

Alchemical Absolute Protein-Ligand Binding Free Energies for Drug Design

Y. Khalak¹, G. Tresadern², M. Aldeghi³, H. M. Baumann⁴, D. L. Mobley^{4,5}, B. L. de Groot¹, and V. Gapsys^{1,*}

¹ Computational Biomolecular Dynamics Group, Department of Theoretical and Computational Biophysics, Max Planck Institute for Biophysical Chemistry, D-37077, Göttingen, Germany.

² Computational Chemistry, Janssen Research & Development, Janssen Pharmaceutica N. V., Turnhoutseweg 30, 2340, Beerse, Belgium.

³ MIT Department of Chemical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

⁴ Department of Pharmaceutical Sciences, University of California, Irvine, CA 92697, USA.

⁵ Department of Chemistry, University of California, Irvine, CA 92697, USA.

* vgapsys@gwdg.de

September 18, 2021

Abstract

The recent advances in relative protein-ligand binding free energy calculations have shown the value of alchemical methods in drug discovery. Accurately assessing absolute binding free energies, although highly desired, remains a challenging endeavour, mostly limited to small model cases. Here, we demonstrate accurate first principles based absolute binding free energy estimates for 128 pharmaceutically relevant targets. We use a novel rigorous method to generate protein-ligand ensembles for the ligand in its decoupled state. Not only do the calculations deliver accurate protein-ligand binding affinity estimates, but they also provide detailed physical insight into the structural determinants of binding. We identify subtle rotamer rearrangements between apo and holo states of a protein that are crucial for binding. When compared to relative binding free

energy calculations, obtaining absolute binding free energies is considerably more challenging in large part due to the need to explicitly account for the protein in its apo state. In this work we present several approaches to obtain apo state ensembles for accurate absolute ΔG calculations, thus outlining protocols for prospective application of the methods for drug discovery.

1 Introduction

Computational techniques for estimating relative differences in protein-ligand binding free energy have now reached remarkable accuracy. Relative binding free energy calculations over a large range of protein-ligand complexes have shown average agreement with experiment to be within 1 kcal/mol (4.184 kJ/mol)^{16,28,52}. These methods have become mature and reliable enough to be included in industrial drug discovery and lead optimization

pipelines^{10,26,35,40}. A substantial limitation of this approach, however, is the requirement for the ligands to be structurally similar to each other: the predictive power decreases for ligands with different scaffolds or binding poses. Evaluation of novel ligand classes, therefore, requires a prior experimental absolute binding free energy as a reference for each new class of mutually similar ligands. Thus, the next qualitative leap for the field of first principles based protein-ligand affinity estimation encompasses the reliable and accurate prediction of absolute binding free energies.

The calculation of relative binding free energies is relatively easy in comparison. The bound ligands are confined to the binding site and only the small subset of atoms that differ between two ligands need to be perturbed. In contrast, absolute binding free energy calculations decouple the entire ligand, meaning it is in principle free to explore the whole simulation volume. Early work on the topic explored various ways of restraining the decoupled ligand and taking into account the resulting contribution to the free energy^{7,18,20,54}. The approach introduced by Boresch et al⁷ has emerged as a rigorous way to resolve this issue via orthonormal relative restraints between the ligand and the protein^{2,3}.

Another challenge for the absolute binding free energy calculations is posed by the need to explicitly sample the apo state of the protein, i.e. the protein without the bound ligand. As this state may substantially differ from the ligand bound (holo) state, the simulation method needs to be capable of capturing the free energy differences between the protein conformers. Non-equilibrium (NEQ) free energy calculations present an elegant solution to this challenge. Such calculations determine the free energy difference by performing rapid out-of-equilibrium ligand coupling/decoupling transitions initialized from the equilibrium protein apo and holo ensembles (Fig S1). This allows one to explicitly include the different apo and holo end-states into the same calculation¹⁷. Several recent applications of the NEQ approach on model host-guest systems showed promising results for the calculation of absolute

binding free energies^{5,27,37}.

The NEQ approach does not offer a free lunch in the sense that the relevant conformations still need to be sampled in the end-state ensembles¹⁷. Compared to the more popular free energy perturbation (FEP) series of methods^{24,51,56}, though, it does offer several advantages in terms of computational efficiency. Namely, such sampling needs to be performed only for physical end-states and can be done with plain molecular dynamics or, if desired, it can also be augmented with enhanced sampling methodologies in a straight forward manner³⁶. Secondly, the out-of-equilibrium portion of the approach, which accounts for the majority of the compute time, is highly parallelizable, requiring no information exchange between individual simulations, unlike modern FEP approaches with replica exchange^{24,51}. Furthermore, the NEQ approach allows for initialization of the two end-states with the distinct apo and holo protein structures, which facilitates obtaining reliable equilibrium ensembles for the cases where experimental structures are available. For an equilibrium FEP approach, incorporation of different conformers in a single ΔG estimation would require decision on the mixing rule for seeding the starting structures¹⁹ and potentially Hamiltonian replica exchanges would be needed to achieve convergence. Finally, when comparing different ligands, protein mutations, or conformational states NEQ allows for reuse of existing equilibrium sampling of end-states, e.g. the same apo state can be used for assessing affinities of different ligands.

In the current work we use the NEQ approach to demonstrate the feasibility of accurate absolute binding free energy calculations for a large number of protein-ligand systems, showing accuracy on par with the relative binding free energy estimates. To achieve this, we introduce methodological advancements that allow for an efficient treatment of the ligand in its decoupled state and careful considerations of the protein in its apo state. This allows for identification of protein states that have a drastic effect on ligand binding affinity, such as e.g. a flip of a single amino acid rotamer. Our calculation strategy also allows identify-

ing the most representative structure for a protein’s apo state for the cases where multiple likely candidates (structures in their local free energy minima, X-ray structures) are available.

2 Results

In this study, we have used an alchemical non-equilibrium free energy calculation approach to calculate absolute protein-ligand binding free energies for 128 complexes. We have developed a novel way of treating the decoupled state of the ligand (see Methods section for details). The large set of investigated systems allows us to have an extensive evaluation of the accuracy that can be achieved with the first principles based calculations. Figure 1 shows the calculated values for the binding free energy plotted against the experimental measurements. When compared to the experimental values, the absolute unsigned error (AUE) of 4.9 ± 0.5 kJ/mol (1.2 ± 0.1 kcal/mol) only marginally exceeds the state-of-the-art accuracy threshold of 1 kcal/mol achievable for relative binding free energy calculations. Accuracies for jnk1 and p38 α are exceptionally good with AUEs of 3.0 ± 0.8 and 3.1 ± 0.7 kJ/mol (0.7 ± 0.2 kcal/mol for both), respectively.

Some systems, however, have a considerably lower accuracy (AUE of 10.8 ± 1.5 kJ/mol for tyk2, 5.5 ± 1.4 kJ/mol for pde2), revealing a particular challenge for affinity estimation in these systems. An accurate evaluation of the offset in the ΔG is critical for obtaining reliable absolute ΔG values: inaccuracies in this case manifest as large shifts of the calculated values with respect to the experimental measurements, e.g. tyk2 in Fig 1. Interestingly, even such offsets do not significantly deteriorate the relative free energy difference estimates (Fig S2).

We have identified this effect to be a consequence of the inadequate representation of the protein in its apo state. While the apo state is not considered in relative free energy calculations, assessment of absolute free energies needs to explicitly account for it. In the following analysis we demonstrate how failure to capture the free energy differences between the apo and

holo protein states affects the absolute binding free energy calculation accuracy.

2.1 Apo and Holo states

For the situations where protein rearrangements are required upon ligand binding, sufficient sampling of the two end states may present a considerable challenge. An accurate quantification of the process of ligand binding to an apo protein and forming a stable holo state requires correctly estimating not only the component of the free energy originating from the ligand interaction with the protein, but also the difference between the apo and holo protein states.

Non-equilibrium free energy calculations offer a particularly convenient approach for the computation of binding affinities, as both states, apo and holo, can be explicitly considered in a single simulation¹⁷. The alchemical ligand decoupling transitions can be started from a holo conformer ensemble, while ligand coupling transitions can start from an apo ensemble.

Among the protein-ligand complexes investigated in this work, 6 out of 7 systems have both their apo and holo structures resolved by means of x-ray crystallography. We have probed two methods of calculating the binding ΔG value: firstly, removing the ligand from the holo state and treating the obtained structure as an apo state. For the second approach we used the crystallographically resolved apo structure directly. Overall, there is a large and significant improvement in the calculated binding ΔG accuracy when an experimentally defined apo state is considered explicitly (Fig 2, S3). A substantial improvement in the AUE (from 7.1 ± 0.6 to 4.4 ± 0.5 kJ/mol) shows that starting the simulations with a corresponding apo structure largely removes an offset which is otherwise present for the calculations initialized with the holo structures only. This indicates that substantial rearrangements occur in the studied proteins upon ligand binding that do not equilibrate at the nanosecond timescale covered in the simulations.

The largest effect from using an experimen-

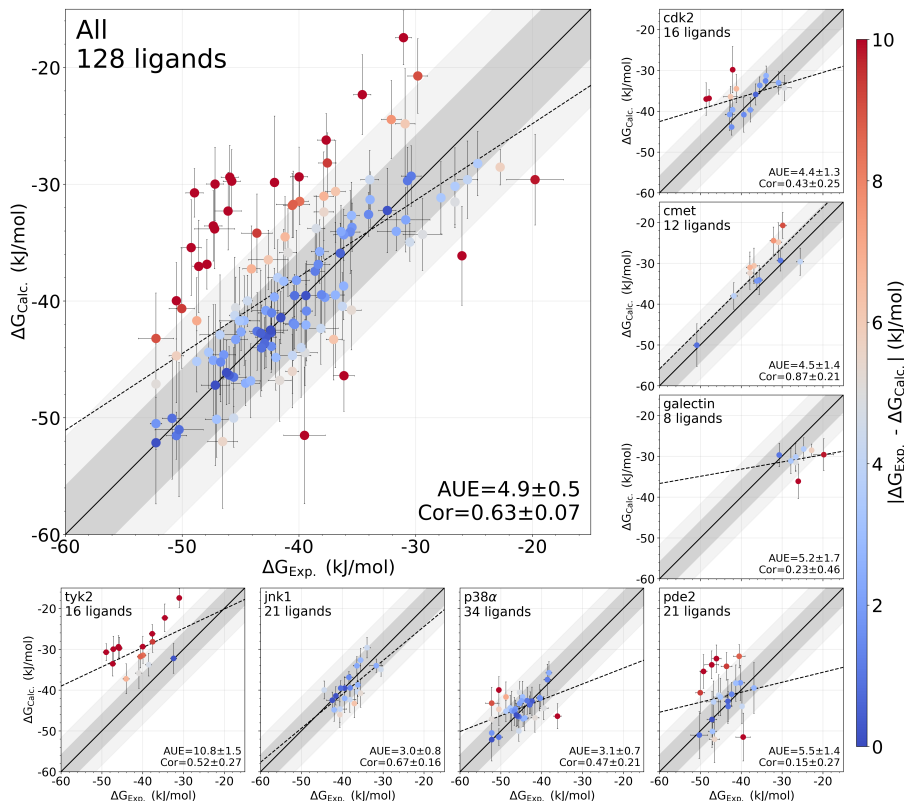


Figure 1: Overview of the accuracy of calculated absolute binding free energies $\Delta G_{\text{Calc.}}$. Error bars represent standard errors for free energies, absolute unsigned errors (AUE, units of kJ/mol), and the Pearson correlation coefficients (Cor). Apo states were initialized with X-ray crystal structures for all systems except pde2 and tyk2 where holo X-ray structures with the ligand removed were used. Dark and light shaded areas represent regions deviating from experiment by at most 1 and 2 kcal/mol.

tally resolved apo structure is observed in the galectin, p38 α and cdk2 protein-ligand complexes, while for the other cases the differences in accuracy are less affected. To understand what structural features are responsible for such pronounced effects, we have further explored the p38 α system for which ΔG had the largest difference among the systems depicted in Fig 2.

2.2 Large effect of a single rotamer

The p38 α protein-ligand complex shows a particularly strong dependence of the calculated ΔG on the starting structure. For this case, we were able to identify the particular structural details that are responsible for more than 9 kJ/mol offset in the calculated ΔG values




Figure 2: Experimental binding ΔG values plotted against the calculated estimates. In the panel on the left, the simulations in the apo state were started from the experimental holo structure after removing the ligand. In the panel on the right, the alchemical simulations of the protein in its apo state were initialized with the experimentally resolved apo structure. The starting structure has a marked effect on the calculated ΔG accuracy. The bottom panel shows a break up of the accuracies by protein system. Probability values measure the significance of the difference between apo and holo absolute unsigned errors via a Welch's t-test. Dark and light shaded areas represent regions deviating from experiment by at most 1 and 2 kcal/mol.

(Fig 3).

One of the main differences between the apo (pdb id 1wfc⁵³) and holo (3fly) structures occurring close to the binding site is a major loop motion: colored in orange and blue in Figure 3. However, it appears that even the short (10 ns) equilibrium simulations that we employed in the current protocol are sufficient to sample this loop transition (Fig S4). We have also explicitly probed whether this structural feature may modulate the accuracy of the calculated ΔG

values. We have filtered the starting structure ensemble for the ligand coupling transitions, retaining only those conformers with a loop position similar to the one from the crystallographic apo structure. This, however, had no effect on the calculated binding ΔG values (Fig S5).

While the large loop motion has no substantial effect on the ΔG accuracy, a single rotamer flip appears to be responsible for the larger than 9 kJ/mol shift in calculated ΔG . The crystallographic structures 3fly (holo) and 1wfc (apo)




Figure 3: A detailed investigation of the p38 α protein-ligand complexes. The apo (1wfc⁵³, orange) and holo (3fly, blue) structures have several structural differences close to the ligand binding site: a substantial loop motion and a different T106 rotamer state. In the simulations, the rotamer T106 retains its initial state: shown in lines, with a sphere marking threonine's oxygen. The calculated ΔG values depend strongly on the starting structure (holo or apo) that is used to initialize protein simulations in its apo state: scatterplots at the bottom. The green structure in the sub-panel and corresponding ΔG scatterplot depict a case, where apo simulations were initialized with a holo structure (ligand removed), but with the T106 rotamer set into its apo state. Dark and light shaded areas represent regions deviating from experiment by at most 1 and 2 kcal/mol.

have different threonine 106 (T106) rotameric states. Initializing apo simulations with either the experimentally resolved apo structure or a holo structure with the ligand removed yields

ensembles where the rotamer never crosses the barrier and remains in its starting state (Fig 3). The barrier crossing for the T106 sidechain rotamer appears to be too high to be sampled in the short (10 ns) equilibrium simulations used in the free energy calculation protocol.

To verify that T106 is truly the cause for this marked difference, we have initialized ligand coupling simulations from the holo structure (with the ligand removed), but setting the T106 rotamer into its apo state (green structures in Fig 3). This single change in the holo structure was sufficient to bring the calculated ΔG to the same accuracy as obtained from simulations started with the true apo structure.

It appears that initializing ligand coupling simulations from holo structures leaves the binding site - in particular the T106 rotamer - pre-arranged to accommodate the ligand. This, in turn, leads to an overly stabilized protein-ligand complex as quantified by the binding ΔG . The missing term in ΔG , in this case, is the free energy required to switch T106 rotamer from its apo to holo state. To demonstrate this, we also computed free energy surfaces for the residue's χ_1 dihedral with well-tempered metadynamics^{4,29} simulations biasing the potential of the dihedral. The free energy surfaces (Fig 4 A) reveal the average free energy difference between the minima of *gauche*- and *trans* conformations (present in the 1wfc and 3fly structures, respectively) of the apo state to be $\sim 8 \pm 2$ kJ/mol. This matches well the observed shift in the binding free energies calculated using 1wfc and 3fly starting structures for the apo state. Due to insufficient end-state sampling and high energy barriers, we do not observe a transition in this rotamer during short 10 ns equilibrium simulations, yet simulations started from the true apo state allowed taking the missing ΔG contribution into account.

2.3 Can longer simulations reveal true apo states?

Undersampling is a frequently encountered shortcoming of simulation-based phase space exploration, e.g. numerous examples are pro-

vided in¹⁴. Naturally, one of the underlying reasons for the inadequate representation of the apo state in the case of p38 α protein could be insufficient equilibration of the system. Therefore, we probed whether longer simulations would be able to cross the energy barrier and arrive in the true apo state when starting from a holo crystallographic structure with the ligand removed. To explore this, we have extended the p38 α apo state simulations started from 3fly by performing 5 independent runs of 1 μ s each.

The longer simulations indeed showed a transition of the T106 rotamer from its *trans* state (3fly holo conformer) to the *gauche*- state observed in the apo 1wfc structure (Fig 4 B). In all 5 independent replicas, the transition occurred within the first 200 ns. After this, no recrossings back to the *trans* rotameric state were observed, only short lived transitions from the *gauche*- to *gauche*+ state occurred.

Binding ΔG calculations where sampling of the decoupled ligand state is initialized with the final structures from 1 μ s simulations show this shift and have the same accuracy as those started with the crystallographic apo state 1wfc (Fig 4 C and D). This confirms our previous observation that the rotameric state of T106 plays a crucial role in the ligand binding to p38 α . All in all, the observations from the long simulations suggest that, at least in some cases, we can rely on longer (or enhanced) sampling to recover a protein's apo state for the subsequent ΔG calculations.

It is important to note, however, that the increased sampling does not automatically translate into a better agreement of the simulated trajectory with the experimentally measured observables. For example, longer simulations of the tyk2 kinase in its apo state (4gih³⁰ with the ligand removed; Fig S6) explore a broader range of conformations. However, as simulations progress, they deviate substantially from the starting crystallographic structure. The substantial drift of simulated trajectories, in turn, results in large uncertainties of the calculated binding affinities and deteriorates the ΔG prediction accuracy. This observation indicates that either the longer sampling reaching

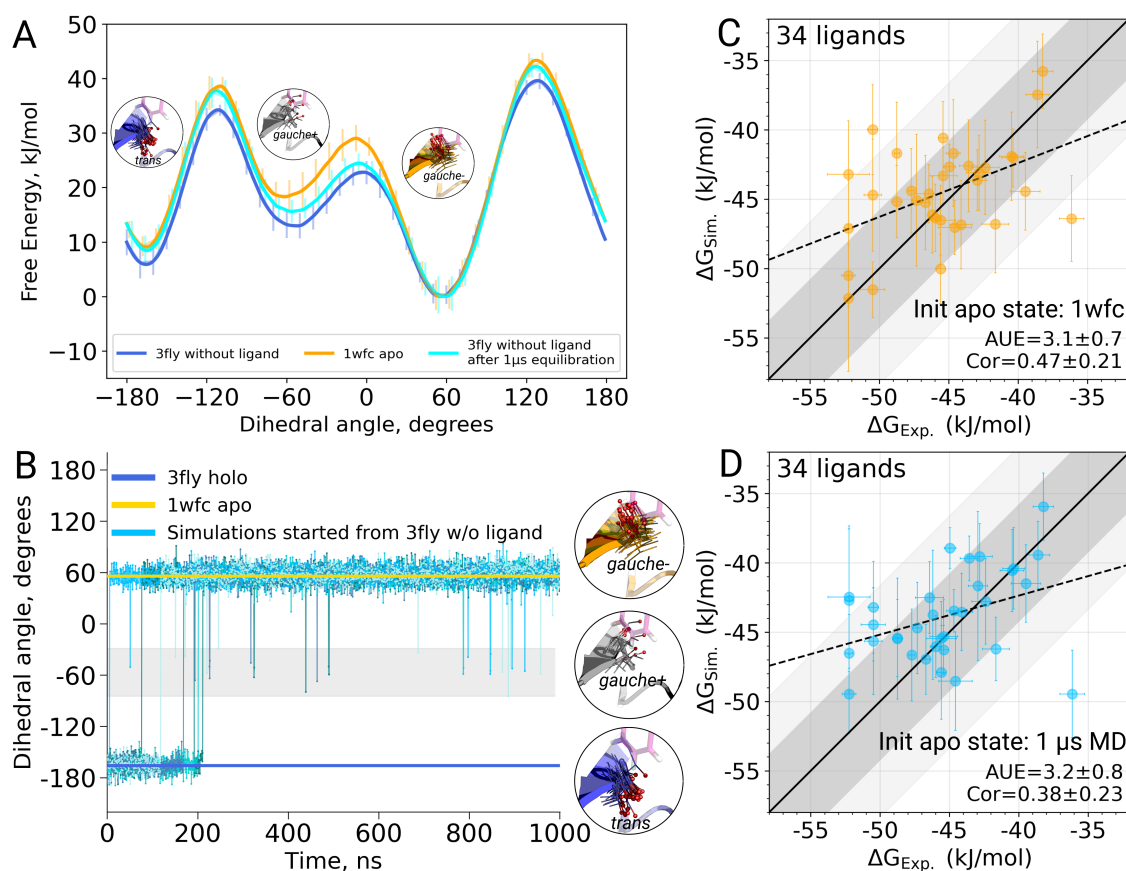


Figure 4: The effect of T106 rotamer states of the p38 α kinase on the calculated ligand binding ΔG . Free energy surfaces (A) of the χ_1 dihedral angle for the T106 residue obtained from well tempered metadynamics simulations starting from 3fly and 1wfc structures as well as from the output of 1 μ s equilibrations started from 3fly. χ_1 dihedral angle for the T106 residue in the crystallographic apo (1wfc) and holo (3fly) states, as well as in 5 independent simulations of 1 μ s each (B). Binding ΔG calculated by initializing apo state simulations with the 1wfc structure (C) and the end-states from 1 μ s simulations (D). Dark and light shaded areas represent regions deviating from experiment by at most 1 and 2 kcal/mol.

1 μ s for each of the 5 repeats is still not sufficient, or the new free energy minima identified by the force field are not representative of the true free energy landscape.

2.4 Using binding ΔG to identify apo states.

For the cases where multiple experimentally resolved structures are available, it may not be

evident which structure would be best suited for initializing simulations to obtain a representative apo state ensemble. It is, however, possible to exploit binding free energy calculations to identify the structure yielding the most probable conformational ensemble. This analysis does not require any knowledge of the actual (experimentally measured) set of binding affinities. It rather relies on multiple calculations of the binding affinities connecting one holo struc-

ture with multiple possible apo states (Fig 5).

We use phosphodiesterase 2 (pde2) complexed with 21 inhibitors³⁵ to illustrate this approach. Numerous experimentally resolved monomeric pde2 structures are available, where the protein is crystallized in its apo state (e.g. 4htz⁵⁵) or in a complex with a ligand (e.g. 6ezf³⁵, 4d08, 4d09⁸). Availability of these structures allows constructing a set of apo states by using either an actual apo conformer from the crystallographic structure or by removing a ligand from a holo structure. In principle, the most likely apo state is at its free energy minimum, i.e. of the multiple candidate conformers, the one with the lowest free energy would be the most populated in the ensemble. However, calculating free energy differences between the apo conformers directly is a computationally highly demanding challenge.

Instead, we can evaluate relative free energies of these conformers by connecting them via a common holo state. We calculate binding affinities for a set of 21 pde2 inhibitors using the structure 6ezf representing the protein-ligand complex and each of the 6ezf, 4htz, 4d08 and 4d09 structures independently representing the apo state. In this way we relate each apo state to one another via a common reference 6ezf holo state. Setting the free energy of the reference to 0 kJ/mol for convenience allows us to directly compare the apo states (Fig 5): the ΔG for an apo state is represented by averaged binding free energies calculated over the whole ligand set. The barrier (denoted with the dashed lines in Fig 5) is not attainable with this approach, as alchemical calculations do not explicitly probe the binding-unbinding pathway.

It is, however, important to understand the limitations of the ΔG values obtained this way. The calculated values should not be interpreted as reporting on the actual free energy differences between the apo conformers, but only on a component of ΔG corresponding to the change in the degrees of freedom relevant for ligand binding. It is likely that the binding site rearrangements are experienced by the ligands and have a strong effect on the ΔG calculated based on this approach. At the same time, substantial conformational rearrangements further

from the binding site may not have a contribution to ΔG if they do not affect the ligand binding affinity. Therefore, the conclusions about the most likely apo state identified with this approach should be limited to the interpretations of the binding affinities for a specific set of ligands.

In the current analysis, the 6ezf holo structure without the ligand was identified as the most likely representation of the apo structure for the set of 21 pde2 inhibitors. Interestingly, this structure is predicted to have a lower free energy than the crystallographically resolved apo state. One reason for that might be particular structural details that could have been resolved in a higher resolution structure 6ezf (1.5 Å)³⁵ as compared to 4htz (2.0 Å)⁵⁵, or larger conformational changes that may be more comparable with the full length apo protein³⁵. Comparison of the experimental binding affinities to the values calculated with the 6ezf structure as a template for the apo state provide further support for this methodology (Fig 5). The estimated ΔG values for this case have the best agreement to experiment (AUE of 5.5 ± 1.4 kJ/mol) in comparison to the calculations using the other structures. The similar correlations between experiment and calculation for all examples in Fig 5 again confirm the effect of the apo state to modify the offset of the calculated binding affinities.

3 Discussion

3.1 Relative Free Energies

It appears that the calculation of the overall offset is one of the major challenges in the absolute binding free energy estimation. Interestingly, given the equivalent simulation conditions for a set of ligands, even such large overall shifts in the calculated absolute ΔG values may have no effect on the relative free energies between the ligands (e.g. the case of tyk2 in Fig 1 and Fig S2). This suggests that the cause of the offset could be largely the same for all the considered ligands and cancels out in calculating the free energy differences. The protein-ligand complexes investigated in this study present a

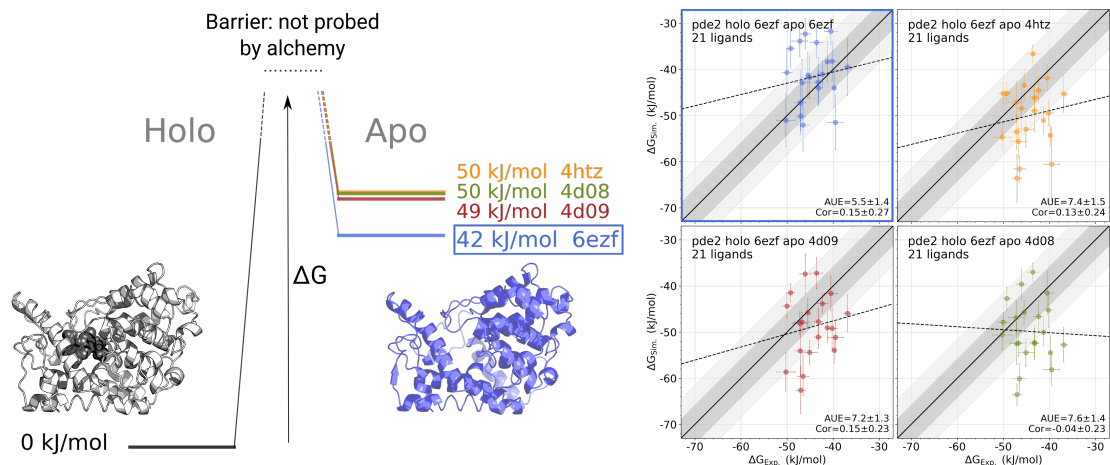


Figure 5: Identification of the most likely apo state for pde2 system based on binding affinities of 21 inhibitors. The absolute binding free energies were calculated using 6ezf holo state and 4 structures without the ligand to represent the apo state (4htz⁵⁵, 6ezf³⁵, 4d08 and 4d09⁸). The common holo state allows comparing the apo states one to another in terms of ΔG (the uncertainty of each estimate is below 1 kJ/mol). 6ezf structure is identified as the most likely apo structure based on the binding free energies for the considered ligand set. The panels on the right compare the experimental binding affinities to those calculated with each of the apo structures. Dark and light shaded areas represent regions deviating from experiment by at most 1 and 2 kcal/mol.

convenient set of systems for testing this hypothesis: the relative free energies for these systems have been previously calculated directly by alchemical transformations between ligand pairs with a non-equilibrium approach and the same force field¹⁶.

In Fig 6 we compare the relative binding free energies constructed from the absolute ΔG calculations to the values from Gapsys et al.¹⁶ obtained by an explicit relative $\Delta\Delta G$ calculation protocol. The absolute ΔG protocol indeed yields relative free energies comparable to those calculated via direct alchemical transformation of the ligands (AUE of 4.0 ± 0.4 kJ/mol) (Fig 6 left), indicating that the relative binding free energies are captured properly even when considering additional challenges of the absolute ΔG estimation. Furthermore, the accuracy of relative free energies obtained from the absolute ΔG protocol is also in good agreement with experimental results (Fig 6 middle)

yielding an AUE of 4.5 ± 0.4 kJ/mol (correlation of 0.54 ± 0.08) in comparison to AUE of 3.6 ± 0.3 kJ/mol (correlation of 0.65 ± 0.05) for explicit relative calculations (Fig 6 right) for the same systems.

This observation is encouraging for the prospective drug design studies. Absolute ΔG calculations can be reliably used for the cases where the main assumptions for estimating relative free energy differences do not hold, e.g. where binding pose changes occur or investigated ligand structures differ substantially. It is, however, important to take into account the computational time required by these methods: absolute ΔG estimates in this work required 10 times longer sampling in comparison to the $\Delta\Delta G$ calculations in¹⁶. The difference in computational cost between these approaches suggests a natural delineation in their application. When exploring large chemical libraries by means of free energy calculations, it




Figure 6: Accuracy of the relative $\Delta\Delta G$ values. Comparison of the relative binding free energies calculated from absolute ΔG values to the $\Delta\Delta G$ values calculated by explicit alchemical ligand modifications¹⁶ (left). $\Delta\Delta G$ values from absolute ΔG compared to experiment (middle). $\Delta\Delta G$ values from explicit alchemical ligand modifications compared to experiment (right). Dark and light shaded areas represent regions deviating from experiment by at most 1 and 2 kcal/mol.

would be most efficient to evaluate structurally similar compounds by computing $\Delta\Delta G$ values, while absolute ΔG calculations could be performed less frequently for the cases that are not tractable by the relative free energy estimation.

It is important to mention that there also exist specialized approaches based on the relative $\Delta\Delta G$ calculations to evaluate free energy differences for structurally highly distinct ligands and different binding poses, e.g. separated topology method³⁹. Using the relative free energy calculations has the advantage of avoiding the requirement to properly represent protein’s apo state, as this state is not explicitly considered. Yet, absolute binding free energy protocol offers a number of additional possibilities. For example, estimation of the absolute ΔG makes it possible to evaluate ligand selectivity against different protein targets, evaluate affinity for various protein conformers, calculate binding affinities for individual molecules without the need to consider them in a relation to other ligands.

3.2 Sources of statistical uncertainty

Calculations of the absolute binding free energies show larger statistical uncertainties when compared to the relative free energy calculations (Fig S10). The increase in statisti-

cal errors arises due to larger perturbations to the system required by an alchemical absolute ΔG calculation. Coupling/decoupling of the whole ligand involves introducing/removing more interactions in comparison to the alchemical transformations of a small number of atoms when morphing ligands to one another for relative free energy estimations. Convergence of the absolute ΔG estimates in pharmaceutically relevant systems can be achieved, yet it requires extending the alchemical transitions to nanoseconds¹⁷. Such slower transitions retain the system closer to equilibrium, dissipating less work along the alchemical path, thus facilitating convergence.

Although lower uncertainties of the estimated ΔG are desired, the long alchemical transition times quickly become intractable for large scale ligand binding affinity scans. Therefore, it is necessary to balance the trade-off between the available simulation time and the attainable precision. This, naturally, requires a robust uncertainty estimation for the ΔG estimates. It has been observed that relying on the statistical uncertainties from the ΔG estimators, either analytical expressions, or bootstrapped values, may not be reliable.^{5,38} Therefore, in this work we rely on independent repeats of the whole free energy calculation procedure to gain access to the variation of the ΔG estimates.^{14,46} Subsequently, we incorpo-

rate both, uncertainties from the independent replicas and statistical uncertainty from the estimator by means of bootstrap into a single uncertainty value.¹⁶

3.3 Apo protein state in absolute ΔG calculations

The major conceptual difference between the absolute and relative binding free energy calculations stems from the need to explicitly consider the apo protein state when computing absolute ΔG . This poses a challenge for a theoretically rigorous treatment of the decoupled ligand that subsequently needs to be coupled to the system in a well-defined binding site of the protein. In the current work we present a novel approach for the construction of the decoupled ligand state ensembles (see Methods) which, in combination with the ligand restraining protocol⁷, provides an efficient solution to the problem. In brief, our method positions and restrains the decoupled ligand in the binding pocket of the apo protein creating a decoupled state ensemble without the need to explicitly simulate it.

Furthermore, explicit consideration of the protein’s apo state also requires accurate quantification of a transition between the protein’s conformational states sampled upon ligand binding. The non-equilibrium free energy calculation approach presents a convenient setting, where the simulations for holo and apo states can be initialized with different starting structures¹⁷. In such a way, the apo and holo state ensembles can be generated by simulations started with the corresponding experimentally resolved structures whenever they are available. The initialization of the simulations with a proper starting structure has a profound effect on the accuracy of estimated ΔG (Fig 2).

This observation, however, could be interpreted merely as a sampling issue: routine free energy calculation protocols use short (5-20 ns) equilibrium simulations^{16,43,52} that may not be sufficient for generating a representative apo state ensemble. Inaccuracies in the estimated free energies due to undersampling have been previously reported for both rel-

ative³¹ and absolute⁵ protein-ligand binding free energy calculations. The issue can be alleviated with longer simulations or enhanced sampling. This appears to be feasible in the case of p38 α kinase, where longer simulation of the protein’s apo state was able to recover the experimentally resolved rotamer T106 which proved essential for accurate ΔG calculations (Fig 4). Yet, the case of tyk2 kinase, for which long (1 μ s) simulations were used for the apo state, demonstrates that the extended sampling does not necessarily lead to higher accuracy in ΔG estimation (Fig S6). This is in line with several previous observations where enhanced sampling showed no improvement in the accuracy of the free energy estimates.^{27,48} In fact, a deterioration in prediction accuracy can be observed in longer or enhanced-sampling simulations, when the ligand explores poses that are less relevant for binding.⁴⁸ In turn, this manifests in an underestimation of the relative binding free energy differences⁴⁸, which we have also observed in our study (Fig 6).

Another approach that we introduced in this study allows to circumvent the need of an exhaustive apo state sampling by probing multiple initial apo states (when they are available) with the absolute ΔG calculation protocol (Fig 5). This method does not require any prior knowledge of the experimentally measured binding affinities and it allows estimating relative free energies for the apo states by relating them one to another via a common holo reference state. The ΔG value for apo state structures calculated this way represents only one component of the overall free energy of the conformers, as only a contribution that is experienced by the ligand binding is considered. Nevertheless, this method allows identifying the most likely apo state for the use in the absolute binding free energy calculations.

In this study we used datasets that have previously been used for relative binding free energy calculations. We observed how the absolute calculations could return good correlations, with the apo state affecting the offset seen in terms of the larger AUE. In other words, the difference in apo state conformation and energy had a consistent effect for the binding free en-

ergy of the ligands for each target. It remains to be seen if that will hold true as the diversity of the ligands increases, even if they are binding in the same site. We anticipate that future studies on an even larger scale will be required to examine these effects.

While in this work we have highlighted the importance of the proper apo state ensemble for the accurate absolute binding free energy predictions, it is essential to reliably represent the holo state as well. Here, we relied on the crystallographic protein holo states and carefully modeled ligand binding poses from previous investigation¹⁶. Naturally, the ligand modelling step introduces additional uncertainty in defining the starting structure for initializing the simulations. Accurate binding ΔG estimates suggest that the holo state representation was proper for most of the investigated cases. The tyk2 kinase, however, is an exception, as the calculated ΔG values significantly underestimate the experimentally measured binding affinities (Fig 1). The apo state representation is unlikely to be solely responsible, as identification of any deeper free energy minima for the apo state would only impose an additional penalty on the ΔG of binding, thus reducing the predicted affinity even further, as illustrated in Fig S6. A deeper free energy minimum for the holo state can lead to the prediction of a lower binding ΔG . This prompts us to assume that the holo state representation for the tyk2 kinase could be improved by exploring additional ligand poses, protein conformations, internal water placement or a combination of these components. This way, tyk2 could serve as an interesting candidate for future investigations possibly presenting a challenging case for the holo state description.

4 Conclusions

In this work we propose methodological advances that enable efficient absolute binding free energy calculations with an accuracy on par with relative free energy calculations. We demonstrate the generality of the protocol across multiple pharmaceutically relevant tar-

gets in a large scale study. Our approach enables the incorporation of both holo and apo structural information for reliable affinity predictions. The key structural determinants of binding can be as small as a single rotamer change between the apo and holo states and appropriate sampling of such determinants can be computationally demanding. When multiple alternative apo structures exist, absolute binding free energy calculations can be used to identify the most likely candidate for a prospective study.

5 Methods

The process of a ligand binding to a protein requires considering two end-states: solvated ligand and ligand bound to the protein. Computationally, these two states can be connected via alchemical paths arranged in a thermodynamic cycle depicted in Fig 7^{2,18}.

Following this thermodynamic cycle, firstly, the ligand located in solvent (state A') is decoupled from its environment (state B'). The decoupled ligand (state B') is allowed to freely sample the whole simulation box. To be able to proceed with the second leg of the cycle, i.e. coupling the ligand to the system in the protein's binding site, the ligand needs to be restrained to the protein (state B). The contribution of the added restraints is taken into account analytically⁷. Finally, the ligand in the protein's binding site is coupled to the system and the restraints are removed (state A). The free energies for the two legs of the thermodynamic cycle are obtained separately by performing multiple non-equilibrium transitions in the ligand coupling and decoupling directions, recording the work distributions and using the Crook's Fluctuation Theorem¹¹ to evaluate the free energy differences.

The absolute ΔG calculations are particularly sensitive to the decoupled ligand restraining method. In our setup we employ a rigorous restraining approach⁷ acting between the protein and the decoupled ligand to anchor it in a narrow range of orientations within the binding pocket (SI Section 1). This restraint scheme




Figure 7: Diagram of the thermodynamic cycle for absolute binding calculations. As inserting a whole protein around a ligand would involve a massive disturbance in the system, the binding free energy ΔG_{bind} is calculated by traveling the long way along the thermodynamic cycle: first decoupling the ligand from the surrounding solvent, applying the analytical correction for the effect of protein-ligand restraints⁷, and then coupling the ligand back in the protein active site. The equilibrium structures for the decoupled ligand in the active site (state B) can be generated by aligning its structures in solvent (state B') into equilibrium frames of the apo protein.

uses six orthogonal relative restraints with harmonic potentials (a distance, two angles, and three dihedrals) acting on three anchor atoms in the ligand and three in the protein. The orthogonality of the potentials restraining the decoupled ligand allows for an analytical expression of the free energy contribution ΔG_{restr} .

5.1 Novel approach for treating the ligand's decoupled state

The alchemical approaches for absolute ligand binding free energies require explicit sampling

of the ligand-protein complex with the ligand in its decoupled state (state B in Fig 7). For that, a definition of restraints prior to starting the simulations is needed. The partition function of the decoupled state, however, can be separated into the independent contributions from the apo protein, the restraints, and the internal degrees of freedom of the decoupled ligand⁷. The simulation trajectories of the decoupled ligand (state B') are readily generated for every considered ligand in the ligand-solvation leg of the thermodynamic cycle. The simulation of an apo protein does not contain the ligand, thus a single trajectory of such a protein can be generated and later used in combination with any ligand of interest.

In the novel proposed approach we suggest generating an equilibrium ensemble of the decoupled ligand in the protein's binding site without the explicit simulation of this state. For that purpose we use the readily available trajectories of the decoupled ligand in water and protein in its apo state. Firstly, each frame of the protein-ligand trajectory (state A in Fig 7) is superimposed onto the corresponding frame of the apo protein trajectory (state B in Fig 7) by aligning their α -carbons. The corresponding frame of the decoupled ligand in solvent (state B' in Fig 7) is then aligned onto the new coordinates of ligand from the protein-ligand complex using all heavy atoms as reference. The now appropriately positioned decoupled ligand atoms are added to the apo protein trajectory. Finally, the six orthogonal restraints are constructed to match the potentials that would have generated equivalent distributions for each restrained degree of freedom in an explicit simulation (SI Section 1).

An ensemble created this way, however, may contain correlations between the restrained degrees of freedom (Fig S11, S12), therefore, to acquire a well defined ensemble of ligand poses we implemented several correction schemes. For the *a priori* correction, we sample from the harmonic protein-ligand restraint potentials at the simulation temperature, thus creating a proper ensemble of ligand orientations. This ensemble can be used for calculating the ligand-protein coupling free energy. The *post*

hoc correction allows performing the calculations starting directly with the superpositioned ensemble which still contains the correlations between the restrained degrees of freedom. In this case, the work values obtained from the ligand coupling simulations are adjusted with the contribution from the correlations as illustrated in the scheme in Fig S13.

We verified the validity of the superpositioning approach and the performance of the proposed decorrelation techniques, by comparing their predictions for a subset of the studied proteins (tyk2, jnk1, and p38 α) to those of a standard protocol where the restrained state was simulated explicitly (Fig S14, S15). For the main results reported in this work we used the *post hoc* decorrelation method.

5.2 Simulation details

All simulations were carried out with the GROMACS 2019.4 molecular dynamics engine¹ modified to correctly handle pair interactions within a decoupled molecule larger than the electrostatic cutoff (bug 3403). The Amber99sb*ILDN^{6,22,32} force field was used for the proteins throughout this work together with the TIP3P²⁵ water model. Ligand parameters were taken from the previous relative free energy study¹⁶ parametrized with the General Amber Force Field (GAFF v2.1)⁴⁹ using AM1-BCC charges²³ assigned with ACPYPE⁴⁵ and AnteChamber⁵⁰. Initial ligand binding poses were reused from a previous relative binding free energy study¹⁶, where the ligands were modeled based on similar compounds in existing holo crystal structures (Table S2). Initial protein structures were stripped of surrounding water (if present) and resolvated in dodecahedral simulation boxes with 1.5 nm of padding between solute and box edges. For apo simulations started from holo structures the ligands were removed before adding water, so that the binding cavities could also be filled with water. Effects of retaining crystallographic water before filling the remaining cavities were also examined (Fig S7). Ions were added to neutralize the system and reach a salt concentration of 150 mM.

Van der Waals interactions were calculated with a 1.1 nm cutoff and a switching function starting at 1.0 nm. Coulomb interactions were computed with Smooth Particle Mesh Ewald^{12,13} and a real space cutoff of 1.1 nm. For temperature regulation a system-wide stochastic velocity rescale thermostat⁹ was used with a time constant of 0.1 ps and a target temperature of 298 K. The pressure was kept at 1 bar with the aid of the isotropic Parrinello-Rahman barostat³⁴ with a time constant of 5 ps and a compressibility of $4.6 \cdot 10^{-5} \text{ bar}^{-1}$. Throughout this work a 2 fs time step was used with all bond lengths constrained via the LINCS²¹ algorithm.

Initial holo structures for ligands coupled to proteins as well as the solo ligands (used for the ligand in water leg of the thermodynamic cycle) were reused from the relative binding free energy study¹⁶. Initial structures for protein apo simulations were constructed from apo crystal structures (PDB IDs 3o17 (jnk1), 1wfc⁵³ (p38 α), 4htz⁵⁵ (pde2), 1h27³³ (cdk2), 1r1w⁴² (cmet), 3zsl⁴¹ (galectin)), where available. For these structures missing residues were modeled in and the amino acid protonation states were adjusted to match those of the holo structures. Energy minimization and equilibration with NVT and NPT simulations in the presence of solvent and ions were also carried out (in the same manner as in the core part of the protocol below) to relax the reconstructed residues. Finally, the apo protein structures were extracted from the last frame of the NPT simulations (or, in cases where no residue reconstruction was necessary, from the protonation adjustment stage) and were used to initialize the protein-only systems. As no apo crystal structure was available for tyk2, a holo structure with the ligand removed was used instead.

To obtain equilibrium distributions of the coupled protein-ligand and protein-only systems at 298 K, harmonic position restraints with a force constant of 9000 kJ/mol/nm² were applied to all protein and ligand atoms of the above initial structures and the energy was minimized followed by a 300 ps simulation in the NVT ensemble, where the first 5 ps were used to bring the temperature of the system from 0

to 298 K with simulated annealing. Position restraints were then relaxed to 500 kJ/mol/nm² for a 50 ps NVT simulation followed by a 10 ns production NPT simulation without any position restraints. For the leg of the thermodynamic cycle of ligands in water, the simulations used no position restraints. Firstly, energy minimization was performed followed by the 10 ps NVT and 10 ns production NPT simulations.

To initialize the non-equilibrium alchemical transitions, the first 2.256 ns of all production simulations were discarded and the equilibrium conformations were sampled every 67 ps yielding 165 conformations for each system. The extracted conformers were used to construct an equilibrium ensemble of the decoupled ligand in the protein’s binding site and generate protein-ligand restraints as described in the Section 1.

The non-equilibrium simulations, each 500 ps long, were run from each conformation to the opposite coupling state of the ligand by linearly interpolating the Hamiltonian between the two end-states. The gradients $\partial H(\lambda, x)/\partial \lambda$ were integrated over the course of each non-equilibrium simulation to obtain the amount of work performed. Free energies were computed from the work distributions in both directions using a maximum likelihood estimator⁴⁴ based on the Crooks Fluctuation Theorem¹¹ by means of pmx¹⁵. Finally, free energy estimates from different legs of the thermodynamic cycle were combined and the contribution of restraining the decoupled ligand to the protein was added⁷ incorporating the correction for the correlations in the restrained degrees of freedom (Section 2).

Well-tempered metadynamics^{4,29} calculations were carried out with GROMACS 2016.3¹ in combination with plumed 2.3.1.⁴⁷ A bias factor of $(T + \Delta T)/T = 11$ with a time constant of $\tau = 10$ ps, a time step of 0.002 ps, and a temperature $T = 298$ K were used for 100 ns simulations started from apo structures previously equilibrated as described above. Every $\tau_{\text{pace}} = 2$ ps Gaussian biases of 5 degree width and an initial height of $k_b \Delta T \tau_{\text{pace}}/\tau = 2k_b T \approx 4.955$ kJ/mol, where k_b is the Boltzmann constant, were deposited onto a periodic grid consisting of 360 points equally distributed along

the χ_1 dihedral. Resulting free energy surfaces and uncertainties are reported as averages and standard deviations across 5 repeats.

Throughout this work uncertainties were computed via bootstrap, unless explicitly specified otherwise, and represent standard deviations when taking into account all available calculations. Bootstrapping was performed for the individual repeats of free energy predictions for each ligand based on the work values, the final free energy prediction for each ligand across multiple repeats, as well as for AUE and Pearson correlation functions across multiple ligands. Actual values around which these uncertainties are reported are still the means (or results of the AUE and correlation functions) of the underlying data, not the means of the bootstrapped distributions.

Author contributions

All authors contributed to the manuscript and gave approval to its final version.

Acknowledgements

YK was supported by the Vlaams Agentschap Innoveren & Ondernemen (VLAIO) project number HBC.2018.2295, "Dynamics for Molecular Design (DynaMoDe)". VG was supported by the BioExcel CoE (<http://www.bioexcel.eu>), a project funded by the European Union (Contract H2020-INFRAEDI-02-2018-823830).

References

- [1] Mark James Abraham, Teemu Murtola, Roland Schulz, Szilárd Páll, Jeremy C. Smith, Berk Hess, and Erik Lindahl. GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX*, 1-2:19–25, September 2015. ISSN 2352-7110. doi: 10.1016/j.softx.2015.06.001.
- [2] Matteo Aldeghi, Alexander Heifetz, Michael J. Bodkin, Stefan Knapp, and Philip C. Biggin. Accurate calculation of the absolute free energy of binding

- p>for drug molecules.
- Chem. Sci.*
- , 7(1): 207–218, 2016. ISSN 2041-6539. doi: 10.1039/C5SC02678D.
- [3] Matteo Aldeghi, Alexander Heifetz, Michael J. Bodkin, Stefan Knapp, and Philip C. Biggin. Predictions of Ligand Selectivity from Absolute Binding Free Energy Calculations. *J. Am. Chem. Soc.*, 139(2):946–957, January 2017. ISSN 0002-7863. doi: 10.1021/jacs.6b11467.
 - [4] Alessandro Barducci, Giovanni Bussi, and Michele Parrinello. Well-Tempered Metadynamics: A Smoothly Converging and Tunable Free-Energy Method. *Phys. Rev. Lett.*, 100(2):020603, January 2008. doi: 10.1103/PhysRevLett.100.020603.
 - [5] Hannah M. Baumann, Vytautas Gapsys, Bert L. de Groot, and David L. Mobley. Challenges Encountered Applying Equilibrium and Nonequilibrium Binding Free Energy Calculations. *J. Phys. Chem. B*, 125(17):4241–4261, May 2021. ISSN 1520-6106. doi: 10.1021/acs.jpcc.0c10263.
 - [6] Robert B. Best and Gerhard Hummer. Optimized Molecular Dynamics Force Fields Applied to the Helix-Coil Transition of Polypeptides. *J. Phys. Chem. B*, 113(26): 9004–9015, July 2009. ISSN 1520-6106. doi: 10.1021/jp901540t.
 - [7] Stefan Boresch, Franz Tettinger, Martin Leitgeb, and Martin Karplus. Absolute Binding Free Energies: A Quantitative Approach for Their Calculation. *J. Phys. Chem. B*, 107(35):9535–9551, September 2003. ISSN 1520-6106. doi: 10.1021/jp0217839.
 - [8] Peter Buijnsters, Meri De Angelis, Xavier Langlois, Frederik J. R. Rombouts, Wendy Sanderson, Gary Tresadern, Alison Ritchie, Andrés A. Trabanco, Greet VanHoof, Yves Van Roosbroeck, and José-Ignacio Andrés. Structure-Based Design of a Potent, Selective, and Brain Penetrating PDE2 Inhibitor with Demonstrated Target Engagement. *ACS Med. Chem. Lett.*, 5(9): 1049–1053, September 2014. ISSN 1948-5875, 1948-5875. doi: 10.1021/ml500262u.
 - [9] Giovanni Bussi, Davide Donadio, and Michele Parrinello. Canonical sampling through velocity rescaling. *The Journal of Chemical Physics*, 126(1):014101, January 2007. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.2408420.
 - [10] Myriam Ciordia, Laura Pérez-Benito, Francisca Delgado, Andrés A. Trabanco, and Gary Tresadern. Application of Free Energy Perturbation for the Design of BACE1 Inhibitors. *J. Chem. Inf. Model.*, 56(9):1856–1871, September 2016. ISSN 1549-9596, 1549-960X. doi: 10.1021/acs.jcim.6b00220.
 - [11] Gavin E. Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Phys. Rev. E*, 60(3):2721–2726, September 1999. ISSN 1063-651X, 1095-3787. doi: 10.1103/PhysRevE.60.2721.
 - [12] Tom Darden, Darrin York, and Lee Pedersen. Particle mesh Ewald: An $N \cdot \log(N)$ method for Ewald sums in large systems. *The Journal of Chemical Physics*, 98(12): 10089–10092, June 1993. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.464397.
 - [13] Ulrich Essmann, Lalith Perera, Max L. Berkowitz, Tom Darden, Hsing Lee, and Lee G. Pedersen. A smooth particle mesh Ewald method. *The Journal of Chemical Physics*, 103(19):8577–8593, November 1995. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.470117.
 - [14] Vytautas Gapsys and Bert L de Groot. On the importance of statistics in molecular simulations for thermodynamics, kinetics and simulation box size. *eLife*, 9:e57589, August 2020. ISSN 2050-084X. doi: 10.7554/eLife.57589.
 - [15] Vytautas Gapsys, Servaas Michielssens, Daniel Seeliger, and Bert L. de Groot. Pmx: Automated protein structure and

- topology generation for alchemical perturbations. *J. Comput. Chem.*, 36(5):348–354, 2015. ISSN 1096-987X. doi: 10.1002/jcc.23804.
- [16] Vytautas Gapsys, Laura Pérez-Benito, Matteo Aldeghi, Daniel Seeliger, Herman van Vlijmen, Gary Tresadern, and Bert L. de Groot. Large scale relative protein ligand binding affinities using non-equilibrium alchemy. *Chem. Sci.*, 11(4): 1140–1152, January 2020. ISSN 2041-6539. doi: 10.1039/C9SC03754C.
- [17] Vytautas Gapsys, Ahmet Yildirim, Matteo Aldeghi, Yuriy Khalak, David van der Spoel, and Bert L. de Groot. Accurate absolute free energies for ligand–protein binding based on non-equilibrium approaches. *Commun Chem*, 4(1):1–13, May 2021. ISSN 2399-3669. doi: 10.1038/s42004-021-00498-y.
- [18] M.K. Gilson, J.A. Given, B.L. Bush, and J.A. McCammon. The statistical-thermodynamic basis for computation of binding affinities: A critical review. *Biophysical Journal*, 72(3):1047–1069, March 1997. ISSN 00063495. doi: 10.1016/S0006-3495(97)78756-3.
- [19] David F. Hahn, Gerhard König, and Philippe H. Hünenberger. Overcoming Orthogonal Barriers in Alchemical Free Energy Calculations: On the Relative Merits of λ -Variations, λ -Extrapolations, and Biasing. *J. Chem. Theory Comput.*, 16(3): 1630–1645, March 2020. ISSN 1549-9618. doi: 10.1021/acs.jctc.9b00853.
- [20] J. Hermans and S. Shankar. The Free Energy of Xenon Binding to Myoglobin from Molecular Dynamics Simulation. *Isr. J. Chem.*, 27(2):225–227, 1986. ISSN 00212148. doi: 10.1002/ijch.198600032.
- [21] Berk Hess. P-LINCS: A Parallel Linear Constraint Solver for Molecular Simulation. *J. Chem. Theory Comput.*, 4(1):116–122, January 2008. ISSN 1549-9618. doi: 10.1021/ct700200b.
- [22] Viktor Hornak, Robert Abel, Asim Okur, Bentley Strockbine, Adrian Roitberg, and Carlos Simmerling. Comparison of multiple Amber force fields and development of improved protein backbone parameters. *Proteins*, 65(3):712–725, November 2006. ISSN 08873585, 10970134. doi: 10.1002/prot.21123.
- [23] Araz Jakalian, Bruce L. Bush, David B. Jack, and Christopher I. Bayly. Fast, efficient generation of high-quality atomic charges. AM1-BCC model: I. Method. *J. Comput. Chem.*, 21(2): 132–146, 2000. ISSN 1096-987X. doi: 10.1002/(SICI)1096-987X(20000130)21:2<132::AID-JCC5>3.0.CO;2-P.
- [24] Wei Jiang and Benoît Roux. Free Energy Perturbation Hamiltonian Replica-Exchange Molecular Dynamics (FEP/H-REMD) for Absolute Ligand Binding Free Energy Calculations. *J. Chem. Theory Comput.*, 6(9):2559–2565, September 2010. ISSN 1549-9618. doi: 10.1021/ct1001768.
- [25] William L. Jorgensen, Jayaraman Chandrasekhar, Jeffry D. Madura, Roger W. Impey, and Michael L. Klein. Comparison of simple potential functions for simulating liquid water. *J. Chem. Phys.*, 79(2): 926–935, July 1983. ISSN 0021-9606. doi: 10.1063/1.445869.
- [26] Henrik Keränen, Laura Pérez-Benito, Myriam Ciordia, Francisca Delgado, Thomas B. Steinbrecher, Daniel Oehlrich, Herman W. T. van Vlijmen, Andrés A. Trabanco, and Gary Tresadern. Acylguanidine Beta Secretase 1 Inhibitors: A Combined Experimental and Free Energy Perturbation Study. *J. Chem. Theory Comput.*, 13(3):1439–1453, March 2017. ISSN 1549-9618, 1549-9626. doi: 10.1021/acs.jctc.6b01141.
- [27] Yuriy Khalak, Gary Tresadern, Bert L. de Groot, and Vytautas Gapsys. Non-equilibrium approach for binding free energies in cyclodextrins in SAMPL7: Force

- p>fields and software.
- J Comput Aided Mol Des*
- , November 2020. ISSN 1573-4951. doi: 10.1007/s10822-020-00359-1.
- [28] Maximilian Kuhn, Stuart Firth-Clark, Paolo Tosco, Antonia S. J. S. Mey, Mark Mackey, and Julien Michel. Assessment of Binding Affinity via Alchemical Free-Energy Calculations. *J. Chem. Inf. Model.*, 60(6):3120–3130, June 2020. ISSN 1549-9596, 1549-960X. doi: 10.1021/acs.jcim.0c00165.
 - [29] Alessandro Laio and Michele Parrinello. Escaping free-energy minima. *PNAS*, 99(20):12562–12566, October 2002. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.202427399.
 - [30] Jun Liang, Anne van Abbema, Mercedesz Balazs, Kathy Barrett, Leo Berezhkovsky, Wade Blair, Christine Chang, Donnie Delarosa, Jason DeVoss, Jim Driscoll, Charles Eigenbrot, Nico Ghilardi, Paul Gibbons, Jason Halladay, Adam Johnson, Pawan Bir Kohli, Yingjie Lai, Yanzhou Liu, Joseph Lyssikatos, Priscilla Mantik, Kapil Menghrajani, Jeremy Murray, Ivan Peng, Amy Sambrone, Steven Shia, Young Shin, Jan Smith, Sue Sohn, Vickie Tsui, Mark Ultsch, Lawren C. Wu, Yisong Xiao, Wenqian Yang, Judy Young, Birong Zhang, Bing-yan Zhu, and Steven Magnuson. Lead Optimization of a 4-Aminopyridine Benzamide Scaffold To Identify Potent, Selective, and Orally Bioavailable TYK2 Inhibitors. *J. Med. Chem.*, 56(11):4521–4536, June 2013. ISSN 0022-2623, 1520-4804. doi: 10.1021/jm400266t.
 - [31] Nathan M. Lim, Lingle Wang, Robert Abel, and David L. Mobley. Sensitivity in Binding Free Energies Due to Protein Reorganization. *J. Chem. Theory Comput.*, 12(9):4620–4631, September 2016. ISSN 1549-9618, 1549-9626. doi: 10.1021/acs.jctc.6b00532.
 - [32] Kresten Lindorff-Larsen, Stefano Piana, Kim Palmo, Paul Maragakis, John L Klepeis, Ron O Dror, and David E Shaw. Improved side-chain torsion potentials for the Amber ff99SB protein force field. *Proteins*, 78(8):1950–1958, June 2010. ISSN 0887-3585. doi: 10.1002/prot.22711.
 - [33] Edward D. Lowe, Ivo Tews, Kin Yip Cheng, Nick R. Brown, Sheraz Gul, Martin E. M. Noble, Steven J. Gamblin, and Louise N. Johnson. Specificity Determinants of Recruitment Peptides Bound to Phospho-CDK2/Cyclin A^{†,‡}. *Biochemistry*, 41(52):15625–15634, December 2002. ISSN 0006-2960, 1520-4995. doi: 10.1021/bi0268910.
 - [34] M. Parrinello and A. Rahman. Polymorphic transitions in single crystals: A new molecular dynamics method. *Journal of Applied Physics*, 52(12):7182–7190, December 1981. ISSN 0021-8979. doi: 10.1063/1.328693.
 - [35] Laura Pérez-Benito, Henrik Keränen, Herman van Vlijmen, and Gary Tresadern. Predicting Binding Free Energies of PDE2 Inhibitors. The Difficulties of Protein Conformation. *Sci. Rep.*, 8(1):4883, March 2018. ISSN 2045-2322. doi: 10.1038/s41598-018-23039-5.
 - [36] Piero Procacci. Methodological uncertainties in drug-receptor binding free energy predictions based on classical molecular dynamics. *Current Opinion in Structural Biology*, 67:127–134, April 2021. ISSN 0959-440X. doi: 10.1016/j.sbi.2020.08.001.
 - [37] Piero Procacci and Guido Guarnieri. SAMPL7 blind predictions using nonequilibrium alchemical approaches. *J Comput Aided Mol Des*, January 2021. ISSN 0920-654X, 1573-4951. doi: 10.1007/s10822-020-00365-3.
 - [38] Andrea Rizzi, Travis Jensen, David R. Slocower, Matteo Aldeghi, Vytautas Gapys, Dimitris Ntekoumes, Stefano Bosio, Michail Papadourakis, Niel M. Henriksen, Bert L. de Groot, Zoe Cournia, Alex Dickson, Julien Michel, Michael K. Gilson,

- Michael R. Shirts, David L. Mobley, and John D. Chodera. The SAMPL6 SAMPLing challenge: Assessing the reliability and efficiency of binding free energy calculations. *J Comput Aided Mol Des*, 34(5): 601–633, May 2020. ISSN 1573-4951. doi: 10.1007/s10822-020-00290-5.
- [39] Gabriel J. Rocklin, David L. Mobley, and Ken A. Dill. Separated topologies—A method for relative binding free energy calculations using orientational restraints. *J. Chem. Phys.*, 138(8):085104, February 2013. ISSN 0021-9606. doi: 10.1063/1.4792251.
- [40] Frederik J. R. Rombouts, Gary Treadern, Peter Buijnsters, Xavier Langlois, Fulgencio Tovar, Thomas B. Steinbrecher, Greet Vanhoof, Marijke Somers, José-Ignacio Andrés, and Andrés A. Trabanco. Pyrido[4,3-*e*] [1,2,4]triazolo[4,3-*a*]pyrazines as Selective, Brain Penetrant Phosphodiesterase 2 (PDE2) Inhibitors. *ACS Med. Chem. Lett.*, 6(3):282–286, March 2015. ISSN 1948-5875, 1948-5875. doi: 10.1021/ml500463t.
- [41] Kadirvel Saraboji, Maria Håkansson, Samuel Genheden, Carl Diehl, Johan Qvist, Ulrich Weininger, Ulf J. Nilsson, Hakon Leffler, Ulf Ryde, Mikael Akke, and Derek T. Logan. The Carbohydrate-Binding Site in Galectin-3 Is Preorganized To Recognize a Sugarlike Framework of Oxygens: Ultra-High-Resolution Structures and Water Dynamics. *Biochemistry*, 51(1):296–306, January 2012. ISSN 0006-2960, 1520-4995. doi: 10.1021/bi201459p.
- [42] N. Schiering, S. Knapp, M. Marconi, M. M. Flocco, J. Cui, R. Perego, L. Rusconi, and C. Cristiani. Crystal structure of the tyrosine kinase domain of the hepatocyte growth factor receptor c-Met and its complex with the microbial alkaloid K-252a. *Proceedings of the National Academy of Sciences*, 100(22):12654–12659, October 2003. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1734128100.
- [43] Christina E. M. Schindler, Hannah Baumann, Andreas Blum, Dietrich Böse, Hans-Peter Buchstaller, Lars Burgdorf, Daniel Cappel, Eugene Chekler, Paul Czodrowski, Dieter Dorsch, Merveille K. I. Eguida, Bruce Follows, Thomas Fuchß, Ulrich Grädler, Jakub Gunera, Theresa Johnson, Catherine Jorand Lebrun, Srinivasa Karra, Markus Klein, Tim Knehans, Lisa Koetzner, Mireille Krier, Matthias Leiendecker, Birgitta Leuthner, Liwei Li, Igor Mochalkin, Djordje Musil, Constantin Neagu, Friedrich Rippmann, Kai Schiemann, Robert Schulz, Thomas Steinbrecher, Eva-Maria Tanzer, Andrea Unzue Lopez, Arielle Viacava Folis, Ansgar Wegener, and Daniel Kuhn. Large-Scale Assessment of Binding Free Energy Calculations in Active Drug Discovery Projects. *J. Chem. Inf. Model.*, 60(11):5457–5474, November 2020. ISSN 1549-9596, 1549-960X. doi: 10.1021/acs.jcim.0c00900.
- [44] Michael R. Shirts, Eric Bair, Giles Hooker, and Vijay S. Pande. Equilibrium Free Energies from Nonequilibrium Measurements Using Maximum-Likelihood Methods. *Phys. Rev. Lett.*, 91(14):140601, October 2003. doi: 10.1103/PhysRevLett.91.140601.
- [45] Alan W. Sousa da Silva and Wim F. Vranken. ACPYPE - AnteChamber PYthon Parser interfAcE. *BMC Research Notes*, 5(1):367, July 2012. ISSN 1756-0500. doi: 10.1186/1756-0500-5-367.
- [46] Miroslav Suruzhon, Michael S. Bodnarchuk, Antonella Ciancetta, Russell Viner, Ian D. Wall, and Jonathan W. Essex. Sensitivity of Binding Free Energy Calculations to Initial Protein Crystal Structure. *J. Chem. Theory Comput.*, page acs.jctc.0c00972, February 2021. ISSN 1549-9618, 1549-9626. doi: 10.1021/acs.jctc.0c00972.
- [47] Gareth A. Tribello, Massimiliano Bonomi, Davide Branduardi, Carlo Camilloni, and Giovanni Bussi. PLUMED 2: New feathers for an old bird. *Computer Physics*

- Communications*, 185(2):604–613, February 2014. ISSN 0010-4655. doi: 10.1016/j.cpc.2013.09.018.
- [48] Shunzhou Wan, Gary Tresadern, Laura Pérez-Benito, Herman Vlijmen, and Peter V. Coveney. Accuracy and Precision of Alchemical Relative Free Energy Predictions with and without Replica-Exchange. *Adv. Theory Simul.*, 3(1):1900195, January 2020. ISSN 2513-0390, 2513-0390. doi: 10.1002/adts.201900195.
- [49] Junmei Wang, Romain M. Wolf, James W. Caldwell, Peter A. Kollman, and David A. Case. Development and testing of a general amber force field. *J. Comput. Chem.*, 25(9):1157–1174, 2004. ISSN 1096-987X. doi: 10.1002/jcc.20035.
- [50] Junmei Wang, Wei Wang, Peter A. Kollman, and David A. Case. Automatic atom type and bond type perception in molecular mechanical calculations. *Journal of Molecular Graphics and Modelling*, 25(2): 247–260, October 2006. ISSN 1093-3263. doi: 10.1016/j.jmgm.2005.12.005.
- [51] Lingle Wang, B. J. Berne, and Richard A. Friesner. On achieving high accuracy and reliability in the calculation of relative protein–ligand binding affinities. *PNAS*, 109(6):1937–1942, February 2012.
- [52] Lingle Wang, Yujie Wu, Yuqing Deng, Byungchan Kim, Levi Pierce, Goran Krilov, Dmitry Lupyan, Shaughnessy Robinson, Markus K. Dahlgren, Jeremy Greenwood, Donna L. Romero, Craig Masse, Jennifer L. Knight, Thomas Steinbrecher, Thijs Beuming, Wolfgang Damm, Ed Harder, Woody Sherman, Mark Brewer, Ron Wester, Mark Murcko, Leah Frye, Ramy Farid, Teng Lin, David L. Mobley, William L. Jorgensen, Bruce J. Berne, Richard A. Friesner, and Robert Abel. Accurate and Reliable Prediction of Relative Ligand Binding Potency in Prospective Drug Discovery by Way of a Modern Free-Energy Calculation Protocol and Force Field. *J. Am. Chem. Soc.*, 137(7):2695–2703, February 2015. ISSN 0002-7863. doi: 10.1021/ja512751q.
- [53] Keith P. Wilson, Matthew J. Fitzgibbon, Paul R. Caron, James P. Griffith, Wenying Chen, Patricia G. McCaffrey, Stephen P. Chambers, and Michael S.-S. Su. Crystal Structure of p38 Mitogen-activated Protein Kinase. *Journal of Biological Chemistry*, 271(44):27696–27700, November 1996. ISSN 00219258. doi: 10.1074/jbc.271.44.27696.
- [54] H.-J. Woo and B. Roux. Calculation of absolute protein-ligand binding free energy from computer simulations. *Proceedings of the National Academy of Sciences*, 102(19):6825–6830, May 2005. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.0409005102.
- [55] Jian Zhu, Qiqi Yang, Dandan Dai, and Qiang Huang. X-ray Crystal Structure of Phosphodiesterase 2 in Complex with a Highly Selective, Nanomolar Inhibitor Reveals a Binding-Induced Pocket Important for Selectivity. *J. Am. Chem. Soc.*, 135(32):11708–11711, August 2013. ISSN 0002-7863, 1520-5126. doi: 10.1021/ja404449g.
- [56] Robert W. Zwanzig. High-Temperature Equation of State by a Perturbation Method. I. Nonpolar Gases. *J. Chem. Phys.*, 22(8):1420–1426, August 1954. ISSN 0021-9606. doi: 10.1063/1.1740409.