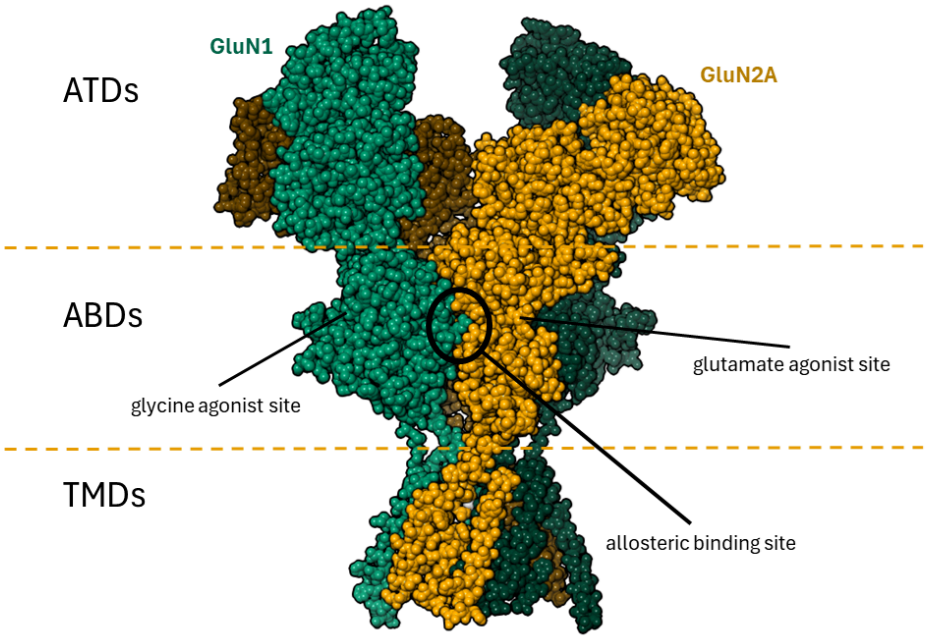

Figure 1: *Crystal structure of the 4 subunits of the human GluN1/N2A NMDA-receptor. Dashed lines separate between functional domains. Carboxy-terminal domains are not present in this structure. N1-subunits are depicted in green, N2A-subunits are depicted in yellow.*

Importantly there is evidence that the structural integrity and characteristics of the ABDs are retained in isolation, meaning that a crystal structure containing only the binding domains can mimic the behaviour of the ABDs in the full receptor complex (Armstrong and Gouaux, 2000). This allows in particular for studying the contact area between GluN1 and GluN2, where the presence of different subunits modulate the function of the receptor. This interface contains a binding site (see figure 1) for both positive and negative allosteric (long-range) modulators with selectivity for N1/N2A, making it an important location for interacting with NMDARs (Hansen et al., 2018).

This subtype-selective binding site is the main attraction in the current work, and the experimental underpinnings of the inquiry have been provided by a high-resolution (1.7 Å) crystal structure of the ABDs present in the GluN1/N2A-complex. This crystal structure is visualised in figure 2. Figure 2 aims to provide an overview of the allosteric binding site and selected residues within it. These residues will be discussed further in the results.

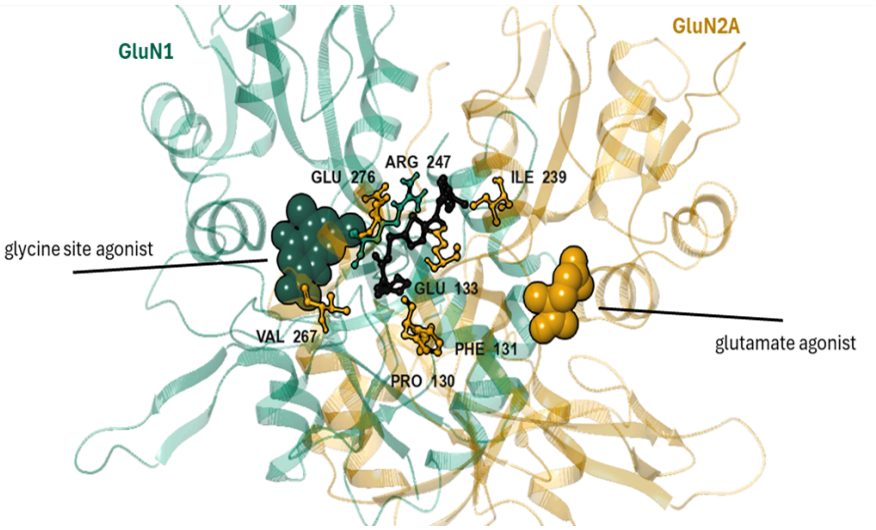

Figure 2: *Crystal structure of agonist binding domains of GluN1/N2A, co-crystallised with ligand AA33 (coloured black). Selected residues relevant to ligand binding and homology modelling have been marked out.*

This work concerns itself primarily with two of the four subcomplexes of N1/N2-receptors, which will be referred to as GluN1A (or subcomplex A, seen in figure 2), and GluN1B (or subcomplex B), in which N2A is exchanged with N2B.

2.2 Molecular Dynamics

On the subject of simulation, one ought to wonder about their accuracy in relation to experiments in the lab. Can these simulated dynamics really capture the real physics and chemistry taking place in a biological system, like those of GluN1A and GluN1B?

More specifically: Are the results of molecular dynamics *really* statistically representative of the *true behaviour* of the system? Answering this demands a closer look at the underlying theory.

In what is referred to as classical molecular dynamics, a numerical solution to the *Hamiltonian* is calculated for each step of the simulation, which tends to be around 10^{-15} seconds. The Hamiltonian is the total energy of a system of particles, as found in the sum of their potential and kinetic energies. Kinetics are given by particle momenta, while potential energies come from interactions between particles, such as chemical bonds, van der Waals forces and electrostatics (Bussi and Branduardi, 2015).

In Hamilton's description of mechanics, the predetermined path of the system lies in the variables (q_i , p_i), positions and momenta. Given all the positions and all the momenta, it is possible to completely describe the time dependent evolution of a system (Fowler, 2014).

When the total energy of a system changes, that is the derivative of the Hamiltonian. For an incremental step, it is given as:

$$dH(p_i, q_i) = -\sum_i \dot{p_i}\, dq_i + \sum_i \dot{q_i}\, dp_i \tag{1}$$

Where $\dot{x}$ is the time-derivative of x, q_i holds all particle coordinates and p_i the momenta of those particles (Fowler, 2014).

Given a system of N_{AT} atoms with the coordinates q_i, momenta p_i and masses m_i, it is possible to define a potential energy as a function of all atomic positions: $U(q)$.

$$\dot{q_i} = \frac{p_i}{m_i} \tag{2a}$$

$$\dot{p_i} = -\frac{\partial U(q)}{\partial q_i} \tag{2b}$$

This potential energy function describes interatomic interactions, often defined in terms of experimentally determined force fields, which give reasonable approximations of $U(q)$ (Bussi and Branduardi, 2015).

With potential energy we get the total energy of the system as:

$$H(p, q) = \sum_i \frac{p_i^2}{2m_i} + U(q) \tag{3}$$

That total energy is conserved for an evolution according to equations (2a) and (2b). The system will only visit states that have the same total energy as the initial one, given that no energy has escaped the borders of the system (Bussi and Branduardi, 2015).

This however is not a very realistic scenario. In the case of biological systems, we know that they are constantly exchanging energy with their surroundings, while maintaining an approximate temperature.

Assuming a distribution of different values of total energy, we find that a given set of p's and q's will be visited by the system with a probability *P(p, q)*. This probability, in the case of a temperature controlled system, corresponds to the *canonical ensemble*

$$P(p,q) \propto \exp\left(-\frac{\sum_i \frac{p_i^2}{2m_i} + U(q)}{k_B T}\right) \tag{4}$$

Here k_B is Boltzmann's constant, which relates average kinetic energies of particles to a temperature T (Bussi and Branduardi, 2015). It is important at this point to discuss what is implied by equation (4), without getting caught up in its appearance.

We are relying on an assumption that at this point is difficult to ignore. If the simulation is going to assign correct statistical values to the different states of the system, then it must be *ergodic* in the timescale of that simulation. This means that the simulation must be long enough for it to construct the canonical ensemble of that system (Barducci, Bonomi and Parrinello, 2011). This is referred to as *the ergodic hypothesis*.

An ergodic system will pass through every point in phase space that is available to it given enough time. Under this assumption the statistical weights of time-averages become equivalent to averages across all possible states of the system. This can however not be proven for most systems, and even when it has been verified it often takes an unpractical amount of resources to reach this point. This has left a mark on MD in the introduction of *enhanced sampling protocols* (Bussi and Branduardi, 2015). These present workarounds to the ergodic hypothesis, and will be the subject of section 2.3.

Moving momentarily past this theoretical objection, the assumption of ergodicity and the introduction of the canonical ensemble leads to a different complexity. The *phase space* of a system (as defined by the variables q_i and p_i), where the canonical distribution is defined, contains all of the possible momenta and positions of every atom in the system. These two 3-dimensional properties give $3N_{AT} + 3N_{AT}$ dimensions to the system. When the dimensionality of the data is six times the number of particles, the canonical distribution is useless when one is trying to extract information about the thermodynamic properties of the system. Addressing this has led to the introduction of *collective variables* (CV's) to the analysis of MD simulations. This is a function of atomic coordinates, *s(q)* that is designed to capture an underlying process (Bussi and Branduardi, 2015).

For example: When trying to model the folding of a protein one might observe distances between residues as collective variables instead of looking at the totality of atomic coordinates (Vojtěch Spiwok, Kurečka and Aleš Křenek, 2022). Defining CVs can be quite difficult, since they are intended to reduce a complex situation to a lower dimensionality, while capturing something of interest. This will be revisited in greater detail in section 2.3.

Turning to the concept of Free Energy, which will be needed to evaluate collective variables *s(q)*, it is defined in the form of Helmholtz free energy as:

$$F = U - TS \tag{5}$$

With U in this equation being the *internal energy* of the system, T the temperature and S its entropy. The free energy here is generally understood as the amount of useful work that can be done by a closed thermodynamic system (Wikipedia, 2020). U is *not* the potential energy in this case, but the potential energy *is* contained within the internal energy. Free energy is a central concept in the interpretation of molecular dynamics. It presents a macroscopic perspective where the Hamiltonian gives insight at the mechanical, microscopic level.

In the context of analysing collective variables in MD, the role of potential energy *U(q)* is taken instead by the free energy *F(s)*. It has been defined so that the likelihood of observing a given value of *s, P(s)*, is given by:

$$P(s)ds \propto \exp\left(-\frac{F(s)}{k_B T}\right) ds \tag{6}$$

In terms of the given CV, the free energy *F(s)* provides insight into the thermodynamic landscape of the system. Given equation (6), we see that a lower free energy correlates with a higher probability of observation (Bussi and Branduardi, 2015). Note the similarity to equation (4).

Now that a given CV can be connected to the canonical ensemble, it is possible to draw up a free energy landscape of different states of the same CV. A theoretical such landscape is shown in figure 3. If the selection of CVs is well thought out, we can expect them to highlight relevant transitions in the system, like that between the hypothetical states A and B.

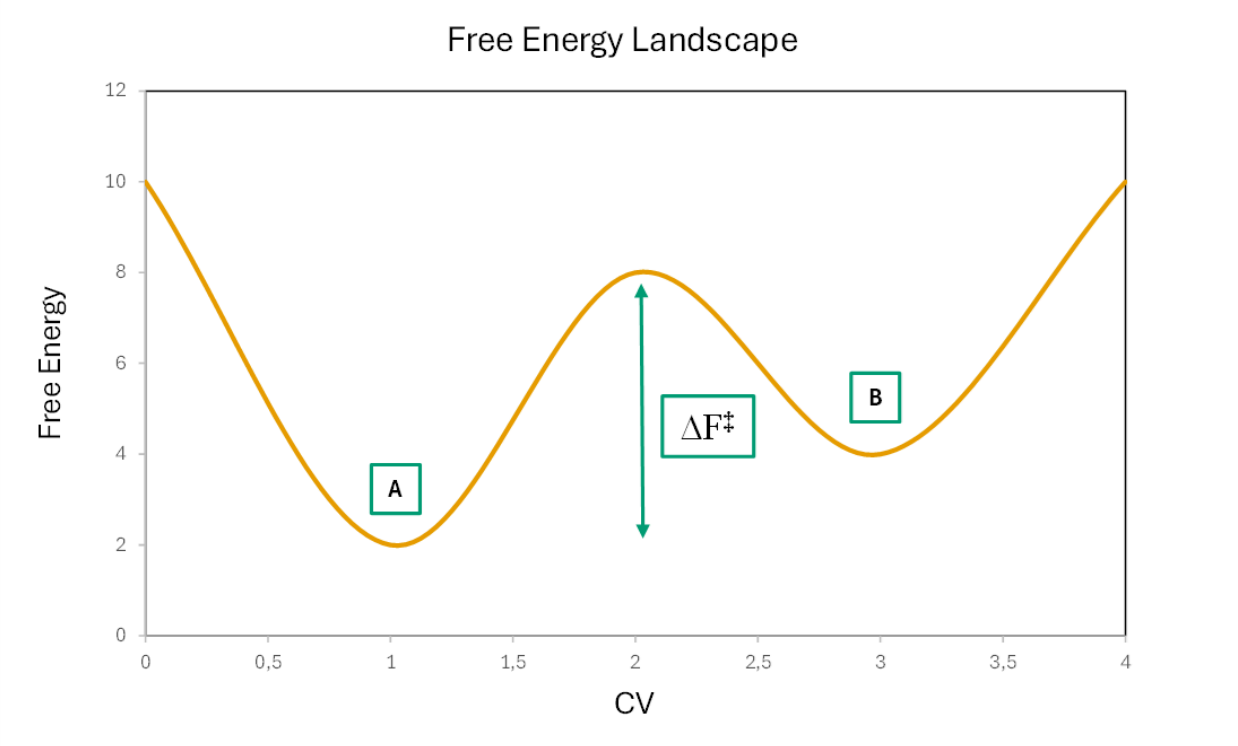

Figure 3: *Free energy model showing a double-well potential depicted as a function of a Collective Variable and Free Energy. $\Delta F^{\ddagger}$ denotes the difference in free energy between state A and the transition state.*

The analysis of the free energy landscapes is a complex problem, which will become clearer from the results and discussion (section 4). *If* a sequence of conformations that follow the canonical ensemble is produced, then computation of the free energy from the histogram of visited states becomes possible as:

$$F(s) = -k_B T \log N(s) \tag{7}$$

Here the free energy of the state *s* is directly related to the histogram count of that state *N(s)* (Bussi and Branduardi, 2015). In classical MD, the creation of such a distribution relies on the earlier mentioned ergodic hypothesis. This is a theoretically problematic exercise and in practice, it is perhaps more productive to turn to the practice of *metadynamics*.

2.3 Metadynamics

Metadynamics (metaD) is an addition to classical MD that really doesn't require any massive leaps of faith. It rests on the introduction of *sampling bias* to simulation. If we take another look at figure 3, one might imagine a world where the simulation "never" leaves the respective low-energy wells in A and B, due simply to the statistical unlikeliness of moving through the energy barrier at $\Delta F^{\ddagger}$. Now the prospect of deciding the system's preference for either state becomes practically hopeless, because the realistic time scale of observing the event is beyond our computational resources.

If an additional potential is added to the potential energy *U(q)*. Which raises the potential energy of places that have been visited by a collective variable. This would produce a different free-energy, but one that is mathematically relatable to the original one, because the added bias can be recorded, and accounted for later. In metaD this increase in potential energy comes in the form of *Gaussians*, also known as normal distributions. They are sometimes referred to as hills, which have a preassigned width and height, determined so that their individual sizes are negligible in the context of a given collective variable. These are deposited at a given interval throughout the simulation, raising the potential energy of configurations where the simulation has already been.

The addition of gaussians to a CV encourages, but does not force the system to move on from a certain state. As an exceptionally relevant example: The CV employed in this work is the *RMSD* (Root Mean Square Deviation) of the ligand compared to its original position. It is quite possible for the ligand to revisit a CV value through a different microscopic configuration, which we want to encourage, as we are interested in the "depth" of that well. Different heights in the normal distribution added can force a ligand to abandon its pose slowly, and more evenly, or faster, but with a less representative number of visits per state.

Crucially: *The total deposition of bias provides a negative rough estimate of the free energy profile* (Bussi and Branduardi, 2015). This gradual addition of a rough negative of free energy to the free energy tends to flatten the landscape, leading to a useful allegory. The process of metadynamics can be likened to filling the free energy landscape in with sand (figure 4), and measuring its depth.

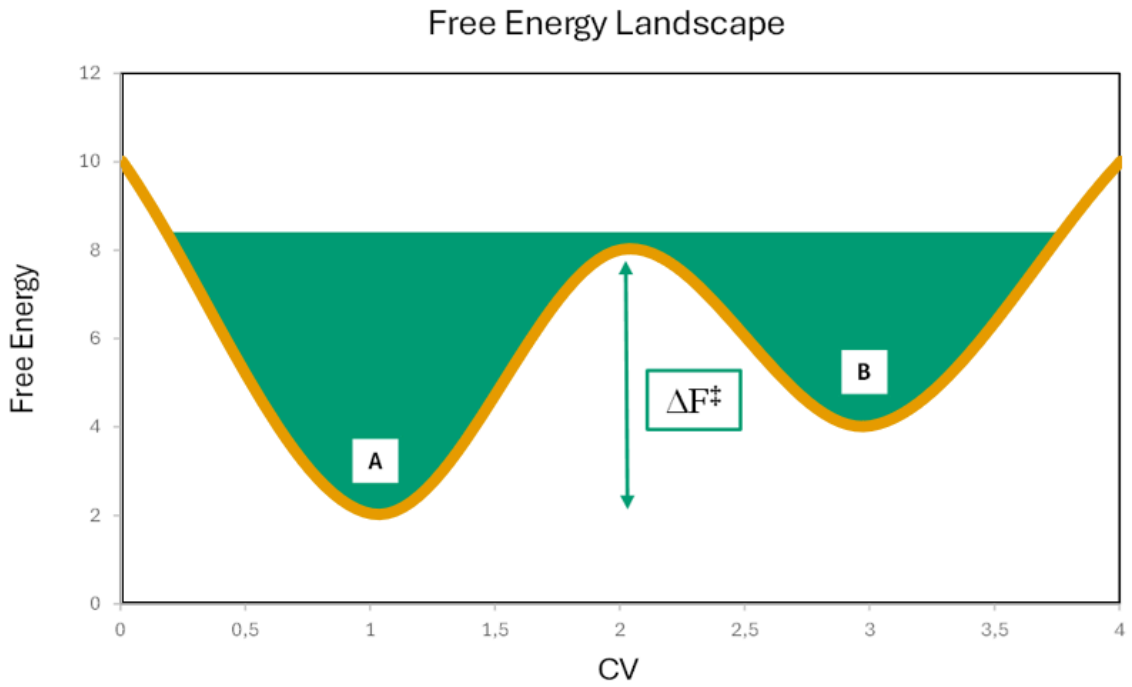

Figure 4: *Depicts the free energy landscape seen in figure 3. The filled in area represents an added bias potential*

The core assumption underlying this method is that the potential defined by the sum of added Gaussians in fact *is* an unbiased estimate of the free energy. This has been shown to be an increasingly accurate assumption, which enables the use of metadynamics to either *escape an energy minimum*, or to *explore a region of CV space* and reconstruct the free energy surface (FES). In the case of *well-tempered metadynamics,* which hasn't been resorted to in this work, a gradual reduction in the height of the added gaussian leads (over time) to a constant (exact) estimation of the free energy (Bussi, Alessandro Laio and Pratyush Tiwary, 2018).

Now is a good time to note that the method employed in this thesis makes no attempt to exactly determine the free energy of the system. Instead it focuses on an exploration of the CV space, made up of ligand-rmsd in relation to its predicted pose. It focuses on the depth of the potential energy wells near to the starting position of the ligand.

There are drawbacks to the application of metadynamics, in spite of its impressive estimation of free energy. In a single run of standard metaD, the total bias does not converge to a constant value of free energy, but rather oscillates around it, which can lead to a sort of overfilling. The uneven addition of gaussians can result in small, artificial energy barriers that push the system to spend time where it otherwise wouldn't (Barducci, Bonomi and Parrinello, 2011) For example, in a protein folding problem, it is possible to reach energies that are high enough to unfold secondary structures, which invalidates the usefulness of the simulation (Bussi and Branduardi,

2015). With this in mind, it is not necessarily easy to know when to stop a simulation, which will be discussed further in the results section.

The perhaps most difficult theoretical consideration is the identification of collective variables that can condense a complex problem in a useful way. It is at the core of running a simulation, and in spite of not having a general solution, there are some guidelines that can be applied. CVs should be able to *distinguish between one state and another*, while capturing all the relevant things that happen between them. They should *include all of the slow modes* of the system, meaning that there can be events that happen on a slower scale (e.g a larger movement in the protein), which might prevent the evolution of metadynamics (e.g a molecule exiting a binding site). If such events are not monitored, then the bias applied to the CV will fail to describe the full behaviour of the system within the timeframe of the simulation. This means that the choice of CVs must be informed by an understanding of the process it is meant to capture, which results in a natural process of trial and error. The number of *CVs should be kept to a minimum*, this reduces computational costs, and makes it far easier to interpret the results in a meaningful way (Barducci, Bonomi and Parrinello, 2011). As was mentioned in the previous section, the purpose of introducing these variables is to *reduce the complexity of the problem*, while preserving the aspects of interest.

As is the case in all scientific endeavours, reality is reduced to a model that *makes useful predictions*. Only within the contextual bounds of its usefulness, can it be said to capture something of reality.

2.3 Homology Modelling

It often happens that a protein structure has been resolved, while the sequences coding for its close relatives are known, but their structures are not. In these cases, researchers can attempt to model structures for *homologous* proteins from a known, related structure. In this thesis, the structure of GluN2B has been modelled on a crystal structure of GluN2A. Specifically, this thesis has relied on energy-based homology modelling (Schrödinger Documentation, 2024). It relies on calculations of potential energy and alignment of sequences to suggest a stable structure for a homologue. For this it uses the *Prime Energy* model, a force field described in section 2.4. It also relies on libraries of side-chain and backbone orientations created using data available in the Protein Data Bank (PDB).

2.4 Force Fields

Making useful predictions in the context of molecular dynamics tends to start with a good guess. Challenges like homology modelling depend on a description of physics that is robust enough to be general, but computationally cheap enough to handle high throughput. Prime Energy is essentially a force field, in that it models atomic interactions. Aside from homology modelling it

has been used in this thesis to prepare 3D-structures of proteins and ligands, and applied in flexible docking to suggest binding modes for ligands in GluN2A and GluN2B.

The authors of this model specifically point to homology modelling as a type of problem that can benefit from improved descriptions of protein forces, as there are many cases where a high quality crystal structure exists for one protein, but not for its close relatives. Deducing with some confidence the dominant structures of these relatives should, in principle, be possible to tackle. We can expect the resulting arrangement to be fairly close to the crystal structure reference (Li et al., 2011).

Upon the publication of Prime (under the less exciting name of *the VSGB 2.0 model)*, Li et al. stated:

"... it is unlikely that the entire universe of biologically relevant protein structural data can be accessed by exclusively experimental means. (Li et al., 2011 p. 1)"

That claim has yet to be challenged, and was in reference to the reality of life on the atomic scale, where biomolecules exist in a state of constant flux, leading to a number of different, yet structurally relevant conformations for any given protein sequence. This fact of life points to the need for development of energy models like Prime (Schrödinger, 2024), which can account for the dynamic motions of molecular interactions. The model was an improvement on their earlier work, and utilises a continuum representation of solvent, meaning that the existence of bulk solvent molecules is implied, rather than explicitly simulated. This simplifies the treatment of water (the common solvent), but introduces some weaknesses with regards to the roles of specific water molecules, which can play integral roles in the structure of a given protein (Li et al., 2011).

The Prime Energy model was created through the fitting of parameters to a large database of crystallographic data, with the express purpose of being able to predict side chain conformations and the arrangement of irregular structures, or loops. Their energy function added to the existing OPLS-AA protein force field with a solvation term, and a group of physics based corrections. Their implicit solvent model is based on the *Surface Generalised Born,* or SGB model, and a variable treatment of dielectric constants, in an attempt to better account for the polarisation of certain residues (Li et al., 2011). In the interest of keeping things condensed, the Schrodinger Inc. support documentation lists the properties involved in determining a structure's Prime Energy as:

Coulomb energy, Covalent binding energy, Van der Waals energy, Lipophilic energy, Generalised Born electrostatic solvation energy, Hydrogen-bonding correction, Pi-pi packing correction, Self-contact correction (Schrödinger Documentation, 2024).

These provide an example of atomic phenomena that need to be taken into account when constructing a generalizable model. Li et al. stress in their paper that their model, in spite of being empirically fitted to training data (solving a *specific problem*), still tended to result in a better description of the underlying physics of the model (a *general problem*).

The authors express confidence in Prime Energy's relationship to the underlying physics:

"... the physics of the model is correct, and the right answers are being obtained for the right reason. (Li et al., 2011 p. 4)"

For MD and metadynamics a different force field has been employed, where the simulations rely on explicit solvent molecules. The methodology has more information on this (see section 3.5).

2.5 Ligand Docking

Prior to running an advanced simulation it is useful to have a place to start. A potential binding mode for the protein-ligand complex. *Docking methods* are a common approach to quickly generating such structures. In this work, such a method is provided by *Glide.*

The name of the Glide method is really a semi-complete acronym, coming from *grid-based ligand docking with energetics*. The computational method is based on a series of hierarchical filters, allowing for the sampling of positional, conformational and orientational space across a large number of ligand structures. Upon its publication, Glide widely outperformed the state of the art docking methods used in drug discovery (Kirkpatrick, 2004). It relies on the creation of a *grid,* which represents the shape and properties of the receptor protein as described by a set of fields. This structure is created once for a given structure. Following the creation of the grid, initial screening takes place. This screening samples energy minima from the torsion-angles of a ligand, generating conformations that are used to search for promising ligand poses across the grid. Beginning with these poses, the ligand goes through minimization within the field of the receptor. Glide relies on available versions of the force field *OPLS* (Optimised Potentials for Liquid Simulations) (Harder et al., 2015). The lowest-energy poses are put through Monte Carlo simulations to examine nearby torsional minima. For selection of the docked pose most likely to match the crystal structure used as training data, the authors of Glide employ a combination of their GlideScore, a ligand-receptor molecular mechanics interaction energy and a ligand strain energy (Friesner et al., 2004). The scoring function is built from modifications of the empirical scoring function ChemScore. It is designed to estimate the free energy of binding for a protein-ligand complex, given a known 3D-structure (Eldridge et al., 1997).

The current work uses the extra precision Glide method. Developed around the same time as its standard precision counterpart, Glide XP is more quality-oriented and computationally

demanding. It introduces large desolvation penalties to polar and charged groups as well as identification of certain structural motifs that provide large contributions to enhanced binding affinity (Friesner et al., 2006).

It is noted by the authors that the largest issue faced when rigidly docking a library of ligands into a receptor this way, is the inability of certain active compounds to fit into the active site. This is due to the site itself or the approach to it being too small to accommodate the ligand. The Glide-protocols reduce steric clashes by scaling Van der Waals radii of nonpolar atoms, resulting in a modest "induced fit" effect.

2.6 Induced Fit Docking

The induced fit docking protocol relies on the previously described Glide and Prime functionalities, and attempts to account for the inherent flexibility of both ligand and receptor. This raises the degrees of freedom in the system, thereby raising the computational cost of docking into the structure. The induced fit docking approach lies between the extremes of softening interaction potentials and running explicit molecular dynamics. It addresses the dual flexibility of the system by combining methods for docking ligands into a rigid receptor, with models of conformational changes in the receptor (Sherman et al., 2006).

On the whole, IFD can be described in four steps. [1] Softened potential docking with a rigid receptor, generating a collection of poses. [2] Sampling of receptor side-chain rotamers (within 5Å) for each pose (Prime). [3] Redocking of the ligand into low-energy structures given by the previous step (Glide). [4] Assignment of scores for each pose, as given by the docking energy (GlideScore), and the terms for receptor strain and solvation (Prime Energy). When the top scoring outputs of step four have close to identical scores, the protocol runs a second round where the ligand is redocked into the top ranking structures of the fourth step (Sherman et al., 2006).

Challenges introduced in these steps include generating at least one plausible docked pose in the first step, independent of scoring. Without an initial guess that is close enough to a realistic pose, it becomes unlikely that the receptor will usefully reorganise within the simulated timespan. The solution implemented by Sherman et al., other than the softening of electrostatic potentials, involves the notion that residues with a high degree of structural flexibility are the ones most likely to be blocking a binding site. They apply a temporary mutation to Alanine to obtain a low-rmsd pose that can serve as an initial guess in the second step.

Part of the reasoning behind the scoring function applied in the IFD-protocol is that ligand binding affinity is primarily driven by *the hydrophobic effect*: The displacement of water in the binding site into the bulk solution, and the movement of hydrophobic parts of the ligand into lipophilic regions of the binding cavity. Prime Energy's continuum solvation model (implicit

water) has its greatest challenge in dealing with the free energy of waters bound to the binding site. For IFD the authors optimised the empirically calibrated Glide binding affinity estimator, to get a better estimation of these quantities (Sherman et al., 2006 p. 537).

The current state of their IFD scoring function is given by:

1.0***Prime_Energy** + 9.057***GlideScore** + 1.428***Glide_Ecoul**
(Schrödinger Inc.)

Indicating GlideScore as the usually dominating term. The Glide_Ecoul term is a measure of electrostatic interaction energy.

2.7 The Case for IFD-MD

There is good reason for imagining a workflow that combines flexible docking (induced fit) with some version of molecular dynamics. Given the availability of a qualified guess as to how a ligand interacts with a protein (IFD), there is the possibility of refining or testing that guess with more thorough methods (MD/metaD).

It is laid out in detail by Clark et al. (2016), in a paper titled: *Prediction of Protein–Ligand Binding Poses via a Combination of Induced Fit Docking and Metadynamics Simulations*. The methods employed in this thesis have been largely inspired by theirs.

They note that the IFD-methodology usually locates a structure with good accuracy within the top 10 poses. This isn't enough to be actionable in many drug-discovery applications. To further refine the guesswork and raise confidence in the binding mode of a ligand, they introduce *metadynamics pose scoring* to assess the stability of the suggestions. Often a normal MD simulation will fail to dislodge the ligand from its binding pose due to large free energy barriers between binding poses (see section 2.3). By instead adding a metadynamics bias to the initial pose, using the rmsd of the ligand to its original position as the collective variable, it becomes viable to investigate pose stability (Clark et al., 2016).

Importantly, the IFD-MD method does not aim to calculate the full free energy surface. It instead aims to use a metadynamics approach to determine relative stabilities for different binding poses (Clark et al., 2016 p. 2991). Ligand-rmsd, used as a collective variable, can be reasonably assumed to correlate with the stability of the protein-ligand interaction. A ligand in a stable configuration is in a relatively deep energy well. It should require a longer application of bias to dislodge it from a stable binding pose, given a constant metadynamics bias being applied.

Their work demonstrates (with a test set of 42 protein-ligand complexes) a dramatic reduction in the distances seen between the top-scoring pose and the crystallographically determined binding mode. This method was later developed further by Miller et al. (2021).

3. Methodology

This section outlines the important details involved in getting the results of this thesis.

3.1 Crystal Structure & Homology Model

The Crystal Structure has a resolution of 1.7 Å and consists of the agonist binding domains from one GluN1-subunit and one GluN2A-subunit, co-crystallised with three ligands: DCKA (a glycine site agonist), Glutamate, and AA33 (a negative allosteric modulator).

The homology model (HM) of subcomplex B was created from the prepared crystal structure (CS) of GluN1/N2A-ABDs. The workflow included three Schrodinger functions: The Multiple Sequence Viewer, Energy Based homology modelling, and the Protein Preparation Wizard. A single non-conserved residue: Phe262, was altered to its highest probability rotamer, based on the internal rotamer library. The non-native ligand present in the crystal structure was removed prior to the model's creation. The protein preparation function was run after the energy based homology modelling, with a pH of 7.4 +(-) 0.05. When compared to the original crystal structure, the RMSD following Protein Structure Alignment was 0.227 Å between HM and CS.

3.2 Ligand Library

10 of the 15 ligands included in this work have been tested experimentally through single-clamp voltage experiments, and have been shown to have a measurable impact on the ion-permeation of either GluN2A, GluN2B or both. They are all analogues of TCN-213. The 5 remaining ligands were recently synthesised in a different masters project with the Praetorius group at The University of Copenhagen. All ligands were processed using Schrödinger's Ligprep, to produce a single, low energy 3D-structure for each ligand. These were then used for Induced Fit Docking into subcomplexes A and B.

3.2 Docking Grid

The grid used for Induced Fit Docking specifies the boundaries of the binding site that is to be docked into. In subcomplex B its positioning is based on a different homology model, where the non-native ligand remained in the structure. Specifically, the centre of the 10Å-box was defined as the coordinates of the non-cyclical sulphur atom in AA33, following structural alignment of this homology model with subcomplex B. In the crystal structure (subcomplex A), the grid has been defined around the positioning of the co-crystallised AA33.

3.4 Induced Fit Docking

The Induced Fit Docking method (see section 2.6) was executed using its Standard Protocol, with Extra Precision Glide (XP) used for redocking. For each of the ligands, within each subcomplex, a total of 20 poses were produced. Depending on the difference between these, some are classified as unique, while others are deemed too similar to be included in the output. For subcomplex A, a total of 176 poses were created. For subcomplex B that number was 177. Following this a selection was made based on IFDscores (see section 2.6) and the similarity of ligand poses to that of AA33 in the original crystal structure (all ligand structures share a central scaffold). The top ranked pose was selected unless they exhibited a large deviation from the crystal structure pose, following alignment of protein structures. In case of deviation, another high-ranking pose was added to the selection which aligned more closely to the crystal structure. CS comparisons were made visually and the selection was partly influenced by which ligands had displayed interesting SAR-data. Following selections there were 23 poses docked into subcomplex A, and 20 for subcomplex B. These were converted into PDB-format before being processed by a script designed to make them Gromacs-compatible.

3.5 Topologies

Force field parameters for any new species (ligands) must be derived and validated in a manner that is consistent with the original force field. In light of this we must *manage protein and ligand topologies separately*. For protein topologies, we are using the force field AMBER99SB-ILDN (Lindorff-Larsen et al., 2010) and the TIP3P-model, originally published by Jorgensen et al. in 1983, for modelling the behaviour of solvent molecules. This choice of solvent model is justified both by its computational efficiency and compatibility with established force fields like AMBER.

Ligands from IFD-outputs were moved into separate files before processing by *ACPYPE* (Sousa da Silva and Vranken, 2012) for the creation of topologies. We can rely on acpype and gromacs to recreate bonds using given coordinates, so the rows designating molecular bonds created by Schrodinger are removed. Solvent molecules are also removed from all structures, since each system is artificially solvated later on.

Topology files are (essentially) independent of molecular coordinates, and thus the appearance of acpype in the workflow is limited to whenever there is a new ligand structure. The topology files of static ligands (glycine-site antagonist DCKA and glutamate) should remain compatible with both subtypes of the N1/N2-NMDA complex. This allows using the same parameters for these ligands topologies, removing a potential source of error. Acpype relies on *antechamber* for the creation of topologies. Antechamber is a part of AMBER, and thus applies *GAFF* (General AMBER Force Field), when designating charges, atom types and bond information (Weiner and Kollman, 1981). The method for assigning partial charges was *AM1-BCC***.** This method approximates quantum mechanical calculations through parameters derived

either from higher level calculation or experimental data. This reduces computational costs at the price of some accuracy.

3.6 Alignment

For the transfer of new pose-coordinates produced by IFD, we are using a superposition of protein alpha carbons to create an alignment. This relies on the existence of a compatible reference structure (e.g the crystal structure). For this the *mdtraj* python library is employed. The script superimposes the protein alpha carbons of a new pose onto those of a reference complex. The method is based on the assumption that the differences in protein backbone coordinates will be negligible around the static ligands compared to the later equilibration of the system.

Having prepared the aligned protein structure file, the *pdb2gmx* tool is applied to generate a protein structure file using the selected forcefields. Histidine protonation states were determined manually.

3.7 Pre-Simulation Protocol

This workflow was carried out once per IFD-pose transferred to MD/metaD.

Define Simulation Box

We begin the equilibration of the system by defining its periodic boundary conditions. To ensure that the complex does not "self-interact", a distance of 0.7 mm has been set between the boundaries and the protein-ligand complex. To reduce total simulation volume a truncated octahedral geometry is used instead of a cubic one.

Solvate & Exchange Ions

Following this, the space remaining in the system is filled up with water molecules whose behaviour will be governed by the TIP3P-forcefield. As a standard procedure the system is neutralised with counterions. In this case the net-positive charge of the complexes is matched with the conversion of water molecules into chloride ions.

Indexing

The indexing for this system has been scripted to generate an accessible category consisting of protein structure and the three ligands present. This is necessary for applying position restraints in the minimisation and equilibration stages.

Minimise

The minimisation is designed to relieve initial steric clashes and overlaps within the system. By optimising potential energies towards local minima, atomic coordinates are optimised in preparation of MD simulation.

Equilibrate

The equilibration is run as a short, restricted molecular dynamics simulation with the intent of allowing system properties like temperature and pressure to approach a steady state. This allows for an early interpretation of the stability of the system and functions as a test case of whether the simulation is running as intended. This stage is run as three sequential simulations, where the output of the first is the input of the next.

Constant volume (NVT), 100 ps, all protein and ligand atoms restrained.
Constant pressure (NPT), 500 ps, all protein alpha-carbons restrained
Constant pressure (NPT), 1000 ps, no restraints applied

3.8 Run MD/RMSD-metaD

The output of the final equilibration is used to compile run-files for MD and RMSD-metadynamics simulations. MetaD-runs have been set to a length of 20 ns and three replicates. These are complemented by a regular MD simulation of 100 ns. Metadynamics runs have been set up to apply bias to the RMSD-value of the ligand to the initial pose held following IFD and equilibration. The consistency of RMSD measurement is based on an alignment step, defined by the alpha carbons of all residues having atoms within 5 Å of the crystal structure ligand-pose. This step ensures that the conformations produced by the simulation are in the same orientation and position, before a comparison is made.

In the case of the homology model of GluN2B, the residues used as an alignment reference are defined as those within 5Å of the GluN2A starting ligand position (AA33y4), after the structure has been aligned to the GluN2B model. rmsd-biased simulations have been run using *PLUMED* (Tribello et al., 2014). The applied hill height is 1.2 kcal/mol, with a width of 0.005 Å. The pacing was set to bias deposition once every 500 steps (simulation step-size of 2 fs).

3.9 Analysis

Analysis of simulation output begins with adjusting the frames of the simulation trajectory, ensuring that the complex is not split across multiple periodic boundary conditions. Trajectories are visually inspected using VMD. VMD has also been applied in the creation of Table 2 (section 4.3). The *hbonds* utility was used with an angle cutoff of 60 degrees and a hydrogen bonding distance cutoff at 3 Å.

The evolution of all systems can be overviewed compactly by plotting values of ligand RMSD against time. To inspect the outcomes of the three metadynamics runs, it is useful to observe the free energy surface (FES) at different values of ligand-RMSD. The triplicate runs of metaD have been averaged when constructing these.

Another measure of interest is a Boltzmann-weighted average of the RMSD, which can be computed from the FES using:

$$\frac{s}{q} = \frac{x \cdot \exp\left(-\frac{e}{kT}\right)}{\exp\left(-\frac{e}{kT}\right)}$$

With **X** here being the values of the collective variable. This weights all values according to the boltzmann factor: **– (e / kT)**.

The weighted average is the closest thing to a single number that might indicate the outcome of metadynamics. In this work, all FES-RMSD values have been separated into four groups, dividing the simulation into successively longer FES-averages. These averages are taken from the first 5, 10, 15 and 20ns. This is meant to enable observation of the pose stability before the ligand leaves the binding pocket, and to distinguish the persistent binding modes from the ones that are quickly dislodged by the added potential. Another way in which the outputs of the simulations have been compacted is via clustering of the trajectories of MD-simulations. Here the *GROMOS* clustering method (Daura et al., 1999), has been applied with a cutoff of 1.3 Å. Alignments prior to clustering were performed on the alpha-carbons of the binding site residues, defined in section 3.8. 3D images of structures have been produced using Maestro. The interactions between ligands and binding site residues that can be seen in section 4.2 have been identified by Maestro, where no parameters were altered. IFD-poses, main-clusters and MD-trajectories were inspected in VMD.

Calculations of RMSD between binding sites and the crystal structure (A) or homology model (B), were performed using a least squares alignment to the alpha carbons of binding site residues (defined earlier), followed by taking an rmsd between all heavy atoms of these residues.

4. Results & Discussion

What are the likely stable binding modes of known allosteric binders and their synthetic analogues in the closely related receptors GluN2A and GluN2B? This has been approached using a high resolution (1.7 Å) crystal structure of the agonist-binding domain of the GluN1/N2A NMDA-receptor. The GluN1/N2B homology model was based on the crystal structure. Fifteen ligands, ten of which are known binders of either A, B or both have been subjected to Induced Fit Docking. 43 High-ranking outputs from IFD have been selected and evaluated with 3 x 20 ns of biased rmsd-metadynamics and 100 ns of unbiased molecular dynamics. Reconstructed free energy surfaces and MD have been used to evaluate pose-stability. Simulations indicating a stable binding mode have been visually evaluated.

4.1 Overview of Simulation Results

An initial indication of pose-stability can be read from glide-scoring of IFD-poses, where a lower score is more stable. In Table 1 these have been paired with the additional stability-indicators retrieved from MD and metadynamics. Averages of ligand-rmsd from unbiased 100 ns MD-runs are reported together with averages from triplicate metadynamics-runs at 10 and 20 ns. Values from metaD are calculated from reconstructed free energy surfaces, which are based on bias potentials added over the simulation. As was discussed in the background (section 2.3), greater added bias is presumed to correlate with binding mode stability. Table 1 also presents measures of structural deviation between clustered MD-trajectories, IFD-outputs and the crystal structure (complex A) or homology model (complex B).

Table 1: *Simulation outcomes of a selected group of IFD-poses for each subcomplex. The first columns contain ligand names, followed by an indexing number if there are two poses for that ligand. The second columns contain glide scores for IFD-poses. Columns for 10 and 20 ns contain ligand-RMSD averages of triplicate metaD-runs, in the form of Boltzmann weighted values calculated from free energy surfaces. MD columns contain the average ligand-RMSD of 100 ns of unbiased MD. Structural RMSDs are measures of deviation in binding site atoms between clustered MD-trajectories (clust), IFD-poses and the crystal structure (CS) or homology model (HM). Filled-in cells mark out simulations that are presented visually and discussed in section 4.2.*

Complex A		Metadynamics RMSDs			structure RMSDs (Å)			Complex B		Metadynamics RMSDs			structure RMSDs (Å)		
Ligand	IFDGlide	MD	10ns	20ns	clust-CS	IFD-CS	clust-IFD	Ligand	IFDGlide	MD	10ns	20ns	clust-HM	IFD-HM	clust-IFD
AA33	-9,7	N.D	2,3	1,7	**0,8**	**0,6**	**0,7**	AA33	-9,8	1,9	2,0	2,7	1,8	0,1	1,5
AA34-1	-8,8	**2,0**	1,8	1,8	1,2	0,7	1,2	AA34	-9,9	2,2	2,4	2,2	1,4	0,7	1,5
AA34-2	-8,3	3,3	3,3	3,2	0,9	0,8	1,2								
HQ09-1	-8,6	2,0	2,8	3,0	1,4	0,7	1,5	HQ09	-8,9	1,3	1,1	2,6	1,3	0,7	1,3
HQ09-2	-8,2	2,0	3,1	3,0	1,3	0,5	1,2								
HQ11-1	-9,0	2,9	3,4	3,5	1,2	0,8	1,4	HQ11-1	-10,7	2,4	2,7	3,8	1,3	1,1	1,3
HQ11-2	-8,1	**2,4**	2,1	3,1	1,2	0,7	1,1	HQ11-2	-9,4	2,2	3,4	3,8	1,5	0,6	1,5
HQ18	-9,5	0,9	0,8	1,2	0,8	0,7	1,0	HQ18	-8,8	**1,5**	1,3	1,7	1,3	1,0	1,3
HQ19	-8,5	1,5	1,5	1,4	0,8	0,6	0,9	HQ19-1	-9,8	**2,1**	1,9	5,6	1,6	1,2	1,5
								HQ19-2	-9,1	2,9	3,2	6,0	1,3	1,2	1,5
HQ3R	-7,4	4,6	2,3	2,8	2,2	1,2	2,0	HQ3R	-9,1	3,9	4,4	4,5	2,4	1,5	2,3
HQ3S-1	-7,3	3,5	3,2	7,8	2,1	0,8	2,2	HQ3S-1	-8,3	3,2	3,1	5,4	1,8	1,3	1,9
HQ3S-2	-7,8	2,4	2,9	2,3	1,9	0,6	2,1	HQ3S-2	-8,1	3,6	5,4	5,4	1,5	1,1	1,7
KI48-1	-9,7	2,2	2,6	2,2	1,3	1,2	1,4	KI48	-8,7	3,0	2,0	3,3	1,5	0,9	1,6
KI48-2	-7,8	3,4	4,0	3,9	1,5	1,2	1,6								
KP3A	-5,9	3,2	1,5	1,5	1,8	0,5	1,8	KP3A	-8,1	3,7	5,7	9,1	1,7	0,7	1,8
KP3B-1	-7,4	2,0	4,4	4,8	1,5	0,7	1,6	KP3B	-8,9	2,1	3,5	3,6	1,6	1,0	1,3
KP3B-2	-7,1	3,0	3,1	4,4	1,2	0,9	1,3								
KP6A	-7,8	1,7	2,4	2,6	1,4	0,6	1,4	KP6A-1	-9,1	1,8	2,4	1,9	1,4	1,1	1,4
								KP6A-2	-9,3	2,2	2,4	2,8	1,3	1,0	1,2
KP6B-1	-9,4	3,5	5,3	5,2	2,2	1,6	1,9	KP6B-1	-9,3	3,0	6,2	7,8	1,3	1,1	1,5
KP6B-2	-7,2	2,4	3,3	3,5	1,1	1,6	1,7	KP6B-2	-9,3	3,4	4,3	4,2	1,8	1,1	1,5
KP6C	-7,7	3,1	5,5	6,4	1,5	0,7	1,6	KP6C	-8,0	5,8	3,0	6,4	1,3	1,1	1,5
TCN-1	-8,8	3,0	3,3	4,1	1,6	0,6	1,5	TCN	-8,9	2,8	2,4	2,6	1,7	0,6	1,7
TCN-2	-8,4	2,5	3,5	3,4	1,5	0,9	1,4								

4.1.1 Glidescores and MD/metaD

The initial selection of poses was based on IFD-scores, which rank different poses of the same ligand. Within the same ligand and complex, the first pose (marked 1) had the lower (better) IFDscore out of the two poses. These scores were left out of Table 1, as they do not translate as

well between different ligands as the glide scoring. Table 1 presents the ligands aligned across the two complexes. The poses are unique, as they have been docked into different complexes. Interpretation of metadynamics-averages presents a challenge, especially considering that these are non-converged simulations. As mentioned in section 2.3, this application of metadynamics only begins to explore the free energy landscape. With three replicates per pose there is still a random element to the appearance of free energy surfaces and their associated RMSD-averages. Computing the average at increasingly longer intervals intended to capture the gradual exit of the ligand from the binding site, without introducing a cutoff. This approach was taken partly because it is difficult to determine at which RMSD a stable pose lies (assuming one exists). If the model works as intended, then the unbiased MD simulation can find a local minimum where it remains, while the metaD-runs will diverge away from that pose given enough time. Assuming that the metaD simulations capture the relevant motion in the structure, a longer duration in the binding site should correlate with a stable binding mode. How far the ligand moves away from the binding site once it has been dislodged has not been given any importance.

A complexity arises when the MD-average diverges from the one suggested by IFD. When these models give widely different answers it is possible that the difference lies in the force fields. There is always some disagreement between the models. The lowest MD/metaD-averages are not necessarily the most stable, but when the observed metaD-averages are low enough to still lie within the binding site, then it is assumed that persistently low values can be associated with a real minimum in potential energy. This kind of energy minimum can be easier to identify from a graph of the free energy surface, these are available in the appendix. The FES-graphs have been used to select which trajectories were inspected visually.

The MD-clusters have been taken as averages of the entire 100 ns simulations. They contain the structures that have contributed most to that averaged RMSD. These have been interpreted as the binding mode that best represents the local minimum in free energy that IFD has suggested for the ligand. Seeing as the MD-runs are the more rigorous and computationally exhaustive method, they are taken as a refined suggestion compared to IFD. Some poses have been investigated further even when MD and metaD diverge from IFD. Especially in the cases where there is a close match in RMSD between MD and metaD, the reasoning being that both might have found a relevant pose that IFD didn't. There exists an opposite possibility, that neither MD or metaD locate a stable arrangement of the ligand in the IFD-suggested energy minimum. Even where such a binding mode exists, the absence of water molecules in specific positions or some oversight in the force field could leave it obscured.

4.1.2 Structure RMSDs

The clearest pattern in the structure rmsds (Table 1) is the differences in cluster-IFD being larger than IFD-CS or IFD-HM. In both subcomplexes a comparison of these columns points to a greater effect on the binding site residues following MD simulations. This is not always the case,

as with HQ3R in site A, but the pattern seems to be that the MD-clusters take a larger step away from their starting positions than the IFD-poses have done, seeing as MD starts from the IFD-pose and the IFD-poses start from the crystal structure or homology model. This is likely influenced by the movement of the protein backbone being unrestricted during MD. Whether this is a good thing depends on the accuracy of the MD force field. Another observation is the low rmsd-values seen for AA33 in site A. This row is different in that it redocks a ligand into the crystal structure. Assuming that the binding mode seen in the CS is a major contributor to the binding of AA33 in GluN2A, we should expect the models to come to a similar conclusion (low divergence). The close agreement of the three approaches should likely be investigated more closely before any conclusions are drawn.

Three MD-cluster rmsd-values in Table 1 have been highlighted. These are visually inspected trajectories whose ligands have an "inverted" central scaffold. The inversion was produced by IFD, and is defined in relation to the orientation of the ligand in the crystal structure. These binding modes are significantly different from the CS-ligand. Whether they contribute to the real binding of these ligands or not, they appear to be stable in the binding site. It is possible that the approach to the binding site would prevent the ligands from aligning to the scaffold this way.

4.2 Visualised Trajectories

This section presents IFD-poses and dominant clusters of MD-trajectories for cases where FES-graphs and metaD-averages have suggested a stable binding mode. Similarities of MD-averages and metaD-averages have also been taken into account. For each subcomplex two residues have been included in every image to provide a frame of reference. These are the non conserved Ile239 and Val267 (complex A). Figure 2 in section 2.1 gives an overview of binding site A. In complex B these residues are instead Val234 and Phe262. The fact that these are not conserved between the crystal structure and the homology model is the reason they are included. Binding modes that rely on either of these side chains present the possibility of designing ligands that interact more strongly with either A or B.

AA33 in site A

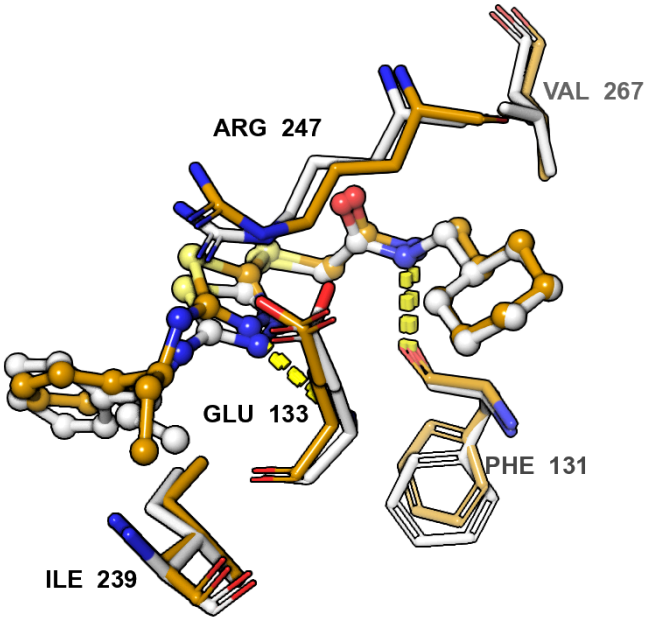

Figure 5: *Suggested binding mode of AA33 in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose. Hydrogen bonds are shown in yellow.*

This ligand, being the one originally found within the allosteric binding site of the crystal structure, one might expect to be highly stable in the GluN2A-complex. The metadynamics runs give a stable indication, with ligand RMSD-average at 20ns being 1.73 Å (see Table 1). The FES-graph indicates two energy minimums under 3 Å (see appendix). These could both be real contributors to the binding of AA33 in GluN2A.

Seeing as it is the original structure, MD has been run for 500 ns, throughout most of which the ligand remains at an average RMSD below 2Å. The shifts noticeable throughout the trajectory are in the ethyl and cyclohexane substituents.

AA33 in site B

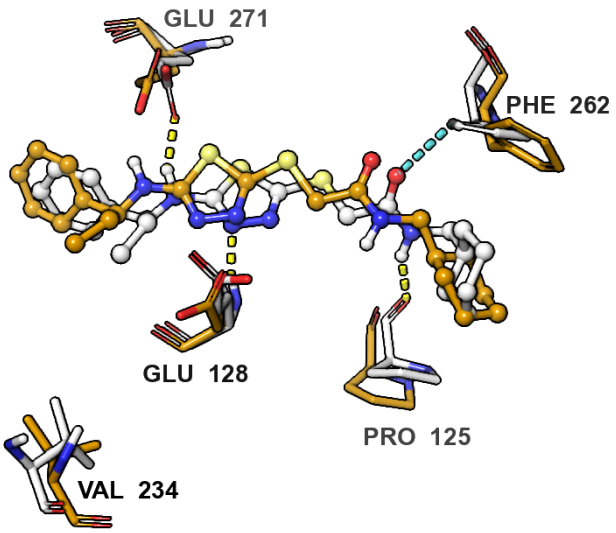

Figure 6: *Selected residues from the binding site of AA33 in subcomplex B. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the IFD binding site. Hydrogen bonds are shown in yellow. Aromatic-hydrogen bonds are shown in teal.*

In the case of GluN2B, whose structure was homology modelled from the GluN2A crystal structure, AA33 still performs quite well. It has an MD-RMSD average of 1.89Å, and MetaD-averages stay below 2Å for the first 10ns (see appendix). The trajectory shows a rotation in the cyclohexane ring, possibly away from Phe262, and a shift of Arg244 towards the ligand. The result of IFD here suggests a number of hydrogen bonds. These bonds may have been elongated in MD and thus no longer recognised by Maestro. Two of these interactions are present for the cluster in Table 2.

AA34 in site A

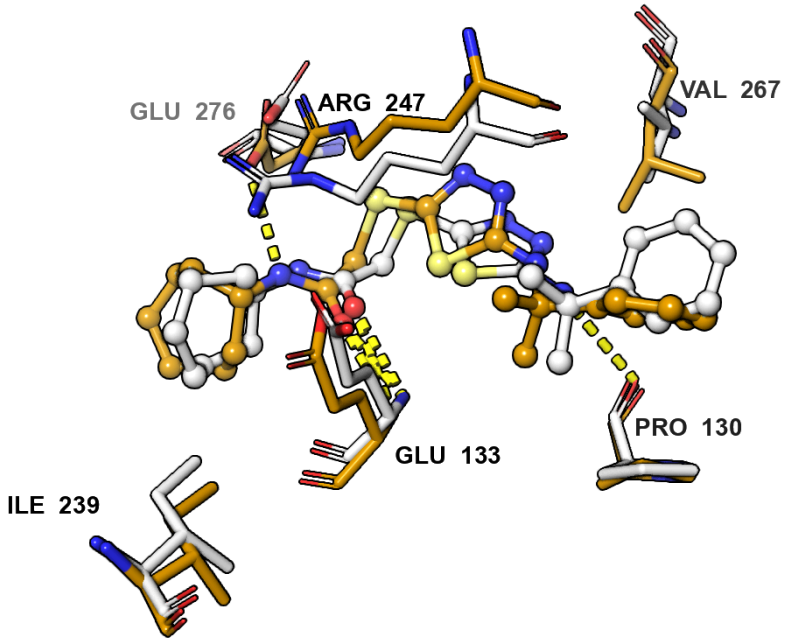

Figure 7: *Binding site of AA34-1 in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose*

The pose on display, AA34-1 has a slight difference in the methyl substituents attached to the benzyl moiety, and is inverted compared to most, with the scaffold rotated 180 degrees in the binding site. The MD-rmsd average is at 2.0 Å, while the average of metaD-outcomes remains below 2 Å for the entirety of the simulations. The cluster to IFD comparison here yields a shift in the rotations of both the cyclohexane and benzyl substituents. Most residue shifts between the poses track the movements of the ligand. There is a hydrogen bond with Pro130 that is suggested by IFD, it however occupies less than 1% of the simulation trajectory. The alternative pose, AA34-2 is not inverted. It has higher rmsd-averages where both metaD and MD seem to agree on a binding mode at 3.3 Å.

AA34 in site B

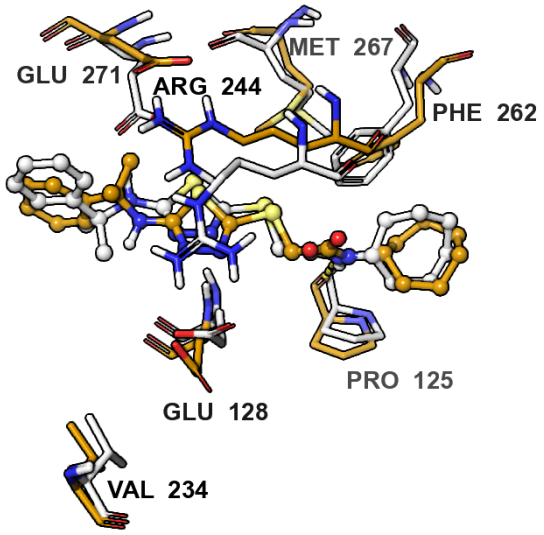

Figure 8: *Selected residues from the binding site of AA34 in subcomplex B. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Hydrogen bonds are shown in yellow.*

In the B-subcomplex, this pose has an MD-rmsd average of 2.16 Å, and is made more interesting by the outcome of metaD. Its averages over the course of metadynamics stay quite close to the MD-average for the entire 20ns. From observing the evolution of the metaD-rmsds (see appendix), only one of the simulations have dislodged the ligand, while the other two have returned to poses around 2 Å after ~15 ns. The trajectory shows the dimethyl-substituent turned upwards in our frame of reference, as well as a rotation in the cyclohexyl group, while the core scaffold remains stable. Met267 and Arg244 have made a clear adjustment towards the ligand, while Glu271 seems to have reoriented itself towards Arg244.

HQ09-1 in site A

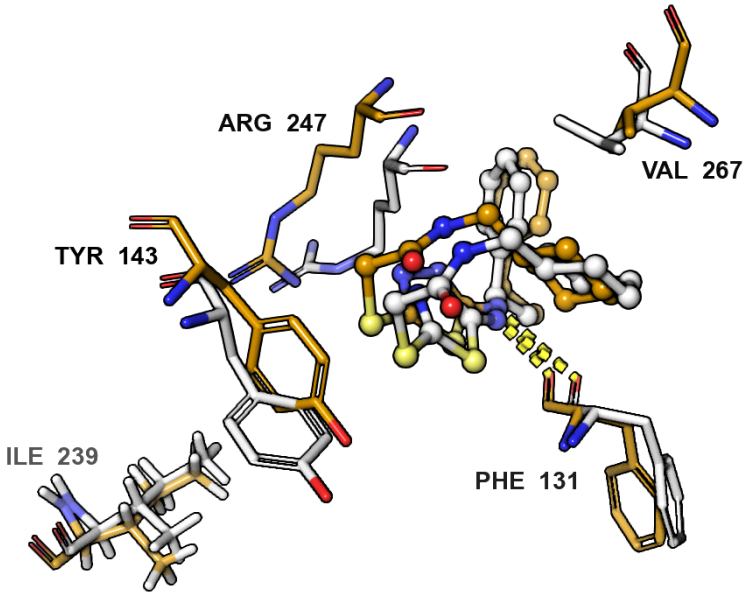

Figure 9: *Selected residues from the binding site of HQ09-1 in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose. Hydrogen bonds are shown in yellow.*

The metadynamics-averages do not give an especially stable prediction for this pose, except for the 5ns measure which is below 2 Å. This matches well to the RMSD average at 2.0 Å. The ligand displays an interesting pose, in that it has folded in on itself, with the benzene and cyclohexane groups being unusually close together. The MD-trajectory maintains this pose, while shifting side chains of residues Arg247 and Tyr143 towards the ligand. The trajectory of the ligand appears largely governed by self-interactions.

HQ09 in site B

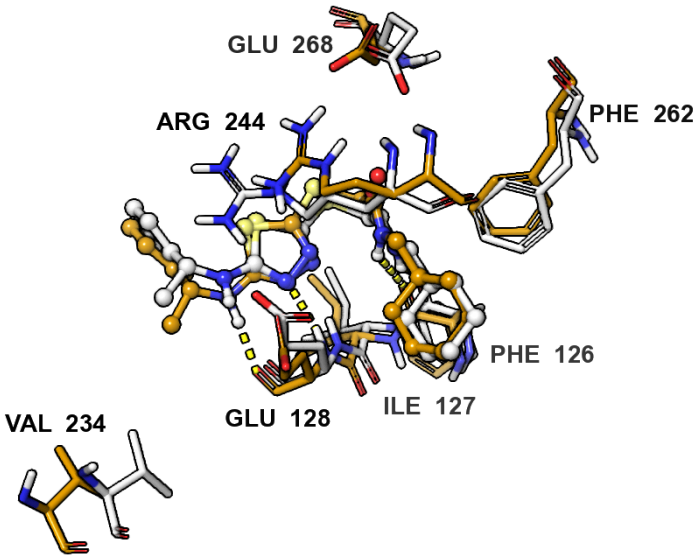

Figure 10: S*elected residues from the binding site of HQ09 in subcomplex B. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Hydrogen bonds are shown in yellow.*

This is another example of a seemingly stable outcome. Largely influenced by contributions of two out of three metaD simulations, the triplicate average of FES-rmsd stays close to 1 Å for 15ns. The average of MD-rmsd is unusually low at 1.3 Å. As is indicated by MD, the ligand trajectory deviates very little from the pose suggested by IFD. In the clustered pose there are small shifts in Glu268, Arg244 and Glu128 away from the ligand. Hydrogen bonding (Table 2) indicates a strong interaction between the central parts of the ligand and the main-chains of Glu128 and Ile127.

HQ19 in site A

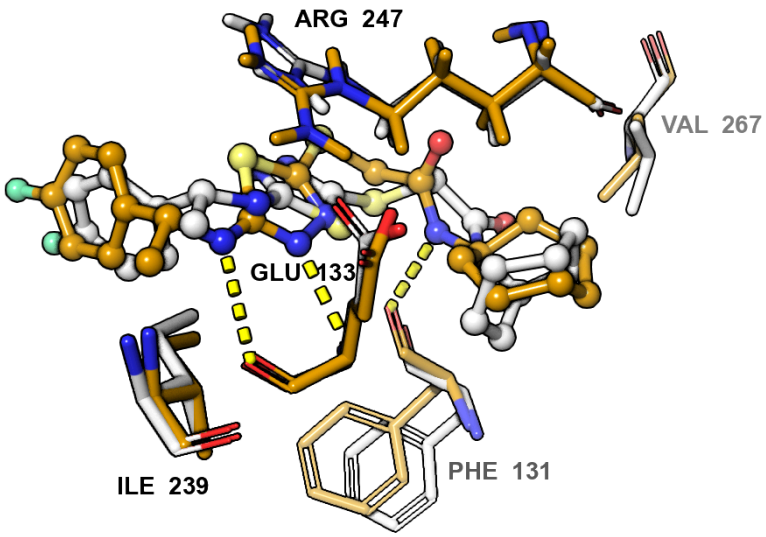

Figure 11: *Selected residues from the binding site of HQ19 in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose*

Similarly to its close relative HQ18, this pose stays below 2 Å across the board, with metaD-averages of rmsd remaining close to that of the MD-average. Free energy surfaces of metaD triplicates show a clear overlap, indicating a stable pose around 1.5 Å (see appendix). The clustering has given a single structure for the trajectory. The MD cluster has the central scaffold rotated on its axis compared to the IFD pose, with the amide group pointing towards the backbone of Arg247. The cyclohexane group is similarly rotated as the amide.

HQ19-1 in site B

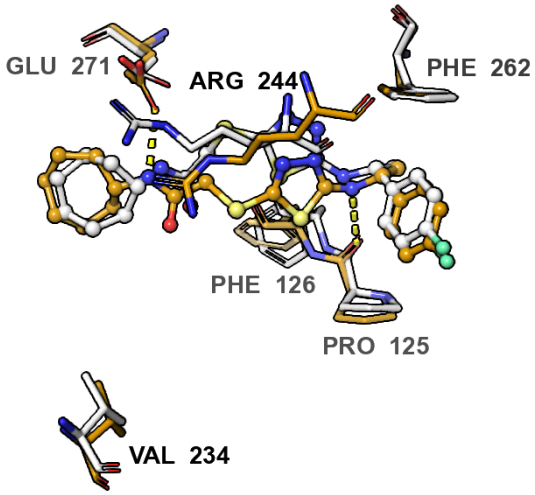

Figure 12: *Selected residues from the binding site of HQ19-1 in subcomplex B. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Hydrogen bonds are shown in yellow.*

MetaD-outcomes for this pose show a clear increase in rmsd after 9ns, with a continuing rise in averages after that point (see appendix). It maintains an average below 2Å for the first two calculated averages (5ns and 10ns, see Table 1) Note that the cyclohexane is on the opposite end of the scaffold compared to most other ligands. The alternative pose for site B, HQ19-2 has a scaffold arrangement that is more common. It has an MD-trajectory with a higher average, at 2.9Å. The MD trajectory of HQ19-1 is influenced by two alternative configurations. One is quite similar to the IFD-pose, while the other has the central scaffold shifted somewhat, with the heterocycle appearing to stabilise over Phe126. The fluorobenzene substituent remains more or less stable. Arg244 is firmly positioned above the amide carbonyl of the ligand.

HQ18 in site A

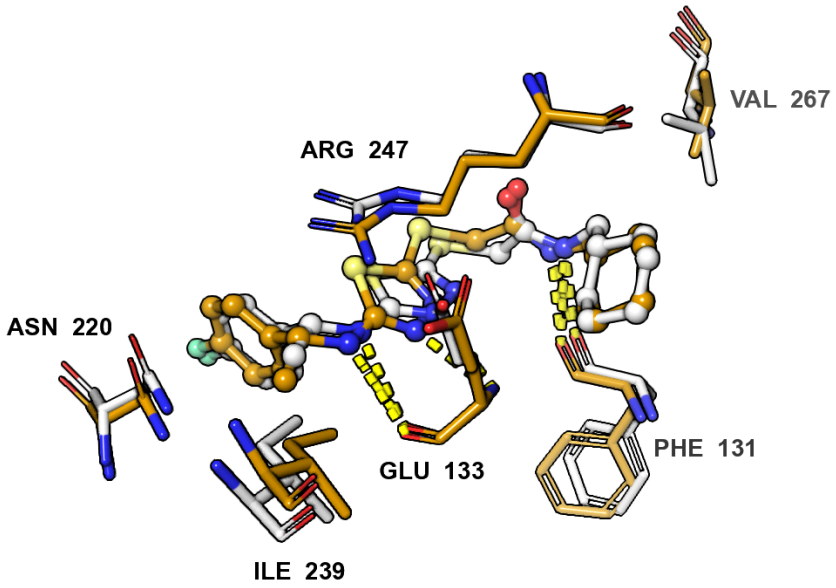

Figure 13: *The binding site of HQ18 in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose*

From the outcome of the simulations there is a clear indication of stability associated with this pose. The free energy surfaces are in agreement, and the metaD-averages remain low throughout. Due to the very stable trajectory of MD (see appendix), the clustering has resulted in a single group. Thus it is the typical arrangement for the entire 100ns simulation. It closely resembles that of the IFD-output, with a smaller shift in Arg247. The trajectory shows no major shifts in the binding mode. Asn220 occasionally orients itself towards the ligand fluorobenzene group.

HQ18 in site B

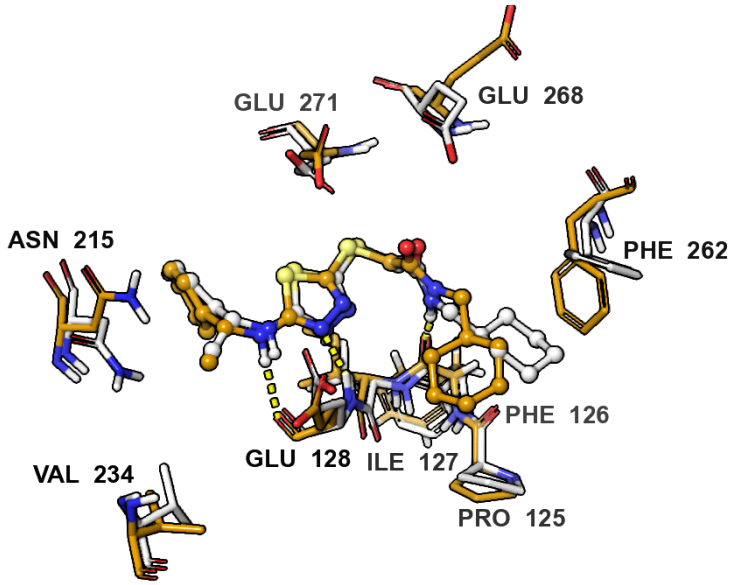

Figure 14: *Selected residues from the binding site of HQ18 in subcomplex B. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Hydrogen bonds are shown in yellow.*

The ligand displays remarkably low rmsd-averages in the metadynamics for site B as well. In two out of three runs, the ligand returns back to a low value, close to that of the MD simulation, after ~17ns (see appendix). As might be suspected from the names, this ligand is very similar in structure to HQ19, save for an additional carbon between the cyclohexane and amide groups. Compared to HQ19-1 this pose has the more common arrangement of the ligand scaffold, which matches with that seen in the crystal structure.

In its MD trajectory we can observe a rotation in Phe262 between MD cluster and IFD-output, as well as a movement of the cyclohexyl group in the general direction of Tyr140. Glu268 is shifted towards the ligand. The poses of IFD and MD cluster match quite well, with the only notable difference being the cyclohexyl group. Phe262 spends a significant part of the simulation pointed towards the ligand, which might explain the cyclohexyl shifting in a similar direction. Asn215 has shifted away from the ligand in the cluster, but oscillates between this and approaching the fluorobenzene substituent. Both Maestro and the hydrogen bonding count from the trajectory (Table 2) indicate strong interactions between the ligand and the backbones of four residues that appear below it in the figure (Pro125, Phe126, Ile127 and Glu128).

HQ11-2 in site A

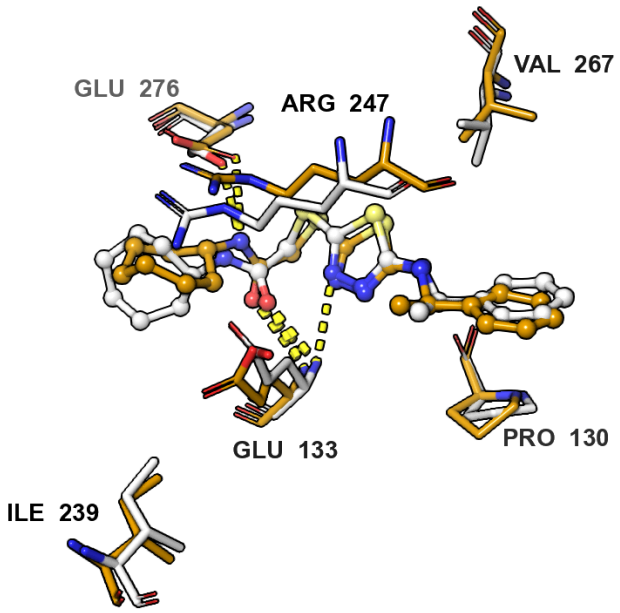

Figure 15: *The binding site of HQ11-2 in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose*

This is one out of two poses that were run for this ligand, with the other (HQ11-1) having such high rmsd-averages that it did not warrant further analysis.

This pose maintains a low metaD rmsd-average over the first 10ns of the simulations. It is another example of an "inverted" pose, where the benzene substituent is on the right side of the image. Compared to the IFD pose, the cluster shifts the 7-membered ring towards Arg247 to some extent.

HQ11-1 and HQ11-2 in site B

Neither of these poses justified closer analysis, given the unstable outcomes of metaD/MD (see table 1 and appendix). HQ11-1 has a remarkably good glide score, but the outcome of MD suggests that it didn't recognize a stable pose close to the one suggested by IFD.

HQ3R in site A

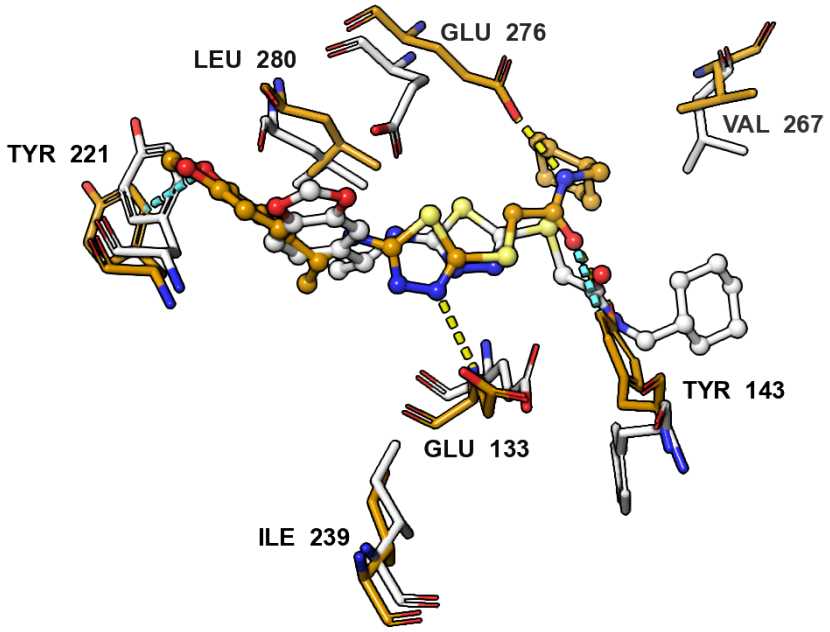

Figure 16: *Selected residues from the binding site of HQ3R in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Hydrogen bonds are shown in yellow, while aromatic hydrogen bonds are shown in teal.*

This pose shows some agreement in the free energy surfaces around 2Å, while metaD averages remain between 2 and 3Å for the 20ns triplicates. The MD run shifts to a pose around 4Å at ~4ns and stays there for the rest of the 100 ns simulation (see appendix). The 4.6Å average of MD is high, but the trajectory indicates a stable pose being found at that position. The dioxolane attached to the benzene group makes the outcome of this simulation more interesting. After the trajectory has shifted, this particular part of the ligand spends its time interacting with the backbones of Leu280 and Tyr221. The cyclohexane group on the other end of the ligand has moved quite a distance from the IFD pose, and seems to have pushed Met272 away from the ligand. Some mentioned residues were not included in the images in an attempt to declutter it.

HQ3R in site B

Due to its low predicted stability in the metadynamics, this pose has received no further attention.

HQ3S in site A

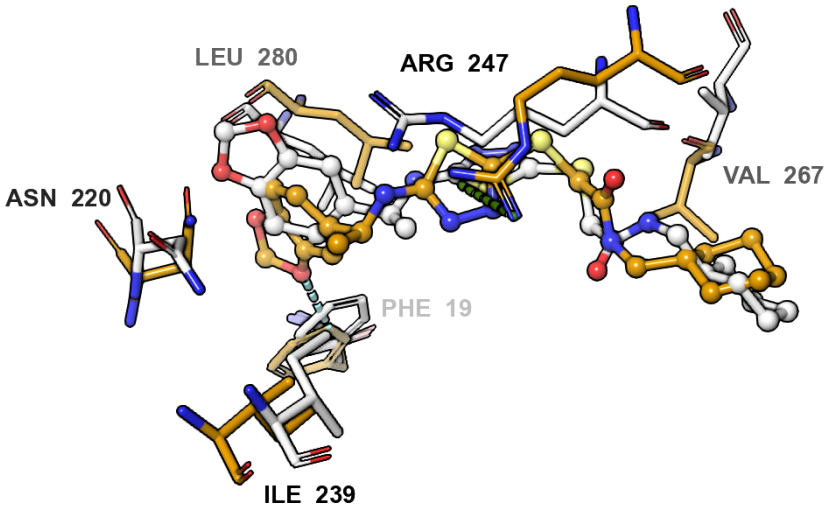

Figure 17: *Selected residues from the binding site of HQ33-S in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Aromatic-hydrogen bonds are shown in teal. π-cation interactions shown in green.*

This pose averages out above the usually required 2Å, but the averages from MD and metaD are quite similar (see Table 1). When compared to the R-isomer, the dioxolane group attached to the benzene is similarly interacting with Leu280 and Asn220. The Cyclohexane is more closely aligned with the IFD-pose, while the central part of the scaffold is rotated on its axis. There is a possible π-cation interaction with Arg247.

HQ3S-1 and HQ3S-2 in site B

This pose has no indicated stability in site B.

KP3A in site A

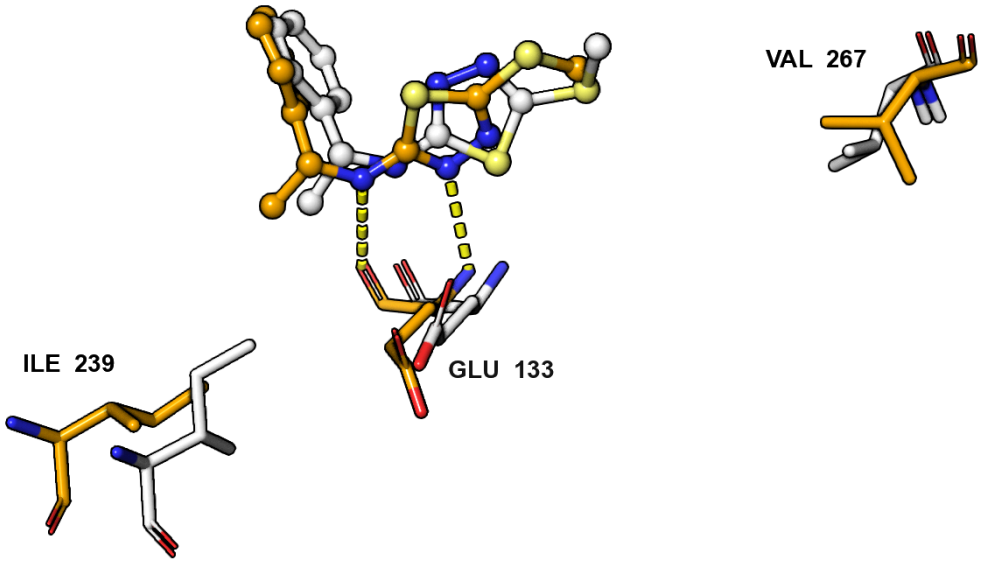

Figure 18: *Selected residues from the binding site of KP3A in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Hydrogen bonds are shown in yellow.*

Given its low glide scoring in IFD, it is surprising to see that the averages from metaD remain below 2Å throughout. The MD simulation gives a higher average, around 3.2Å. It should be mentioned that the apparent stability indicated by metaD could be an artefact related to how rmsd is calculated. A smaller ligand would likely yield a smaller rmsd given a relatively stable trajectory. Another possibility is that metaD has located a stable binding mode different from that of the MD trajectory. The clustering of MD returns a pose that has the central heterocycle rotated, but that otherwise matches with the IFD-pose. By the end of the trajectory it appears that the ligand is in the process of leaving its binding mode.

KP3A in site B

Here both the metadynamics and MD indicate an unstable pose (see Table 1)

KI48-1 in site A

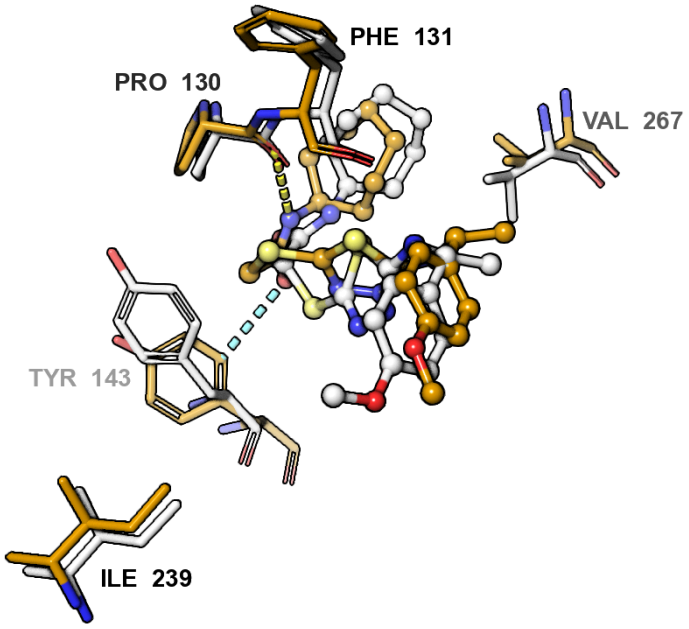

Figure 19: *Selected residues from the binding site of KI48-1 in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Hydrogen bonds are shown in yellow, while aromatic hydrogen bonds are shown in teal.*

This pose has metaD/MD averages in rmsd that are close enough to one another, and generally low enough to warrant interest (see Table 1). One of the metaD replicates seems to follow the trajectory of MD rather closely (see appendix). There is a contracted ligand pose, with the 7-membered ring bending towards Val267. Between the MD cluster and IFD-pose, the central heterocycle rotates and points towards the backbone of Phe131. The anisole (R-benzene-methoxy) substituent has shifted somewhat towards Glu276. There is a possible aromatic interaction between Tyr 142 and the amide carbonyl of the ligand.

KI48 in site B

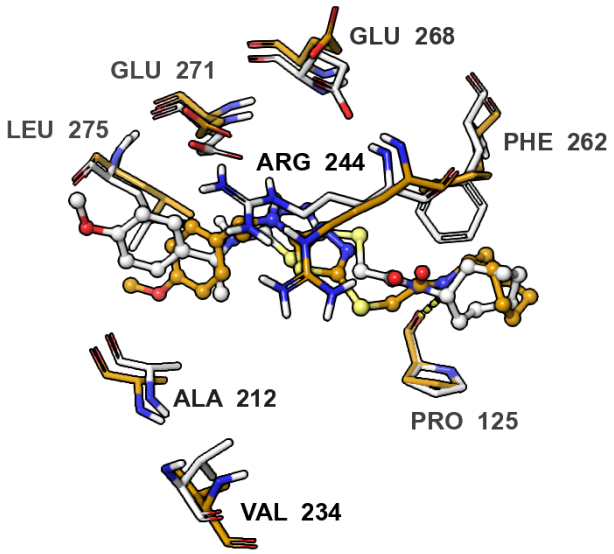

Figure 20: *Selected residues from the binding site of KI48 in subcomplex B. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Hydrogen bonds are shown in yellow.*

The course of MD is here characterised by a slight contraction of the structure that bends the methoxy-group downwards, towards Leu275. Glu268 and Arg244 are turned away from the ligand as compared with the IFD structure. It is an overall stable trajectory, with most of the deviations in MD-trajectory rmsd coming from adjustments of the cyclohexane and anisole substituents. This is a binding site that has the great majority of residues shifted somewhat in the same directions between the cluster and IFD. That could indicate some larger movement in the protein being responsible for the observed difference.

KP6A-1 in site A

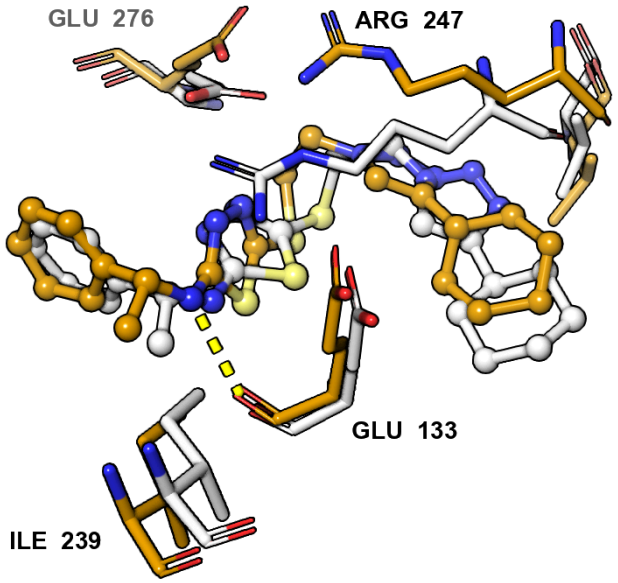

Figure 21: *Selected residues from the binding site of KP6A in subcomplex A. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Hydrogen bonds are shown in yellow.*

This run of metaD-triplicates seems to have two out of three agreeing on a stable pose close to 2Å (see FES-graph in appendix). The metaD-averages of rmsd begin to rise around 10ns. A somewhat stable pose appears to be indicated around the MD-average of 1.8Å. Compared to IFD, the MD-cluster contains a large shift in Arg247 towards the triazole-substituent and Glu276. The overall ligand pose is somewhat more contracted in the cluster. The triazole ring appears mostly locked in place during the MD trajectory.

KP6A in site B

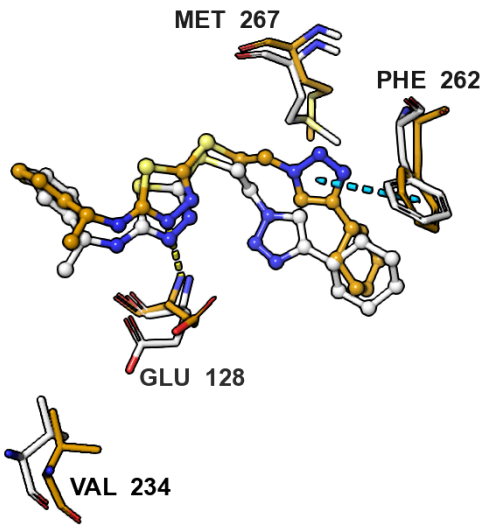

Figure 22: *Selected residues from the binding site of KP6A in subcomplex B. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the MD trajectory cluster. Hydrogen bonds are shown in yellow.* ***π-π*** *stacking is shown in blue.*

This ligand is quite stable over the course of the MD simulation, with much of the relative movements within the binding site stemming from interactions between the triazole-group and Met267. What is most notable about these simulations is that all three replicas of metaD return to a pose below 2Å by the end of 20ns, in spite of the bias being applied. This could indicate that a refined, stable pose has been found in close proximity to the IFD-output. Maestro indicates ***π-π*** stacking here with the non-conserved Phe262.

TCN-1 and TCN-2 in site A

Simulations of both poses for this ligand indicate that no stable pose has been found close to that suggested by IFD. It is possible that TCN-1 has found a stable pose around the MD-rmsd average of 3Å, given the resembling values of metaD-averages. Two of the metaD-triplicates stay close to the MD trajectory (3Å) for most of the 20 ns.

TCN in site B

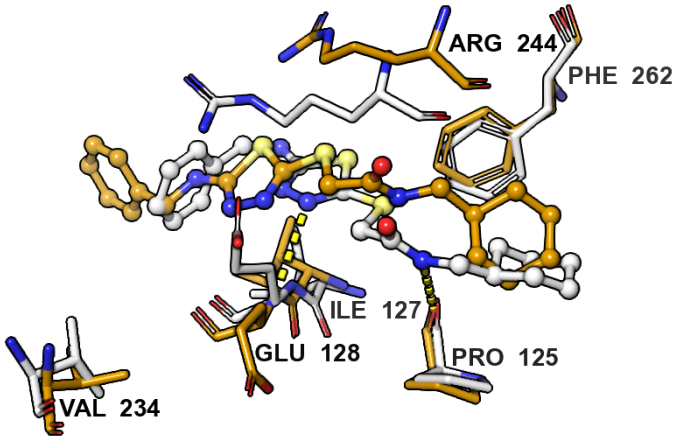

Figure 23: *Selected residues from the binding site of TCN in subcomplex B. Orange: MD trajectory cluster. White: IFD-pose. All interactions displayed are within the IFD binding site. Hydrogen bonds are shown in yellow.*

The main contributors to variation of ligand rmsd in the MD simulation can be seen to result from rotations of the benzene and cyclohexane rings at the ends of the molecule. The central part of the scaffold remains largely in place. There are frequent reorientations of the cyclohexane group, some of which could possibly be attributed to interactions with the non-conserved Phe262. The amide group remains stable more or less throughout. This remains true for the heterocycle and amine substituents in the central part of the ligand.

4.3 Hydrogen Bonding

As an additional inquiry into which binding site residues play the dominant roles in ligand interactions, the hydrogen bonds over the course of all above visualised trajectories have been mapped in Table 2. Any interactions lasting less than 1% of the trajectory have been left out. Interactions have been categorised by main-chain (M) and side-chain (S).

Table 2: *Hydrogen bonds of selected MD trajectories, for residues with a hydrogen bonding occupancy above 1%. M: Main chain S: Side chain.Note that in GluN2A, AA33 has a trajectory of 500 ns instead of the normal 100 ns. Brighter colour indicates a greater persistence. In the case of Glu 133/128, the two numbers occupying some cells are two different hydrogen bonds between main-chain and ligand.*

Ligand	Occupancy (%):	0	1 --> 9	10 --> 29	30 --> 49	50 --> 69	70&
Site A:	**ARG 247**	**GLU 133** M	**PHE 131** M	**GLU 276** S	**PRO 130** M	**ASN 220** S	**TYR 143** S	**GLY 249** M
AA33	34 S	58 & 36	75			1		
AA34-1	7 M	54		52				
HQ09-1	10 M		72					
HQ11-2	6 M	14		78	34			
HQ18		57 & 38	85					
HQ19		65 & 36	90					
HQ3R		10	8	59			82	
HQ3S-2	10 S	6	18	11	12			
KI48-1			4		30			6
KP3A		35 & 10						
KP6A		40 & 10						
Site B:	**ARG 244**	**GLU 128** M	**PHE 126** M	**GLU 271** S	**PRO 125** M	**ASN 215** S	**TYR 140**	
AA33		10		36	1	2		
AA34	2 S	26		1	66			
HQ09	16 S	54 & 34	63			1		
HQ18		37 & 32	74					
HQ19-1		2 & 3	28	60	33			
KI48				32	39			
KP6A-1		8 & 5						
TCN		10		2				

If ligands are compared between complexes A and B, there are several cases of notable shifts in hydrogen bonding. In AA33 the vast majority of interactions disappear when going from site A to B, while a previously non-existent one appears with the side-chain of Glu 271. AA34 sees a highly persistent bond (66%) with Pro 125 (B). HQ09 introduces a significant new interaction with Glu 128 (B) while maintaining the ones seen in A. HQ18 seems to rely on the same hydrogen-bonding pattern in both A and B. Both HQ19 and KI48 give new importance to Glu271 in site B.

On the whole it is mostly the same residues that provide significant hydrogen bonding to these binding modes. Especially notable are Glu 133/128 and Phe 131/126. The backbones of these residues seem to play an integral role in stabilising the central scaffolding that is shared by all ligands. From the trajectories that have been looked at more closely these residues appear to interact primarily with the central heterocycle and the amide-group. It should also be said that there are interaction types not represented by Table 2, such as hydrophobic interactions, aromatic hydrogen bonds and π-π stacking.

4.4 Discussion

This final section will focus on a more general discussion of the results and methods, beginning with the level of agreement between IFD and MD/metaD.

It is on the whole interesting to see that there is a tendency for many MD-trajectories to "visit" the poses that have been suggested by induced fit docking. This is not something that can be discerned from the presented results, but becomes clear from directly observing the paths taken by MD simulations. The clusters rarely match with IFD, since these visits come at a low frequency compared to some other, often similar pose. MD occasionally returning the ligand to its original position indicates doubly that there exists a free energy minimum there, which likely contributes to the binding mode of the ligand. When both methods identify the same binding mode it hints that they share some understanding of the underlying physics. There are other hints in the trajectories that point to the opposite, one of these is partially visible in the binding site depictions in section 4.2. It is rare for Arg247/244 to align with the position it is assigned by IFD. Ths residue is the only one central to the binding site that belongs to the N1 subunit, which is the same for GluNA and GluNB. In most trajectories there is an immediate and persistent realignment of this residue. It would seem that the two methods have different ideas. Table 2 tells us that it has a recurring role in hydrogen bonding with many ligands, while the MD trajectories tend to position this Arginine side-chain somewhere above the heterocycle or the amide group of the central scaffolding. Binding modes produced by IFD don't seem to capture this behaviour, hinting that its top-ranking suggestions lend weight to other interactions. All in all, disparities between IFD and MD could come down to MD being able to capture slower modes of the system. Adjustments in the protein backbone might take longer, or require more movement than what is captured in IFD. It could be such adjustments that lead to a new dominating binding mode, which would reduce the frequency of visits to IFD-suggestions. Another possibility is that parameters like topologies and sensitivities to interactions in the force fields are different enough between these methods to produce consistently different outcomes. The cluster-IFD columns in Table 1 present adjustments to the binding site heavy-atoms that are consistently higher than the 0.7Å seen for AA33 in the crystal structure. If the methods have a large overlap in their physics, then we could expect that they would occasionally agree "almost completely" about the arrangement of the binding site. There are cases where this is true for the ligands. For HQ18 in

site A and HQ09 in site B, IFD-poses and MD clusters overlap almost entirely. In the interest of robustness, these observations should be formalised numerically, in a way that is less susceptible to the interpretation of the beholder. This divergence between the models could possibly be made simpler by redefining the binding site based on the residues that have prominent interactions with the ligand. One interesting approach to determining the importance of specific residues would be in *Free Energy Perturbation* (FEP) (Fratev and Sirimulla, 2019). The FEP-approach has been quite successful in drug discovery, and can be used to examine changes in free energy when one residue is replaced with another, an in-silico mutation study of sorts.

Comparisons between models highlight the absence of training data for this specific application of IFD-MD. Even if there is a strong case to be made for the method as a general approach (see section 2.7), there is still room for improvement when it comes to building confidence in this method. A simple way of adding confidence would be to include a set of crystal structures as training data, and redock already present ligands using the same IFD-MD methodology. No number of crystal structures can overcome the weaknesses inherent to that approach. The static nature of a crystal structure provides little insight into the functional dynamics of molecular interactions. At the end of the day we must grapple with the incompleteness of our models, while recognising the need for modelling molecular interactions in a non-static fashion.

There is a consideration to be had in that the paths taken by the metadynamics haven't been observed closely. It is possible that these never visited some binding modes suggested by IFD. With the right set-up, it would be possible to investigate if IFD-suggestions are represented as local minima on the free energy surfaces mapped out by metaD. Below a certain cutoff of rmsd two poses could be thought of as the same.

Drawing conclusions from the metadynamics would be easier given a standardised treatment of the non-converged trajectories of the collective variable. It is a recurring problem within non-converged molecular dynamics to decide what information can be interpreted from the data. Choosing a cutoff presents a challenge, as it allows one to cherry pick the approach to the method, which could be leveraged to produce wanted outcomes. In a recent paper Antoniu Bjola and Matteo Salvalaglio (2024) present a standardised approach to using *Mean Force Integration* when evaluating free energy surfaces reconstructed from metaD. Their method intends to improve the convergence of thermodynamic properties like the free energy, which is different from the non-converged metaD simulations here. However, a weighted approach to CV-values would likely assist in simplifying the framing of simulation outcomes. As was brought up in section 2.3, the depths of free energy minimums is what is essential to the current approach.

One approach that could simplify things would be to introduce a "wall" into the metadynamics. Essentially a distance beyond which the ligand would be forced back. This runs the risk of introducing artefacts, since the wall would result in unrealistic dynamics. It could be helpful in

that it would constrict the range of rmsd-values seen to a predetermined range. The strength of this use of IFD-MD could definitely be improved on by narrowing down the investigation to a few ligands of special interest. Given similar resources it should be possible to take a group of ~5 ligands and rerun this method with more exhaustive simulations. These could be selected from ligands in this group that have indicated interactions with the non-converged residues of the binding sites, such as KP6A in site B. If 10 metaD-replicates are used instead, as is done by Clark et al. (2016), one should begin to see the free energy surfaces converge on clear energy minimums. Running the method again from the top would also provide an opportunity to standardise selection of IFD-poses. The choices made here relate to the investigation of the "inverse" binding modes seen in 4 poses displayed in section 4.2. As was mentioned before, these seem to display stable binding modes, and it would be interesting to know whether they contribute to the actual binding of these ligands. Answering that would require additional crystal structures or simulating the entry of the ligands into the binding site.

For improvements in future applications of IFD, the clusters produced here could be assembled into a group of binding sites for each complex. These groups of clusters could then be used as binding sites in *ensemble docking*. Since these structures have been adjusted via MD to accommodate some ligands that are known binders, the docking of future analogues could result in stable binding modes faster when applied to the ensemble.

4.5 Conclusions

This work has at its heart revolved around the construction of a method. It has taken inspiration from existing publications of IFD-MD and provided a unique combination of software to test this method across the boundaries of different platforms. That was easier said than done, but as far as we can see the method works as intended, and has led to some qualified guesswork regarding the stability of binding modes in 15 ligands. When correctly understood, these results can provide a path forward in choosing which ligands to investigate further, whether that be through rational design of analogues, computational methods or measurements of ion-channel permeability in NMDA-receptors.

There are plenty of assumptions that have to be made for biochemistry to translate into molecular dynamics. It is increasingly true that the real difficulties are in building the models and interpreting the outcomes. Executing fairly advanced simulation really only requires a modern GPU and the open-source softwares supplied by researchers like the developers of Gromacs. This increasing accessibility promises interesting developments, and is a great opportunity for molecular dynamics and its related methods. Nevertheless, it remains non-trivial to draw relevant conclusions from the data that is produced. It relies on a merger of knowledge that spans as many fields as the many assumptions that have to be made to design relevant simulations. No matter the underlying complexities there are many successful examples of these methods being applied to advance the understanding of complicated phenomena. It is possible to apply methods like

IFD-MD towards developing meaningful insights into systems as complex as the glutamate receptors. Wielded in the right hands, dynamics-simulations are a valuable addition to experimental designs, providing their own, unique view into the world as it exists on the molecular scale.

5. Sources

Antoniu Bjola and Matteo Salvalaglio (2024). Estimating Free-Energy Surfaces and Their Convergence from Multiple, Independent Static and History-Dependent Biased Molecular-Dynamics Simulations with Mean Force Integration. *Journal of Chemical Theory and Computation*, 20(13). doi:https://doi.org/10.1021/acs.jctc.4c00091.

Armstrong, N. and Gouaux, E. (2000). Mechanisms for Activation and Antagonism of an AMPA-Sensitive Glutamate Receptor. *Neuron*, 28(1), pp.165–181. doi:https://doi.org/10.1016/s0896-6273(00)00094-5.

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), pp.1752–1780. doi:https://doi.org/10.1002/jcc.20292.

Barducci, A., Bonomi, M. and Parrinello, M. (2011). Metadynamics. *WIREs Computational Molecular Science*, 1(5), pp.826–843. doi:https://doi.org/10.1002/wcms.31.

Bussi, G., Alessandro Laio and Pratyush Tiwary (2018). Metadynamics: A Unified Framework for Accelerating Rare Events and Sampling Thermodynamics and Kinetics. *Handbook of Materials Modeling*, pp.1–31. doi:https://doi.org/10.1007/978-3-319-42913-7_49-1.

Bussi, G. and Branduardi, D. (2015). Free-Energy Calculations with Metadynamics: Theory and Practice. *Reviews in Computational Chemistry*, 28(1). doi:https://doi.org/10.1002/9781118889886.ch1.

Clark, A.J., Tiwary, P., Borrelli, K., Feng, S., Miller, E.B., Abel, R., Friesner, R.A. and Berne, B.J. (2016). Prediction of Protein–Ligand Binding Poses via a Combination of Induced Fit Docking and Metadynamics Simulations. *Journal of Chemical Theory and Computation*, 12(6), pp.2990–2998. doi:https://doi.org/10.1021/acs.jctc.6b00201.

Coveney, P.V. and Wan, S. (2016). On the calculation of equilibrium thermodynamic properties from molecular dynamics. *Physical Chemistry Chemical Physics*, 18(44), pp.30236–30240. doi:https://doi.org/10.1039/c6cp02349e.

Daura, X., Gademann, K., Jaun, B., Seebach, D., van Gunsteren, W.F. and Mark, A.E. (1999). Peptide Folding: When Simulation Meets Experiment. *Angewandte Chemie International Edition*, 38(1-2), pp.236–240. doi:https://doi.org/10.1002/(sici)1521-3773(19990115)38:1/2%3C236::aid-anie236%3E3.0.co;2-m.

Eldridge, M.D., Murray, C.W., Auton, T.R., Paolini, G.V. and Mee, R.P. (1997). Empirical scoring functions: I. The development of a fast empirical scoring function to estimate the binding affinity of ligands in receptor complexes. *Journal of Computer-Aided Molecular Design*, 11(5), pp.425–445. doi:https://doi.org/10.1023/a:1007996124545.

Fowler, M. (2014). *Hamilton's Equations*. [online] Virginia.edu. Available at: https://galileoandeinstein.phys.virginia.edu/7010/CM_06_HamiltonsEqns.html [Accessed 26 Sep. 2024].

Fratev, F. and Sirimulla, S. (2019). An Improved Free Energy Perturbation FEP+ Sampling Protocol for Flexible Ligand-Binding Domains. *Scientific Reports*, 9(1). doi:https://doi.org/10.1038/s41598-019-53133-1.

Friesner, R.A., Banks, J.L., Murphy, R.B., Halgren, T.A., Klicic, J.J., Mainz, D.T., Repasky, M.P., Knoll, E.H., Shelley, M., Perry, J.K., Shaw, D.E., Francis, P. and Shenkin, P.S. (2004). Glide: A New Approach for Rapid, Accurate Docking and Scoring. 1. Method and Assessment of Docking Accuracy. *Journal of Medicinal Chemistry*, 47(7), pp.1739–1749. doi:https://doi.org/10.1021/jm0306430.

Friesner, R.A., Murphy, R.B., Repasky, M.P., Frye, L.L., Greenwood, J.R., Halgren, T.A., Sanschagrin, P.C. and Mainz, D.T. (2006). Extra Precision Glide: Docking and Scoring Incorporating a Model of Hydrophobic Enclosure for Protein−Ligand Complexes. *Journal of Medicinal Chemistry*, 49(21), pp.6177–6196. doi:https://doi.org/10.1021/jm051256o.

Gareth A. Tribello, Massimiliano Bonomi, Davide Branduardi, Carlo Camilloni, and Giovanni Bussi. Plumed 2: New feathers for an old bird. *Comput. Phys. Commun.*, 185(2):604–613, 2014.

G. Brent Dawe, Musgaard, M., Aurousseau, M., Nayeem, N., Green, T., Biggin, P.C. and Bowie, D. (2016). Distinct Structural Pathways Coordinate the Activation of AMPA Receptor-Auxiliary Subunit Complexes. *Neuron*, [online] 89(6), pp.1264–1276. doi:https://doi.org/10.1016/j.neuron.2016.01.038.

Götz, A.W., Williamson, M.J., Xu, D., Poole, D., Le Grand, S. and Walker, R.C. (2012). Routine Microsecond Molecular Dynamics Simulations with AMBER on GPUs. 1. Generalized Born. *Journal of Chemical Theory and Computation*, 8(5), pp.1542–1555. doi:https://doi.org/10.1021/ct200909j.

Hansen, K.B., Wollmuth, L.P., Bowie, D., Furukawa, H., Menniti, F.S., Sobolevsky, A.I., Swanson, G.T., Swanger, S.A., Greger, I.H., Nakagawa, T., McBain, C.J., Jayaraman, V., Low, C.-M., Dell'Acqua, M.L., Diamond, J.S., Camp, C.R., Perszyk, R.E., Yuan, H. and Traynelis, S.F. (2021). Structure, Function, and Pharmacology of Glutamate Receptor Ion Channels. *Pharmacological Reviews*, 73(4), pp.298–487. doi:https://doi.org/10.1124/pharmrev.120.000131.

Hansen, K.B., Yi, F., Perszyk, R.E., Furukawa, H., Wollmuth, L.P., Gibb, A.J. and Traynelis, S.F. (2018). Structure, function, and allosteric modulation of NMDA receptors. *The Journal of General Physiology*, 150(8), pp.1081–1105. doi:https://doi.org/10.1085/jgp.201812032.

Harder, E., Damm, W., Maple, J., Wu, C., Reboul, M., Xiang, J.Y., Wang, L., Lupyan, D., Dahlgren, M.K., Knight, J.L., Kaus, J.W., Cerutti, D.S., Krilov, G., Jorgensen, W.L., Abel, R. and Friesner, R.A. (2015). OPLS3: A Force Field Providing Broad Coverage of Drug-like Small Molecules and Proteins. *Journal of Chemical Theory and Computation*, 12(1), pp.281–296. doi:https://doi.org/10.1021/acs.jctc.5b00864.

Hilger, D., Masureel, M. and Kobilka, B.K. (2018). Structure and dynamics of GPCR signaling complexes. *Nature Structural & Molecular Biology*, [online] 25(1), pp.4–12. doi:https://doi.org/10.1038/s41594-017-0011-7.

Hollingsworth, S.A. and Dror, R.O. (2018). Molecular Dynamics Simulation for All. *Neuron*, [online] 99(6), pp.1129–1143. doi:https://doi.org/10.1016/j.neuron.2018.08.011.

Jacobson, M.P., Kaminski, G.A., Friesner, R.A. and Rapp, C.S. (2002). Force Field Validation Using Protein Side Chain Prediction. *The Journal of Physical Chemistry B*, 106(44), pp.11673–11680. doi:https://doi.org/10.1021/jp021564n.

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), pp.926–935. doi:https://doi.org/10.1063/1.445869.

Kirkpatrick, P. (2004). Gliding to success. *Nature Reviews Drug Discovery*, 3(4), pp.299–299. doi:https://doi.org/10.1038/nrd1364.

Lewerenz, J. and Maher, P. (2015). Chronic Glutamate Toxicity in Neurodegenerative Diseases—What is the Evidence? *Frontiers in Neuroscience*, [online] 9. doi:https://doi.org/10.3389/fnins.2015.00469.

Li, C.-T., Yang, K.-C. and Lin, W.-C. (2019). Glutamatergic Dysfunction and Glutamatergic Compounds for Major Psychiatric Disorders: Evidence From Clinical Neuroimaging Studies. *Frontiers in Psychiatry*, 9. doi:https://doi.org/10.3389/fpsyt.2018.00767.

Li, J., Abel, R., Zhu, K., Cao, Y., Zhao, S. and Friesner, R.A. (2011). The VSGB 2.0 model: A next generation energy model for high resolution protein structure modeling. *Proteins: Structure, Function, and Bioinformatics*, 79(10), pp.2794–2812. doi:https://doi.org/10.1002/prot.23106.

Lindorff-Larsen, K., Maragakis, P., Piana, S., Eastwood, M.P., Dror, R.O. and Shaw, D.E. (2012). Systematic Validation of Protein Force Fields against Experimental Data. *PLoS ONE*, 7(2), p.e32131. doi:https://doi.org/10.1371/journal.pone.0032131.

Lindorff-Larsen, K., Piana, S., Palmo, K., Maragakis, P., Klepeis, J.L., Dror, R.O. and Shaw, D.E. (2010). Improved side-chain torsion potentials for the Amber ff99SB protein force field. *Proteins: Structure, Function, and Bioinformatics*, [online] 78(8), p.NA-NA. doi:https://doi.org/10.1002/prot.22711.

Miller, E.B., Murphy, R.B., Sindhikara, D., Borrelli, K.W., Grisewood, M.J., Ranalli, F., Dixon, S.L., Jerome, S., Boyles, N.A., Day, T., Ghanakota, P., Mondal, S., Rafi, S.B., Troast, D.M., Abel, R. and Friesner, R.A. (2021). Reliable and Accurate Solution to the Induced Fit Docking Problem for Protein–Ligand Binding. *Journal of Chemical Theory and Computation*, 17(4), pp.2630–2639. doi:https://doi.org/10.1021/acs.jctc.1c00136.

Morris, R.G.M. (2013). NMDA receptors and memory encoding. *Neuropharmacology*, [online] 74, pp.32–40. doi:https://doi.org/10.1016/j.neuropharm.2013.04.014.

Promoting transparency and reproducibility in enhanced molecular simulations. (2019). *Nature Methods*, 16(8), pp.670–673. doi:https://doi.org/10.1038/s41592-019-0506-8.

Reiner, A. and Levitz, J. (2018). Glutamatergic Signaling in the Central Nervous System: Ionotropic and Metabotropic Receptors in Concert. *Neuron*, 98(6), pp.1080–1098. doi:https://doi.org/10.1016/j.neuron.2018.05.018.

Schrödinger Documentation. (2024). *Documentation*. [online] Available at: https://learn.schrodinger.com/private/edu/release/current/Documentation/html/Home.htm [Accessed 17 Oct. 2024].

Schrödinger. (2024). *Prime | Schrödinger*. [online] Available at: https://www.schrodinger.com/platform/products/prime/ [Accessed 22 Oct. 2024].

Sherman, W., Day, T., Jacobson, M.P., Friesner, R.A. and Farid, R. (2006). Novel Procedure for Modeling Ligand/Receptor Induced Fit Effects. *Journal of Medicinal Chemistry*, 49(2), pp.534–553. doi:https://doi.org/10.1021/jm050540c.

Sousa da Silva, A.W. and Vranken, W.F. (2012). ACPYPE - AnteChamber PYthon Parser interfacE. *BMC Research Notes*, 5(1), p.367. doi:https://doi.org/10.1186/1756-0500-5-367.

Stone, J.E., Hallock, M.J., Phillips, J.C., Peterson, J.R., Luthey-Schulten, Z. and Schulten, K. (2016). Evaluation of Emerging Energy-Efficient Heterogeneous Computing Platforms for Biomolecular and Cellular Simulation Workloads. *Europe PMC (PubMed Central)*. [online] doi:https://doi.org/10.1109/ipdpsw.2016.130.

Tribello, G.A., Bonomi, M., Branduardi, D., Camilloni, C. and Bussi, G. (2014). PLUMED 2: New feathers for an old bird. *Computer Physics Communications*, 185(2), pp.604–613. doi:https://doi.org/10.1016/j.cpc.2013.09.018.

Vojtěch Spiwok, Kurečka, M. and Aleš Křenek (2022). Collective Variable for Metadynamics Derived From AlphaFold Output. *Frontiers in Molecular Biosciences*, 9. doi:https://doi.org/10.3389/fmolb.2022.878133.

Weiner, P.K. and Kollman, P.A. (1981). AMBER: Assisted model building with energy refinement. A general program for modeling molecules and their interactions. *Journal of Computational Chemistry*, 2(3), pp.287–303. doi:https://doi.org/10.1002/jcc.540020311.

Wikipedia. (2020). *Helmholtz free energy*. [online] Available at: https://en.wikipedia.org/wiki/Helmholtz_free_energy.