The maximal and current accuracy of rigorous protein-ligand binding free energy calculations

Gregory A. Ross
Schrödinger Inc,
New York, NY, 10036

Chao Lu
Schrödinger Inc,
New York, NY, 10036

Guido Scarabelli
Schrödinger Inc,
New York, NY, 10036

Steven K. Albanese
Schrödinger Inc,
New York, NY, 10036

Evelyne Houang
Schrödinger Inc,
New York, NY, 10036

Robert Abel
Schrödinger Inc,
New York, NY, 10036

Edward D. Harder
Schrödinger Inc,
New York, NY, 10036

Lingle Wang
Schrödinger Inc,
New York, NY, 10036

October 12, 2023

gregoryaross@isomorphiclabs.com
lingle.wang@schrodinger.com

Abstract

It is well recognized that computational techniques can speed up the identification of hits and accelerate the optimization of such hits to lead series and development candidate molecules. A class of rigorous physics-based methods known as free energy perturbation (FEP) have emerged as the most consistently accurate relative affinity prediction tool available to support such efforts. Yet, there remains uncertainty about how accurate these techniques are and how accurate they can ever be. Here, we assemble what we believe to be the largest publicly available dataset of proteins and congeneric series of small molecules and assess the accuracy of the leading FEP workflow in predicting relative binding affinities. To ascertain the limit of achievable accuracy, we survey the reproducibility of experimental relative affinity measurements by comparing chemical series that have been assayed by two or more different techniques. We find a wide variability in experimental accuracy and a general correspondence between binding and functional assays. When the protein and ligand structures are prepared with care, we find that FEP can achieve a level of accuracy close to what we find in our experimental survey. Throughout, we highlight reliable protocols that can help maximize the accuracy of FEP in prospective studies.

1 Introduction

Two critical objectives in drug discovery are developing molecules that bind tightly to a target protein and weakly – or not at all – to off-target proteins. There is a growing consensus that computational methods can help identify early promising compounds and aid the otherwise slow and expensive stage of lead development. In recent years, alchemical free energy calculations[1] – a family of rigorous, physics-based methods – have emerged as the most consistently accurate method available[2–4]. The purpose of this work is to quantify how accurate these types of method has become and to ascertain how close the predictions currently are to experimental accuracy.

Relative alchemical binding free energy calculations consist of a series of simulations in which the interaction and internal energies of pairs of molecules are interpolated. Statistics collected during the course of these simulations are used to produce estimates of the difference in binding free energy between the two molecules. While robust methods for computing absolute binding free energies are emerging, these also produce relative binding free energies if the free energy difference between the *apo* and *holo* protein conformations is unknown[5;6]. The term free energy perturbation (FEP) often refers to a particular class of alchemical method[7] but has in recent years been applied to alchemical binding

free energy methods more generally; henceforth, we will refer to all alchemical binding free energy methods, including thermodynamic integration, as FEP. Out of all the software programs that can perform FEP[8–15], the FEP+ computational workflow has seen the widest adoption in industry[3;16]. Although FEP+ is one program, the accuracy that it can achieve is frequently taken as an indication of what FEP methods can achieve as a whole[2].

There are numerous studies that report the successful application of FEP+ in live drug discovery projects[17–25]. Although physics-based affinity prediction methods are intrinsically computationally expensive, large clusters of graphical processing units and cloud computing enable the evaluation of relative binding free energies for potentially thousands of pairs of compounds within the typical time-constraints of medicinal chemistry deadlines. The accuracy of FEP methods have been steadily increasing due to force field improvements[26;27] and the application of enhanced sampling techniques[28;29]. Along with the increase in accuracy there has been an increase in the domain of applicability. While FEP methods are more often associated with R-group modifications[16;30], advances in methodology have meant these methods, particularly FEP+, can be applied to chemical modifications involving macrocyclization[31], scaffold-hopping[32], covalent inhibitors[33], and buried water displacement[34].

As input, FEP requires the three-dimensional structures of the protein and the putative binding geometries of the chemical series of whose absolute or relative binding free energy will be assessed. A historically difficult aspect of preparing structures has been the determination of the protonation and tautomeric states of the ligands and the protein binding residues[35]. Ambiguities in the protein structure, such as missing loops and flexible regions also pose a challenge and require careful consideration by users as to how they will be modeled. Because of these uncertainties, practitioners commonly perform a retrospective study of FEP on previously assayed compounds to test the reliability of the structural models before moving on to prospective predictions[36].

Although the main utility of FEP lies in its ability to aid live drug discovery projects, retrospective assessment studies of accuracy are, in general, vital to help identify how accurate an FEP method is and to highlight methodological aspects that require further improvement. The two largest protein-ligand benchmark data sets for FEP are the set assembled for the OPLS4 force field by Lu et al. (with a total of 512 protein-ligand pairs)[27] and the set by Hahn et al. (with 599 protein-ligand pairs)[37]; the latter of which was designed to be the community standard and has been used in the validation of the Open Force Field 2.0.0[38]. Unlike data sets such as PDBbind[39], these data sets do not contain only experimentally determined structures, and instead consist of congeneric series of ligands where all ligand binding modes and protonation states have been modeled. While these benchmarks have proved to be very useful they currently do not cover the domain of applicability of FEP, lacking, for instance, any membrane proteins, scaffold hopping transformations, and macrocyclic transformations to name a few deficiencies. Although the OPLS4 benchmark set includes charge-charging and buried water displacing transforms, the benchmark set by Hahn et al. does not. With the benchmark set by Hahn et al., the omission of certain data sets may be by design, as only data sets that met certain quality standards were included.

The apparent accuracy of FEP is fundamentally limited by the accuracy of experimental affinity measurements. The most appropriate observables with which to compare FEP predictions against are *in vitro* measurements of dissociation constants (K_ds), inhibition constants (K_is), or ligand concentrations that achieve 50% inhibition (IC_{50}s). While FEP predictions can be used to complement other experimental measurements that are not directly relatable to binding free energy, such as temperature shifts and percent inhibitions values, these can cloud the apparent accuracy of the predictions. By definition, the K_i is the K_d of an enzyme inhibitor. However, these quantities differ subtly in the way they are measured; K_is are typically measured in functional enzymatic inhibition assays whereas K_ds come from experiments that more directly measure binding[40]. Although IC_{50}s are dependent on the concentrations of the protein and ligand as well as the reaction Michaelis constant, under common assay conditions and reaction mechanisms (but not all), the ratio of the IC_{50} of two ligands is equivalent to the ratios of the K_is[41].

There is a ladder of meaning associated with experimental accuracy that has informed prior studies on this topic. At the two lowest rungs of this ladder, one can consider accuracy as being tied to the 'repeatability' of a measurement (i.e. the same assay under the same conditions with the same equipment, usually conducted by the same team), or the reproducibility of the assay itself (i.e. the same assay set up and run by different experimenters[42;43]). After reviewing an internal repository of protein-ligand activities that had been measured multiple times, Brown et al. reported a median standard deviation between measurements of 0.3 pKi units (0.41 kcal/mol)[44]. Studies have consistently found that the variance between affinity measurements made by different teams is higher compared to the variance encountered when the same team repeats the assay[45]. At the highest rung of the ladder of experimental error is the reproducibility of an affinity measurement, which is what - ideally - independent experiments could observe using different assays. To quantify the reproducibility of binding affinity measurements, Kramer et al. previously surveyed the ChEMBL database for protein ligand complexes that had affinities measured at least twice by different groups[46]. They found that the root-mean-square difference between these independent measurements ranged from 0.56 pKi units (0.77 kcal/mol) to 0.69 pKi units (0.95 kcal/mol) depending on how the data was curated.

There are numerous factors that drive the observed differences in measured binding free affinities between laboratories. These can range from concentration errors in, for instance, the reagents used in isothermal titration calorimetry experiments[47], to the difference in material of the assay containers; one study reported that a particular compound was absorbed by glass and not plastic, which artificially raised the apparent Ki of one compound[48]. Significant differences in measured binding affinities can also occur when assays are repeated with alternate instruments or when data is reanalyzed with different software[49]. Data fitting methods can be made more robust when using analysis methods that explicitly model uncertainties in the experiment[50;51].

While the efforts of by Brown et al. and Kramer et al. to quantify experimental error are informative, of particular interest to FEP is the reproducibility of relative binding affinities: the difference in absolute binding free energy between two molecules. One could produce such an estimate from the error from absolute affinity measurements by assuming every experimental measurement is unbiased with a Gaussian error distribution. These assumptions are significant, and, ideally, one would estimate experimental uncertainty of relative affinity measurements of series of compounds binding to the same protein – not single ligands – that come from different and independent assays. To date, no such study has been conducted despite a notable previous attempt[52].

In this study, we conduct two surveys to ascertain the accuracy of FEP. The first of these two surveys is concerned with the reproducibility of relative (as opposed to absolute) binding affinity measurements. We collect binding data from studies where the affinity of a series of compounds was measured in at least two different assays. The deviation between the relative binding affinities sets a lower bound to the error we can expect from any prediction method on large and diverse data sets. In the second survey, we collect protein-ligand structures and binding data from as many previous FEP validation studies as possible. In the assembly of this benchmark data set, we sought to include systems that cover the current domain of applicability of FEP methods.. We re-assess the quality of many of the modeled structures and simulation inputs to quantify the best possible accuracy that could be achieved with the FEP+ program on these data. Presented together, we hope that both surveys provide a comprehensive picture of the maximal accuracy one could expect from FEP, the current state of accuracy of FEP, and what kinds of setup procedures are robust in prospective drug discovery settings.

2 Methods

2.1 Experimental reproducibility survey

We set out to quantify the maximal accuracy that a relative affinity prediction method could achieve on large and diverse data sets that comprise of many assay types. To do so, we searched for studies where at least two different assays were used to measure the binding affinities of the same set of compounds to the same protein. We are not concerned with the difference in the absolute binding free energies as these can include systematic errors that are ultimately irrelevant for determining which ligands bind stronger than the others. Instead, we evaluated how well pairs of assays agreed in terms of relative binding free energies and rank ordering of the compounds.

Preferably, each chemical series would be evaluated in different assays conducted by independent groups. However, as this data is hard to come by, we collected comparative assay data where the measurements formed part of a single study. Table 1 summarizes the publicly available data that we analyzed which compared the affinity of the same compounds in at least two different assays. Expanded versions of this table, which include more assay information, are in tables S1, S2, and S3 of the Supporting Information (SI). To compliment this survey, we also collected comparative assay data from our own internal drug discovery projects. This data is summarized in table S4 of the SI. We were interested to see whether binding assays, such as those that directly measure dissociation constants were, on average more or less reproducible with another binding assay compared to a functional assay, such as those that measure inhibition constants.

We only considered assays that reported the affinity of ligands using dissociation constants (K_d), inhibition constants (K_i) or the ligand concentration that achieve fifty percent inhibition (IC_{50}). Between any two ligands a and b, we assumed that

$$\frac{IC_{50}^a}{IC_{50}^b} = \frac{K_i^a}{K_i^b} = \frac{K_d^a}{K_d^b},$$ (1)

which is true for a wide variety, but not all, inhibitory mechanisms and assay conditions[41;53]. To aid the comparison with predictions from FEP, the pairwise error between the relative binding free energies in the different assays was also calculated. If X is either an IC_{50}, K_d, or K_i from the same assay, the relative binding free energy between two ligands a and b is given by

$$\Delta\Delta G_{ab} = -kT \ln\left(\frac{X^b}{X^a}\right),$$ (2)

3

Table 1: The proteins where the same set of ligands had binding affinities measured in at least two different assays from our survey of publicly available data. Targets from our own drug discovery efforts are not listed but contribute to the total number of compounds in each category. The assays are categorized in terms of being 'binding' or 'functional'. In some cases, like CaII, there are multiple studies that compare different series of ligands. The number of compounds includes duplicates in the sense that the same compounds can appear in multiple assay comparisons. More details on the provenance of these data can be found in the Supporting Information.

Comparison type	Proteins in set	No. compounds
Binding vs binding	SH2, Herg, CaI, CaII, Lectin, HCV polymerase, FAK, trypsin, DPPIV, bromodomains, BPTF, galectin 3	250
Binding vs functional	Xiap, thrombin, HCV polymerase, Hsp90, FAK, DPPIV, MAPK13, AChE	1207
Functional vs functional	COT kinase, ACE, SARS-Cov2 main protease	410
	Total	1867

Table 2: The data sets where the initial structures and affinities were based on. The proteins and number of ligands are shown. The citations on the data set names show the study where the initial protein and ligand structures were taken from.

Data set name	Proteins in data set	No. compounds
FEP+ R-group set [16]	BACE1, CDK2, JNK1, Mcl1, p38, PTP1B, thrombin, TYK2	199
FEP+ charge-change [54]	CDK2, DLK, EGFR, EPHX2, IRAK4, ITK, JAK1, JNK1, PTP1B, TYK2	53
OPLS stress set [27]	BACE1, CHK1, Factor Xa	114
OPLS drug discovery [27]	A, B, C, D, E	93
Water displacement [34]	BRD4(1), CHK1, Hsp90, scytalone dehydratase, TAF1(2), thrombin, urokinase	76
FEP+ Fragments [55]	T4 lysozyme, LigA, Mcl1, MUP-1, JAK-2, hsp90, p38	79
FEP+ macrocycles [31]	BACE1, CHK1, CK2, MHT1, HSP90	34
FEP+ scaffold-hopping [32]	BACE1, β-tryptase, CHK1, ERα, Factor Xa,	17
Merck sets [56]	CDK8, cMet, Eg5, HIF-2α, PFKFB3, SHP-2, SYK, TNKS2	264
GPCRs [57;58]	A2A, OX2, P2Y1	98
Bayer macrocycles [59]	Ftase, BRD4	8
Janssen BACE1 [36;60]	BACE1	74
MCS docking [61]	HNE, Renin	49
Miscellaneous	CDK8 [62], Galectin [10;63], BTK [64], HIV1 protease [65], FAAH [66]	79
	Total	1237

where k is the Boltzmann constant and T is temperature, assumed throughout to be 300K. To compare two assays with affinity data on the same protein and set of ligands, the relative binding free energies (i.e. $\Delta\Delta G$s) were computed between all pairs of ligands within each assay, creating two sets of relative binding free energies. The root-mean-square error or mean unsigned error between these two sets of relative binding free energies provides a measure of the reproducibility of the assays. The comparison was performed for all pairs of ligands sets shown in Table 1 and Tables S1-S4.

To quantify the agreement between the rank ordering of chemical series between two different assays, we calculated the coefficient of determination (R^2) and the Kendall rank correlation coefficient (Kendall's τ) of absolute binding free energies $\Delta G_a = kT \ln(X^a)$. When X is an IC$_{50}$, these absolute binding free energies include an unknown additive constant that has no effect on rank and correlation measures.

2.2 FEP accuracy benchmark

We aimed to establish the most comprehensive FEP publicly available data set to date by collecting congeneric series from as many previously published FEP studies as possible. Two aims for this benchmark data set were to include as many ligands as possible and to cover the range targets and ligand perturbations that occur in drug discovery, such as charge changing and/or buried water displacing perturbations. By meeting these aims, we hoped that the resultant benchmark would provide a thorough test for current and future FEP methodologies. While the majority of the systems in this benchmark come from previous FEP studies, additional data sets with well resolved protein structures and ligand binding affinities that were encountered during the assembly of this data set were also included.

For inclusion in this study, we required that an X-ray structure of the protein-ligand complex exists for at least one ligand in the congeneric series, there were no significant structural ambiguities in the protein structure, and the binding data for the congeneric series was measured as K$_d$s, K$_i$s, or IC$_{50}$s. Unlike Hahn et al. [37], we did not omit chemical series on the grounds of having too few compounds or too small a dynamic range of the experimental binding free

4

energies. While having such requirements leads to tighter prediction error and correlation metric confidence intervals on individual chemical series, we are primarily concerned with the aggregated performance of FEP on the whole data set. Metrics such as overall RMSE or Kendall's tau can be calculated with high confidence even if the data sets they aggregate over are small or have a low dynamics range. We were motivated to include some of the smaller data sets because they covered specific chemical types or transformations that are not present in other larger benchmark data sets. Examples of these data sets include those focused on macrocycles[31;59], charge changing transformations[54], and buried water displacement[34].

Table 2 lists the protein-ligand data sets that were collected in this study along with the publications where they first appeared. The name given to each data set is the same as used in the Supporting Information (SI).

As the accuracy of FEP is dependent on the input protein and ligand structures, we endeavored to ensure all of the structures were of a consistently high quality. For the majority of systems, we reviewed the ligand binding geometries, the protonation states of ligand and binding site protein residues, the structure of the flexible regions of the protein, as well as the protein crystal structures themselves. Examples of where we changed the protein crystal structure include the CKD8 and SYK systems from Schindler et al.[56] and the JAK2 fragment set from Steinbrecher et al.[55]. The SI details in-depth the modifications we made to the original inputs. We note that the work by Hahn et al.[37] lists alternative crystal structures for many of the systems used in this benchmark which could be considered in future development of our benchmark set.

As the accuracy of FEP is dependent on the input protein and ligand structures, we endeavored to ensure all of the structures were of a consistently high quality. For the majority of systems, we reviewed the ligand binding geometries, the protonation states of ligand and binding site protein residues, and the structure of the flexible regions of the protein. The SI details in-depth the modifications we made to the original inputs.

Where possible, modifications were made to systems that would be plausible in a prospective setting. For instance, if it was unknown which protonation or tautomeric state the ligand was in, all states were added to the FEP map so that the relative free energies could be calculated and accounted for using our pKa correction protocol[54;67]. Similarly, if the rotameric state of a ligand chemical group was unknown, either the chemical group was decoupled in the calculations to facilitate sampling or multiple rotameric states were added to the FEP map and corrected for as described in the Appendix. Examples of the kinds of changes and modifications we made to the systems are shown in figure 1.

2.2.1 Simulation details

All simulations were conducted using FEP+ within the Schrödinger software suite (versions 21-3 and 21-4) with the OPLS4 force field and the modified SPC water model[27]. FEP+ uses replica exchange with solute tempering[28] where exchanges between neighboring replicas are attempted every 1.2 ps. By default in FEP+, the number of lambda windows in a calculation depends on the type of perturbation; charge-changing perturbation use 24 lambda windows, scaffold hopping and macrocyclization perturbations use 16 lambda windows, and all others use 12 lambda windows. In some cases from our benchmark (see the SI) more lambda windows were used.

For alchemical transformations that changed the charge of the ligands, the total charge of the simulation box was kept constant by transmuting a Na+ or Cl- ion either to water or vice versa, depending on the charge difference and perturbation direction using the scheme previously described[54]. Neutralizing counterions and a 0.15 M concentration of NaCl were to the simulation box for charge-changing perturbations; all other perturbation types had no counterions or salt added. Unless otherwise stated, each lambda window was simulated for 20 ns. Integration was performed using the multiple time-stepping RESPA integrator[68] and hydrogen mass repartitioning using the following time steps: 4 fs for bonded interactions, 4 fs for nonbonded interactions within the distance cutoff, and 8 fs for electrostatic interaction in reciprocal space.

For all simulations, the temperature was maintained at 300 K using the Nose-Hoover chains thermostat[69]. All complex leg simulations were run in the μVT ensemble whereby water molecules were sampled with grand canonical Monte Carlo; all solvent simulations were run in NPT using the Martyna-Tobias-Klein barostat[70] was used to maintain the pressure at 1 atm. The use of this enhanced water sampling procedure in the complex legs negated the need to assess the sensitivity of the predictions on the starting positions of water molecules[34].

When preparing proteins and ligands for FEP+, the Schrödinger protein preparation wizard was used. All crystallographic water molecules were retained and missing side chains or loops were added with Prime. Protonation state assignment was carried out with PROPKA[71] and manual inspection. A detailed description is provided in the SI for systems that required more involved preparation and analysis.

5

2.2.2 Metrics used for analysis

Relative FEP methods calculate the binding free energy difference ($\Delta\Delta G$) between pairs of structurally similar ligands. For a series of N ligands binding to the same receptor, in principle, simulations can be performed for all of the $N \times (N-1)/2$ pairs of ligands. However, this is prohibitively expensive when N is large. Instead, it is standard practice to calculate the $\Delta\Delta G$s for a small subset of the possible pairings. The set of perturbation pairs and ligands makes a graph with each perturbation pair forming an edge.

Many previous studies have reported FEP accuracy using either the mean unsigned error (MUE) or root-mean-square error (RMSE) between the calculated and experimental $\Delta\Delta G$s for each edge. However, these so-called edgewise errors are dependent on the topology of the perturbation graph. As perturbation graphs are usually constructed to have edges between similar ligands, edgewise errors are correspondingly limited to quantifying FEP accuracy for similar ligands. A more robust alternative is to instead report the inferred predicted $\Delta\Delta G$s between every pair of ligands in the graph. To ensure a consistent set of $\Delta\Delta G$ between all pairs of ligands, FEP+ uses the cycle closure correction algorithm[72]; there exist other methods for doing so[73;74]. In addition to the edgewise RMSE, we also calculated these pairwise RMSEs for every FEP map in our benchmark. For a collection of M FEP graphs, we calculated the weighted average of the RMSEs using the following formula:

$$\text{RMSE} = \sqrt{\frac{1}{\sum_i^M w_i} \sum_i^M w_i \text{RMSE}_i^2}, \tag{3}$$

where w_i was the weight applied to the ith graph RMSE. For edgewise RMSEs, w_i was set equal to the number of edges in each graph and for pairwise RMSEs w_i was set equal to the number of compounds in the graph. The latter weighting was also used when computing the aggregate RMSE in the experimental survey. We note that this weighting scheme, along with the use of pairwise errors, was used in the FEP assessment by Schindler et al.[56].

Using the cycle closure correction algorithm, the absolute binding free energies (ΔGs) - up to an unknown constant - for the ligands in each graph were determined. Metrics such as the R^2 or Kendall's τ of these ΔGs do not depend on this unknown constant. Unlike correlation statistics for the predicted $\Delta\Delta G$s, correlation statistics for ΔGs provide a direct measure of the rank order ability of FEP. We do not report the correlation statistics for the predicted $\Delta\Delta G$s because, as illustrated by Hahn et al[37], these statistics are dependent on the sign of the arbitrary sign of the $\Delta\Delta G$s. A weighted average of R^2 or Kendall's τ of the ΔGs was calculated across every graph in our data set, where the weight was equal to the number of ligands in the graph.

3 Data Availability

The experimental error survey and the FEP benchmark data sets that do not come from our internal drug discovery projects, along with the code to reproduce the analyses and plots, are available on a open access Github repository[75].

4 Results

4.1 Using FEP to resolve ambiguities in the structural inputs

During the course of our FEP benchmarking exercise, we found that significant gains in accuracy could be made when different structural inputs could be all treated within the FEP workflow. The SI provides extensive details of structural modifications and additional FEP calculations that were applied to each system. Figure 1 shows four examples where FEP was used to resolve ambiguities in the input structures. The top of figure 1 shows MCL1, where several ligands (such as ligand 35 in the figure) had additional rotamers added to the map. Other groups have also considered alternate rotamer states for this ligand series[76]. Augmenting the map with additional rotamer states in MCL1 and applying the binding mode correction as detailed in Appendix A reduced the pairwise RMSE from 1.41 kcal/mol to 1.24 kcal/mol. In cases such as this, replica exchange solute tempering was not sufficient to lower the sampling barriers enough to facilitate complete rotamer sampling; extra rotamers were added to the perturbation map when poor rotamer sampling was observed in prior simulations using the automatically generated FEP+ analysis panel. Second from the top of figure 1 shows TNKS2, where a subset of the ligands had titratable amines (such as ligand 8a in the figure). Adding both the protonated and deprotonated forms of these ligands to the map reduced the pairwise RMSE from 2.10 kcal/mol to 1.60 kcal/mol.

The protonation states of side chains were also validated using the protein residue mutation functionality of FEP+. This approach is exemplified by PTP1B (second from bottom in figure 1). In this system, the sulfur atom of a cysteine

Figure 1: Examples from the FEP benchmark where the calculations were augmented. Where possible, FEP was used to automatically determine the preferred rotamer state of R-group modifications (like with MCL1, top), protonation state (like with TNKS2, second from top), the protonation state of side chain residues (like with PTP1B, second from bottom), and where atom mappings were altered to enhanced conformation sampling (like with CK2, bottom). MCL1 exemplifies cases where multiple orientations of R-groups were added to the FEP map and post-processed based on equation 4 in the Appendix. TNKS2 exemplifies cases where multiple protonation states were post processed based on our pKa correction workflow[54]. In PTP1B, the sulfur atom is in close proximity to a number of backbone NH groups (not shown) which are likely to be the the main drivers of the predicted negatively charged protonation state.

residue sits within a bowl of backbone NH groups but it is also in close proximity to the carboxylic acid group of the ligands; this mixed electrostatic environment makes the determination of the cysteine protonation state nontrivial. In our previous FEP+ validation studies, we treated this binding site cysteine as being deprotonated (i.e. negatively charged). This decision has since been called into question[10]. Rather than choosing a particular protonation state for CYS 215, we calculated its pKa using FEP+ in the presence of 4 representative ligands. The pKa of CYS 215 ranged from 0.92 to 1.66 across the four ligands so we continued to treat it as deprotonated in our calculations (the pairwise RMSE of the map was 0.74 kcal/mol).

As previously noted by Paulsen et al.[77], perturbations involving macrocyclization can benefit from reducing torsion barriers to enhance sampling and thereby more fully predict strain energy differences. We also observed such benefits in our macrocyclization calculations, such as CK2 shown at the bottom of figure 1. The acyclic ligand binds with an amine bond in the relatively high energy *cis* conformation, but the high torsional barrier height prevents the switching between *cis* and *trans* conformations in solvent. An atom mapping was chosen to place the amine in the alchemical

7

region and torsion energies in the alchemical region were scaled to zero in the intermediate λ windows. Without this particular atom mapping, the error between the acylic ligand and macrocycle was 3.4 kcal/mol. With the scaled torsion angle, the error reduced to 1.25 kcal/mol. The improvement in error comes from better capturing the relative strain energy between the acyclic and macrocyclic ligands. While the changes described here came from visual inspection and manual intervention, an automated approach for these kinds of perturbations would be preferable.

4.2 Comparing the accuracy of FEP+ with experimental reproducibility

To evaluate the maximal accuracy one could ever expect from relative binding free energy prediction methods, we conducted a survey on the reproducibility of experimental affinity assays using publicly available data as well as data from our own drug discovery projects. This survey focused on how well different assays agreed with regards to rank ordering and relative binding free energies (i.e. how much stronger one compound bound compared to another). Assay comparisons involved both binding assays and functional assays. Table 3 summarizes the reproducibility in terms of root-mean-square error (RMSE) and mean unsigned error (MUE) of the relative measurements, as well as with the coefficient of determination (R^2) and Kendall rank coefficient (Kendall's τ). These metrics were aggregated by weighting each assay comparison by the number of compounds; confidence intervals were calculated by bootstrap sampling over each assay comparison.

Table 3 also shows the aggregated statistics from our large FEP benchmark after curation. The pairwise errors and correlation metrics of the experimental survey and FEP+ benchmark are directly comparable. whereas the edgewise errors only apply to the FEP predictions as they are dependent on the topology of the perturbation graphs. Given the sampling noise, uncertainty in the structural inputs, and force field error, the pairwise RMSE and MUE of FEP+ is, perhaps surprisingly, close to the experimental survey weighted average. The weighted experimental error values are more uncertain, however, which are indicated by the larger range in the bootstrap confidence intervals. The correlation and rank statistics are higher by a statistically significant degree in the experiment survey than in the FEP+ benchmark. Figure 2 shows correlation plots from the experimental assay comparisons and FEP+ predictions against experiment that are representative of the best, average, and worst RMSE.

Table 3: Summarizing the reproducibility of the experimental relative binding free energies and the accuracy of FEP+. The value of every metric, such as RMSE or R^2, is a weighted average. For the pairwise, R^2, and Kendall τ metrics, the weighting is equal to the number of compounds in the assay (in the experimental survey) or FEP graph. For the edgewise errors, the weighting is equal to the number of edges in each FEP graph. Square brackets encompass 95% confidence intervals that have been calculated by bootstrap sampling over the pairs of experimental series or FEP+ graphs. As the edgewise error is dependent on the topology of an FEP+ graph, there is no equivalent metric in the experimental survey.

Accuracy metric	Experimental survey	FEP+ benchmark
Pairwise RMSE (kcal/mol)	0.91 [0.83, 1.11]	1.25 [1.17, 1.33]
Pairwise MUE (kcal/mol)	0.67 [0.61, 0.83]	0.98 [0.91, 1.05]
Edgewise RMSE (kcal/mol)	N/A	1.17 [1.08, 1.25]
Edgewise MUE (kcal/mol)	N/A	0.91 [0.84, 0.98]
R^2	0.79 [0.75, 0.82]	0.56 [0.51, 0.60]
Kendall τ	0.71 [0.65, 0.74]	0.51 [0.48, 0.55]

In both the experimental survey and FEP benchmark, we endeavored to remove measurements that were clearly below or above the assay detection limits. These kinds of data points were identifiable as vertical or horizontal lines in scatter plots. In the Covid Moonshot study[78], the largest publicly available data set in the experimental survey, this removal only modestly reduced the calculated pairwise RMSE from 0.85 kcal/mol (using all 528 pairs of measurements) to 0.79 kcal/mol (using 324 pairs of measurements). Although measurement accuracy is generally lower close to, but not beyond, assay detection limits, we did not try to correct for these more subtle cases as we expect these will affect both the experimental survey and FEP benchmark similarly.

FEP accuracy varies depending on the method used for the calculation. To quantify this variability and the degree to which this is affected by the size and heterogeneity of the data set, we re-analysed the predictions from FEP+ predictions from 2015 up to 2021, which encompasses sampling and force field improvements. The results, in section 3 of the SI, indicate that larger, more diverse data sets are more able to discriminate between the accuracy of different approaches.

4.2.1 The distribution of errors

Figure 3 shows boxplots of the RMSEs from each assay comparison in the experimental survey and each FEP graph from the benchmark. Notably, the experimental survey contains one assay comparison that has an RMSE greater than 2.5 kcal/mol, which is larger than any of the errors encountered in the FEP benchmark. This data point is a comparison
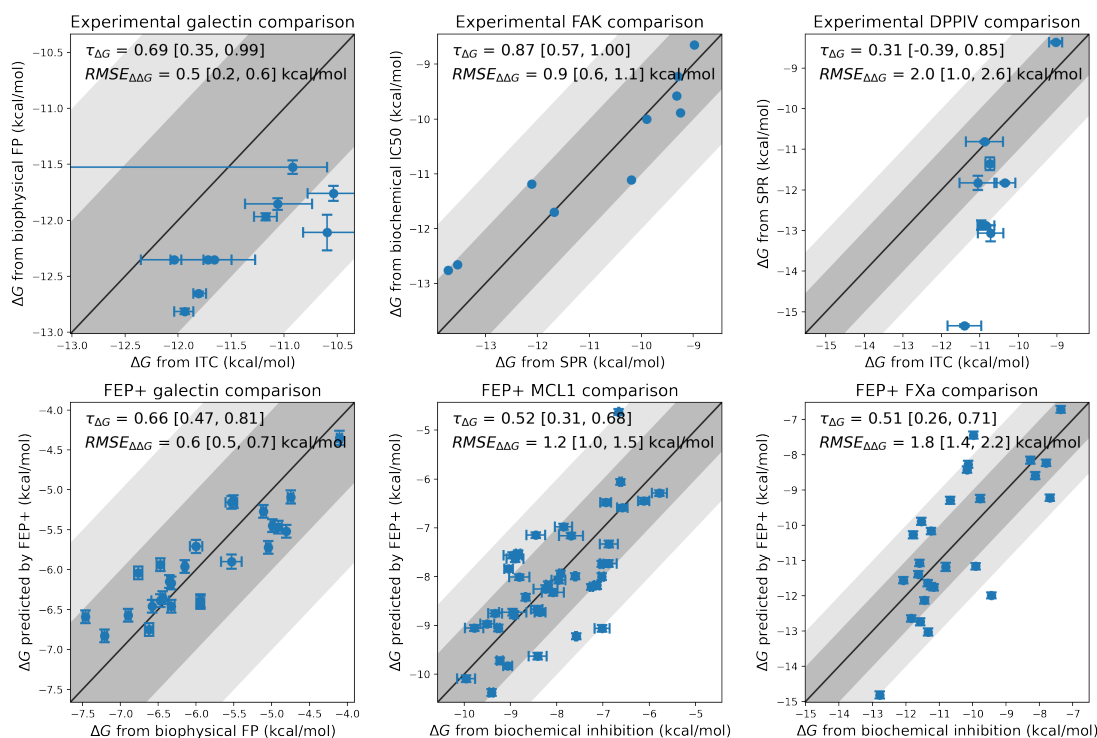
8

Figure 2: Scatter plots showing the range of agreement of ΔGs between different experimental assays (top row) and the agreement between FEP+ predictions and experiment (bottom row). The leftmost column shows examples where the pairwise RMSE of relative binding free energies was much better than average, the middle column shows examples where the RMSE was close to the average, and the rightmost column shows examples where the RMSE was worse than average. The top left of each plot shows Kendall τ and pairwise RMSE for each data set. Points in the dark gray area are measurements or predictions that are within 1 kcal/mol of each other, and points in the light gray area agree within 2 kcal/mol. The top left plot shows that isothermal titration calorimetry (ITC) and fluorescence polarization (FP) binding free energy measurements of galectin ligands are offset by around 1 kcal/mol - this offset is irrelevant for rank ordering and does not affect the correlation or the pairwise RMSE metric. The offset of the FEP+ ΔG predictions was determined by ensuring the mean of the ΔGs was equal to the mean of experimental ΔGs on the x-axis. Where the data was available, we included the reported error of the experimental measurements; the standard errors from FEP+ as calculated with the cycle-closure algorithm are also indicated on the bottom row of plots.

between surface plasmon resonance (SPR) and mass spectrometry on a series of carbonic anhydrase I ligands. Clearly, some assays can differ widely in their measured relative affinities, and so, one should expect some FEP graphs to have a large apparent error based on experimental error alone. While some graphs have RMSEs near 2 kcal/mol in our FEP benchmark (see figure 2), none approach 2.5 kcal/mol as in the experimental survey. The boxplots show that, while Table 3 states that the mean pairwise error of FEP+ across the whole benchmark data set is 1.26 kcal/mol, the error of FEP+ on individual graphs may be lower or higher than this value.

The right panel of figure 3 shows the histograms of all pairwise (not edgewise) errors from both the experimental survey and the FEP benchmark. In all 13,732 FEP+ relative binding free energy predictions, 57.5% were within 1 kcal/mol of the experimental value and 12.9% differed by more than 2 kcal/mol, which are close – but not equal to – the percentages one would expect from a Gaussian distribution that has a standard deviation equal to the pairwise RMSE of FEP+. Out of 314,535 relative binding free energies in the experimental survey, 83.0% were within 1 kcal/mol of each other and 2.1% differed by more than 2 kcal/mol; there are almost twice as many errors above 2 kcal/mol than would be expected from a Gaussian distribution. The experimental error distribution therefore has a "fatter" tail than a Gaussian distribution that is better modeled using a t-distribution. The maximum likelihood estimate of the *degrees of freedom* parameter of the t-distribution for the experimental data was 6.0, which was lower than the estimated value of 25.7 for FEP+ errors. A smaller degrees of freedom value implies fatter tails, meaning that, although the experimental error distribution is tighter than the FEP+ error distribution, there was a higher propensity for significant, non-Gaussian differences in the experimental survey.
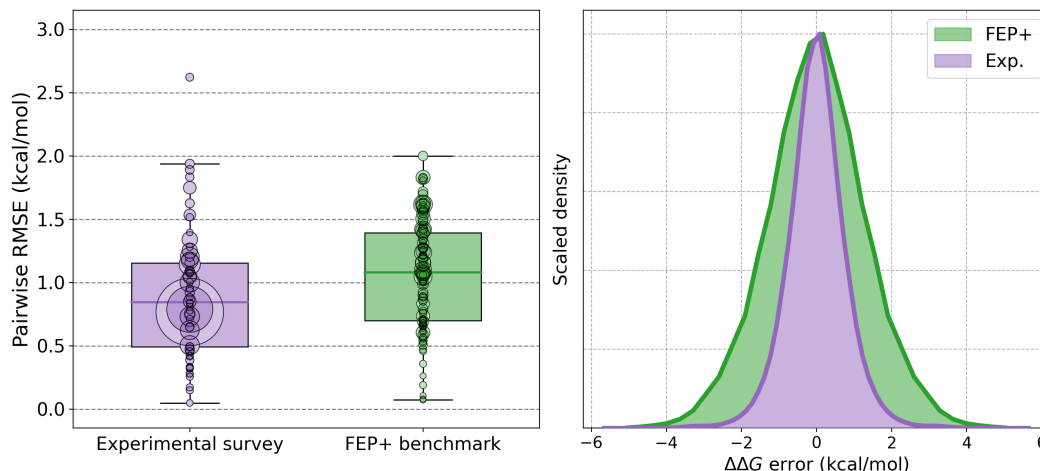
9

Figure 3: The left panel shows boxplots comparing the root-mean-square error (RMSE) between relative binding free energies from different experimental assays (left) and the FEP+ predictions against experimental data (right). In the boxplots on the left, the boxes represent the 25th and 75th percentiles and the dark line represents the median. The whiskers extend to a maximum of 1.5 times the interquartile range. The size of each data point is proportional to the number of ligands in the series in either an assay comparison or perturbation graph. The two largest data points in the experimental survey are from the COVID moonshot project[78] and project A from table S4. The median RMSE in the experimental survey is 0.85 kcal/mol and the median in the FEP+ benchmark is 1.08 kcal/mol. The right plot shows the all pairwise relative binding free energy differences from the experimental survey and all pairwise FEP+ errors. The histograms were symmetrized about the x=0 line in the sense that all $N \times (N-1)$ pairs of compounds were used. The error distributions are bell-shaped and can be approximated by t-distributions.

4.2.2 Validation of the experimental error estimates

It is worth quantifying how our aggregated experimental RMSE compares to previous estimates. After trawling the ChEMBL database for binding affinities that had been measured by at least two different groups, Kramer et al. arrived at two estimates for experimental error[46]. When all pairs of measurements were included in their estimate, they calculated the reproducibility error of absolute binding free energies to be 0.69 pKi (0.95 kcal/mol). Their second error estimate of 0.56 pKi (0.77 kcal/mol) was arrived at by discarding all differences that were greater than 2.5 pKi units. If one assumes experimental error to be unbiased (i.e. there is no offset between two sets of ΔGs) and Gaussian distributed, then the estimates of Kramer et al. imply pairwise RMSEs of 1.34 kcal/mol and 1.09 kcal/mol, which are arrived at by multiplying their estimates by the square root of 2. The second of these estimates is in close agreement with our estimated value in table 3 but the first is above our upper 95% confidence limit. Nevertheless, these assumptions used to derive pairwise RMSEs from Kramer's estimates should be viewed cautiously as we have found that in some of our assay comparisons, such as in top-left scatter plot in figure 2, one set of experimental ΔGs can be offset from the other. Also, as described above, we have found our pairwise error distribution is better modeled as t-distribution rather than a Gaussian. Indeed, Kramer et al. originally noted that the absolute ΔG error distribution from the ChEMBL set was poorly approximated by a single Gaussian distribution.

Our publicly available experimental survey contains comparative assay studies from a wide range of academic and industrial laboratories, all of which may have different levels of quality control. On the other hand, data from our own drug discovery projects should be more consistent in terms of quality control. To see whether this is reflected in the experimental error estimates, we can split the data sets and calculate the experimental error separately. Using only the data from our internal projects, we calculate the weighted RMSE to be 0.88 [0.80, 1.13] kcal/mol, compared to 0.96 [0.83, 1.24] kcal/mol from the publicly available data. As the bootstrap confidence intervals overlap substantially, we cannot distinguish between the quality of the data sources. The broad consistency of these estimates suggest that experimental error values shown in table 3 are generally representative.

The primary interest of our experimental survey is to quantify the reproducibility of experimental binding affinities, which we have approached by quantifying the difference of (relative) binding affinity measurements from different experiments. Nevertheless, it is of interest to place the experimental RMSEs in context of the uncertainty that arises from multiple repeats of the same experiment. In the table S5 of the SI, we collect the reported standard deviations from repeats from a total of 350 compounds spread over 15 experimental assays. The overall root-mean-square of these reported standard deviations is 0.23 [0.18, 0.33] kcal/mol, where the square brackets denote 95% confidence intervals that have been calculated with bootstrap sampling over the different assays. The uncertainty from the repeats of a

single experiment contributes to RMSE when comparing different experiments. Assuming the repeatability error we have calculated is Gaussian distributed and applies to all experiments, the RMSE from taking the difference from two measurements is approximately 0.33 kcal/mol (from the square root of 2 times 0.23 kcal/mol). This value is roughly a third of the 0.91 kcal/mol reproducibility RMSE from table 2, which implies that two thirds of the reproducibility RMSE comes from intrinsic differences from the different experiments.

4.2.3 The variability between different assay types

Our experimental survey also permitted an assessment on the agreement between binding assays, such as surface plasmon resonance (SPR), and functional assays, such as enzymatic activity assays. As FEP predicts relative binding free energies, one could expect a better apparent accuracy with binding assays than with functional assays. Previously, Schindler et al. found that with one protein and chemical series, the predictions from FEP+ were in closer agreement to the measurements from SPR than from a functional assay [56].

Table 4 compares the experimental reproducibility between binding and functional assays and combines the data from our experimental error survey (tables S1-S3) and from our internal program analysis (table S4). When ignoring the confidence intervals, the pairwise RMSE between binding assay derived $\Delta\Delta G$s and functional assay derived $\Delta\Delta G$s appears lower than the pairwise RMSE between two binding assays. This is a result of project A from table S4 dominating the weighted mean because of its large number of compounds. The confidence intervals, as with all others in this study, have been calculated by bootstrap sampling over each assay comparison and minimize the effect of any one comparison. Clearly, there is substantial overlap between the confidence intervals between the 'binding vs binding' and 'binding vs functional' comparisons. Thus, these data support the hypothesis – summarized in equation 1 – that $\Delta\Delta G$s from binding assays are consistent with $\Delta\Delta G$s from functional assays when considering a diverse range of experiments and targets. With regards to what data is best for validating relative binding free energies predictions from FEP, these results show that, in the main, functional measurements of affinity are as appropriate as binding measurements. Having more comparative assay data may change these conclusions, and these results do not preclude the existence of large differences between binding and functional assays that occur on a case-by-case basis, such as with project B in table S4.

Table 4 suggests that functional assays are more consistent with each other than they are with binding assays. However, there are far fewer comparisons in the "functional vs functional" category, which makes these weighted means and confidence intervals less meaningful than the others.

Table 4: Comparing the agreement between binding and functional assay measurements of relative binding free energies. As each assay type differs in what is measured, in the sense that binding may not always result in inhibition, one may expect a larger disagreement between the two types than within the types. The confidence intervals, calculated by bootstrap sampling over the different assay comparisons, show that the differences we have obtained are not statistically significant. The "No. comparisons" column shows how many assays were compared to estimate the reproducibility and bootstrapped over to estimate the confidence intervals.

	No. comparisons	Pairwise RMSE (kcal/mol)	R^2	Kendall τ
Binding vs binding	26	1.10 [0.85, 1.34]	0.76 [0.65, 0.84]	0.69 [0.61, 0.76]
Binding vs functional	30	0.93 [0.82, 1.21]	0.81 [0.73, 0.83]	0.76 [0.62, 0.75]
Functional vs functional	6	0.75 [0.53, 0.79]	0.78 [0.75, 0.91]	0.70 [0.58, 0.81]

5 Discussion

We have assembled what is to our knowledge the largest benchmark data set for free energy perturbation (FEP) calculations of relative binding free energies to date. As prediction accuracy is only meaningful in the context of experimental accuracy, we conducted a survey of experimental reproducibility alongside our FEP benchmark. While other studies have looked at the experimental differences in absolute binding free energies [44,46], our survey focused on reproducing rank ordering and relative binding free energy measurements. Our survey used sets of ligands that had binding affinities determined in at least two different assays.

We found that the accuracy of FEP+ was close to experimental accuracy in terms of relative binding free energies while experimental rank ordering ability was found to be markedly higher than FEP+. As extensively detailed in the supporting information and summarized in figure 1, our FEP+ results were obtained by preparing and studying the systems as thoroughly as time constraints allowed. These involved augmenting the FEP graphs with additional protomeric and tautomeric states, as well as with adding additional rotamer states of R-groups, or choosing perturbation mappings that enhanced sampling. Other changes involved improving the binding modes or protein structure. These kinds of treatments, while sometimes straightforward in retrospect, can be difficult to tease out in drug discovery programs with

time constraints, where predictions must come out in a timely fashion for synthesis and assaying. However, FEP is an "all-in-one" method that can predict protonation states, ligand poses, and binding affinities, so a strategy that is appropriate for drug discovery is enumerate FEP maps with all likely states and use post-processing, as we have here, to correct for the different protonation and rotamer states.

For our FEP benchmark, we primarily collected data sets from previous publications on FEP. This potentially biases our estimates of FEP error as prior publications are more likely to contain systems where a particular FEP method appeared at least satisfactory in terms of accuracy. It should also be noted that our estimate for the error of FEP+ with respect to experiment was estimated retrospectively and that reported errors are typically higher in prospective applications [56]. These biases may partly reflect the results in figure 3, where the largest pairwise error was found in a comparison of different experimental assays and not FEP graphs. Our experimental survey clearly demonstrates that some experimental affinity measurements can differ substantially from others. FEP will have an apparently high error if compared against irreproducible experimental affinity measurements, and these cases will surely be encountered by increasing the amount of data in the FEP benchmark. It should also be noted that given the heterogeneity of the assay quality used both in the experimental survey and FEP benchmark, our estimates of the maximal and current accuracy primarily apply to this same regime of assay quality. While our benchmark is the largest to date, we will continue our efforts to improve its size and representation of drug discovery targets.

FEP is most useful when it can be applied to the kinds of designs that are made by medicinal chemists, which is why we have included a wide variety of transformation types in our benchmark, such as those involving fragments, scaffold-hopping, charge-change, macrocyclization, and water displacement. The range of systems and transformations contained in our benchmark covers a wider domain of applicability than previous benchmark sets, namely the OPLS4 set [27] and the set by Hahn et al. [37]. Our FEP benchmark set also contains over twice as many ligands as either of these data sets. With these considerations in mind, we hope our benchmark dataset provides the most comprehensive test of an FEP method to date. It is also our hope that FEP benchmark datasets will continue to grow in size, ligand diversity, and target coverage. The benchmark data set we present here is a step forward in this respect, although there remain areas for improvement. The number of membrane proteins in this benchmark set could be increased in future iterations to better reflect the distribution of drug targets, but we note that the aforementioned prior benchmark sets did not contain any membrane proteins. Other types of ligand transformations, such as those involving transition-metal complexes which are actively being developed were omitted from this dataset [79] and should be included in future efforts. Although we reviewed the quality of many of the protein and ligand structures in this work, this aspect of the data is an evolutionary process and the structures should remain under review in future versions of this benchmark.

While the focus of this work has been on the analysis of binding free energies for protein-ligand complexes, it is important to acknowledge that these free energy calculations have a wide range of other applications, such as estimating covalent reaction kinetics [33], small molecule solubilities [80], protein and ligand pKas [67;81], and the stability of protein mutations [13;82;83]. In each application, we believe it is important to frame reported accuracies in the context of experimental uncertainty. To help encourage more complete validations of free energy methods in future, as part of this manuscript, we are releasing the publicly available protein and ligand structures that were used in this benchmark as well as the data from our experimental accuracy survey.

As our experimental reproducibility survey indicated that experimental RMSE is on average around 1 kcal/mol, it would be an extraordinary challenge to ever have an FEP method that achieves an error truly statistically indistinguishable from the experimental error on a large and diverse data set, such as we have produced here. However, given the continual development of new methods and protocols, it is an open question as to how close predictions could ever get to the limits set by the inconsistent quality of the target experimental data. In the meantime, given our experiences in assembling and curating this benchmark, the greatest gains in accuracy that present day users can achieve will come from following the best practices that we have attempted to further codify here.

6 Supporting Information

The supporting information contains many details relating to the experimental error survey and FEP benchmark. The experimental survey data includes tables for each assay comparison time that state the source, protein, and assay types for the publicly available data. Summary tables of Schrödinger's in-house drug discovery experimental error data are also summarized in table S4, although the raw data itself has not been made available on the Github repository.

The supporting information on the FEP benchmark contains discussions and validation results on all systems and chemical series that were modified relative to the original publication. The SI contains figures that summarize our work on each system, tables that summarize the accuracy of FEP+ on each series, and tables that contain additional analyses, such as ligand and protein pKa calculations.

7 Author contributions

G.A.R. and L.W. conceived the concept with input from R.A. and E.D.H. G.A.R. assembled and simulated the FEP benchmark data set with assistance from C.L and G.S. FEP dataset curation was performed by G.A.R with guidance from L.W. and E.D.H.. G.A.R assembled the publicly available experimental reproducibility data; S.K.A and E.H. provided the experimental reproducibility data from in-house drug discovery programs. G.A.R performed all the analyses. G.A.R. wrote the manuscript with help from L.W.. All authors discussed the results and commented on the manuscript.

8 Competing interests

The authors declare no competing interests.

9 Acknowledgements

10 Data and code availability

The FEP benchmark data set, the publicly available experimental survey data, and analysis code are freely available on the Github repository github.com/schrodinger/public_binding_free_energy_benchmark. The results in this manuscript use v1.0 of the repository.

A The expanded conformer approach

Before running FEP, users have to decide on the binding geometry of the ligands. Typically, users opt for a one single structure per ligand that has been predicted by docking or by similarity to a crystal structure. If there are multiple possible poses to choose from, one could use FEP to select which pose to use, either via a relative binding free energy protocol[84] or absolute binding free energy protocol[85]. Selecting the binding geometry using free energy calculations benefits from the increased sampling and accuracy afforded by all-atom simulations. A significant drawback to selecting binding geometries in this way is that they require restraints that maintain the geometries somewhat fixed during the calculations. The use of restraints can be cumbersome and requires the application of corrections in the calculated free energies.

For ligands that have more than one possible binding conformation, we sought an approach to avoid using restraints that could easily be incorporated into the FEP+ workflow. We restricted our attention to cases where the binding site of the protein naturally restricts the conformation of the ligand in the complex simulations but where the ligand can freely sample conformations in the solvent simulations. Here, we treat additional binding modes as extra ligands within the FEP+ workflow. This way, as well as outputting the calculated binding free energies between the ligands, FEP+ will also calculate the relative binding free energies between different binding modes. Rather than selecting a single pose to use, we apply a statistical mechanical correction, outlined below, that handles all binding conformations at once.

Consider an edge that connects 2 different binding conformations of the same ligand. If, in the solvent simulations, the ligand adopts a distribution of conformations that is independent of the starting conformation, then the calculated solvent free energy difference would be zero. If, in the complex simulations, the ligand cannot adopt the alternate conformation when the simulation was started in the other, then the binding free energy difference is equal to the relative binding free energy of the two confirmations.

Without additional processing, FEP+ will treat each binding conformation as having a different predicted binding free energy and so a correction is required to account for the additional binding conformations. With the cycle closure correction algorithm, FEP+ produces absolute binding free energies up to an unknown additive constant. Let ΔG_i denote this absolute binding free energy for a particular ligand in binding conformation i and let there be a total of n binding conformations for that ligand. If ΔG denotes the absolute binding free energy for the ligand that fully accounts for all the binding modes, then ΔG_i and ΔG are related to one another by the following equation:

$$\Delta G = \Delta G_i - kT \ln \left(1 + \sum_{i \neq j}^{n} e^{-\Delta\Delta G_{ij}/kT} \right), \tag{4}$$

where $\Delta\Delta G_{ij}$ is the relative binding energy between conformations i and j. This equation is valid as long as the binding conformations do not interconvert in the complex simulations but can readily interconvert in the solvent simulations. The use of equation 1 is only approximate if the solvent simulations do not produce conformational ensembles that are independent of the starting conformation. Additionally, the estimate ΔG will be erroneously biased to more negative free energies if the conformations in the complex can interconvert.

A.1 Correcting under-sampled symmetric rotamers

Equation 4 is general and can be applied to multiple conformations of any type. One type of conformational degree of freedom are R-groups that have rotational torsion symmetry. A common example is a terminal phenyl group, which has a two-fold rotational symmetry. An estimate of the binding free energy that is too positive will occur if an R-group with rotational symmetry samples its symmetric rotamer torsions freely and equally in solvent but is stuck in its starting torsion angle in the protein. As the $\Delta\Delta G_{ij}$ between two symmetric rotamer states i and j is zero, equation 3 tells us that we have underpredicted the binding free energy by $kT \ln(2)$. Thus, cases like this need to have their binding free energy lowered by $kT \ln(2)$.

References

[1] Antonia S.J.S. Mey, Bryce K. Allen, Hannah E. Bruce Macdonald, John D. Chodera, David F. Hahn, Maximilian Kuhn, Julien Michel, David L. Mobley, Levi N. Naden, Samarjeet Prasad, Andrea Rizzi, Jenke Scheen, Michael R. Shirts, Gary Tresadern, and Haufeng Xu. Best practices for alchemical free energy calculations [article v1.0]. *Living Journal of Computational Molecular Science*, 2, 2020. ISSN 25756524. doi:10.33011/livecoms.2.1.18378.

[2] Bernd Kuhn, Michal Tichý, Lingle Wang, Shaughnessy Robinson, Rainer E. Martin, Andreas Kuglstatter, Jörg Benz, Maude Giroud, Tanja Schirmeister, Robert Abel, François Diederich, and Jérôme Hert. Prospective

Evaluation of Free Energy Calculations for the Prioritization of Cathepsin L Inhibitors. *J. Med. Chem.*, 60(6): 2485–2497, 2017. ISSN 15204804. doi:10.1021/acs.jmedchem.6b01881.

[3] Robert Abel, Lingle Wang, Edward D. Harder, B. J. Berne, and Richard A. Friesner. Advancing Drug Discovery through Enhanced Free Energy Calculations. *Acc. Chem. Res.*, 50(7):1625–1632, 2017. ISSN 15204898. doi:10.1021/acs.accounts.7b00083.

[4] Katharina Meier, Joseph P. Bluck, and Clara D. Christ. Use of free energy methods in the drug discovery industry. In Kira A. Amarcost and David C. Thompson, editors, *Free Energy Methods in Drug Discovery: Current State and Future Directions*, chapter 2, pages 39–66. ACS Publications, 2021.

[5] Wei Chen, Di Cui, Robert Abel, Richard A. Friesner, and Lingle Wang. Accurate calculation of absolute protein-ligand binding free energies. *ChemRxiv*, 2022. doi:10.26434/chemrxiv-2022-2t0dq-v2.

[6] Y Khalak, G Tresadern, M Aldeghi, H M Baumann, D L Mobley, B. L. de Groot, and V. Gapsys. Alchemical absolute protein-ligand binding free energies for drug design. *Chem. Sci.*, 12(41):13958–13971, 2021. ISSN 20416539. doi:10.1039/d1sc03472c.

[7] William L. Jorgensen and C. Ravimohan. Monte carlo simulation of differences in free energies of hydration. *The Journal of Chemical Physics*, 83:3050–3054, 9 1985. ISSN 0021-9606. doi:10.1063/1.449208.

[8] Lin Frank Song, Tai Sung Lee, Chun Zhu, Darrin M. York, and Kenneth M. Merz. Using AMBER18 for Relative Free Energy Calculations. *J. Chem. Inf. Model.*, 59(7):3128–3135, 2019. ISSN 15205142. doi:10.1021/acs.jcim.9b00105.

[9] Willem Jespers, Mauricio Esguerra, Johan Åqvist, and Hugo Gutiérrez-De-Terán. Qligfep: An automated workflow for small molecule free energy calculations in Q. *J. Cheminform.*, 11(1):1–16, 2019. ISSN 17582946. doi:10.1186/s13321-019-0348-5. URL `https://doi.org/10.1186/s13321-019-0348-5`.

[10] 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, 2020. ISSN 20416539. doi:10.1039/c9sc03754c.

[11] 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, 2020. ISSN 15205142. doi:10.1021/acs.jcim.0c00165.

[12] Tai Sung Lee, Bryce K. Allen, Timothy J. Giese, Zhenyu Guo, Pengfei Li, Charles Lin, T. Dwight McGee, David A. Pearlman, Brian K. Radak, Yujun Tao, Hsu Chun Tsai, Huafeng Xu, Woody Sherman, and Darrin M. York. Alchemical binding free energy calculations in AMBER20: Advances and best practices for drug discovery. *J. Chem. Inf. Model.*, 60(11):5595–5623, 2020. ISSN 15205142. doi:10.1021/acs.jcim.0c00613.

[13] Zhixiong Lin, Junjie Zou, Shuai Liu, Chunwang Peng, Zhipeng Li, Xiao Wan, Dong Fang, Jian Yin, Gianpaolo Gobbo, Yongpan Chen, Jian Ma, Shuhao Wen, Peiyu Zhang, and Mingjun Yang. A Cloud Computing Platform for Scalable Relative and Absolute Binding Free Energy Predictions: New Opportunities and Challenges for Drug Discovery. *J. Chem. Inf. Model.*, 61(6):2720–2732, 2021. ISSN 1549-9596. doi:10.1021/acs.jcim.0c01329.

[14] Vytautas Gapsys, David F. Hahn, Gary Tresadern, David L. Mobley, Markus Rampp, and Bert L. de Groot. Pre-exascale computing of protein–ligand binding free energies with open source software for drug design. *Journal of Chemical Information and Modeling*, 62:1172–1177, 3 2022. ISSN 1549-9596. doi:10.1021/acs.jcim.1c01445.

[15] Abir Ganguly, Hsu-Chun Tsai, Mario Fernández-Pendás, Tai-Sung Lee, Timothy J. Giese, and Darrin M. York. Amber drug discovery boost tools: Automated workflow for production free-energy simulation setup and analysis (professa). *Journal of Chemical Information and Modeling*, 62:6069–6083, 12 2022. ISSN 1549-9596. doi:10.1021/acs.jcim.2c00879.

[16] 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, feb 2015. doi:10.1021/ja512751q. URL `http://dx.doi.org/10.1021/ja512751q`.

[17] Mariela Bollini, Robert A. Domaoal, Vinay V. Thakur, Ricardo Gallardo-Macias, Krasimir A. Spasov, Karen S. Anderson, and William L. Jorgensen. Computationally-guided optimization of a docking hit to yield catechol diethers as potent anti-HIV agents. *J. Med. Chem.*, 54(24):8582–8591, 2011. ISSN 00222623. doi:10.1021/jm201134m.

[18] Frank Lovering, Cristina Aevazelis, Jeanne Chang, Christoph Dehnhardt, Lori Fitz, Seungil Han, Kristin Janz, Julie Lee, Neelu Kaila, Joseph McDonald, William Moore, Alessandro Moretto, Nikolaos Papaioannou, David Richard, Mark S. Ryan, Zhao Kui Wan, and Atli Thorarensen. Imidazotriazines: Spleen Tyrosine Kinase (Syk) Inhibitors Identified by Free-Energy Perturbation (FEP). *ChemMedChem*, 11(2):217–233, 2016. ISSN 18607187. doi:10.1002/cmdc.201500333.

[19] Jérémie Mortier, Anders Friberg, Volker Badock, Dieter Moosmayer, Jens Schroeder, Patrick Steigemann, Franziska Siegel, Stefan Gradl, Marcus Bauser, Roman C. Hillig, Hans Briem, Knut Eis, Benjamin Bader, Duy Nguyen, and Clara D. Christ. Computationally Empowered Workflow Identifies Novel Covalent Allosteric Binders for KRASG12C. *ChemMedChem*, 15(10):827–832, 2020. ISSN 18607187. doi:10.1002/cmdc.201900727.

[20] Holly Freedman, Juthika Kundu, Egor Petrovitch Tchesnokov, John Lok Man Law, James A. Nieman, Raymond F. Schinazi, D. Lorne Tyrrell, Matthias Gotte, and Michael Houghton. Application of Molecular Dynamics Simulations to the Design of Nucleotide Inhibitors Binding to Norovirus Polymerase. *J. Chem. Inf. Model.*, 60(12): 6566–6578, 2020. ISSN 15205142. doi:10.1021/acs.jcim.0c00742.

[21] Gary Tresadern, Ingrid Velter, Andrés A. Trabanco, Frans Van Den Keybus, Gregor J. MacDonald, Marijke V.F. Somers, Greet Vanhoof, Philip M. Leonard, Marieke B.A.C. Lamers, Yves E.M. Van Roosbroeck, and Peter J.J.A. Buijnsters. Triazolo[1,5-a]pyrimidine Phosphodiesterase 2A Inhibitors: Structure and Free-Energy Perturbation-Guided Exploration. *J. Med. Chem.*, 63(21):12887–12910, 2020. ISSN 15204804. doi:10.1021/acs.jmedchem.0c01272.

[22] María Majellaro, Willem Jespers, Abel Crespo, María J. Núñez, Silvia Novio, Jhonny Azuaje, Rubén Prieto-Díaz, Claudia Gioé, Belma Alispahic, José Brea, María I. Loza, Manuel Freire-Garabal, Carlota Garcia-Santiago, Carlos Rodríguez-García, Xerardo García-Mera, Olga Caamaño, Christian Fernandez-Masaguer, Javier F. Sardina, Angela Stefanachi, Abdelaziz El Maatougui, Ana Mallo-Abreu, Johan Åqvist, Hugo Gutiérrez-De-Terán, and Eddy Sotelo. 3,4-Dihydropyrimidin-2(1 H)-ones as Antagonists of the Human A2BAdenosine Receptor: Optimization, Structure-Activity Relationship Studies, and Enantiospecific Recognition. *J. Med. Chem.*, 64(1):458–480, 2021. ISSN 15204804. doi:10.1021/acs.jmedchem.0c01431.

[23] Daniel H. O' Donovan, Clare Gregson, Martin J. Packer, Ryan Greenwood, Kurt G. Pike, Sameer Kawatkar, Andrew Bloecher, James Robinson, Jon Read, Erin Code, Jessie Hao Ru Hsu, Minhui Shen, Haley Woods, Peter Barton, Shaun Fillery, Beth Williamson, Philip B. Rawlins, and Sharan K. Bagal. Free energy perturbation in the design of EED ligands as inhibitors of polycomb repressive complex 2 (PRC2) methyltransferase. *Bioorganic Med. Chem. Lett.*, 39(February):127904, 2021. ISSN 14643405. doi:10.1016/j.bmcl.2021.127904. URL https://doi.org/10.1016/j.bmcl.2021.127904.

[24] Deyan Wu, Xuehua Zheng, Runduo Liu, Zhe Li, Zan Jiang, Qian Zhou, Yue Huang, Xu Nian Wu, Chen Zhang, Yi You Huang, and Hai Bin Luo. Free energy perturbation (FEP)-guided scaffold hopping. *Acta Pharm. Sin. B*, 12 (3):1351–1362, 2022. ISSN 22113843. doi:10.1016/j.apsb.2021.09.027. URL https://doi.org/10.1016/j.apsb.2021.09.027.

[25] Maryam Jama, Marawan Ahmed, Anna Jutla, Carson Wiethan, Jitendra Kumar, Tae Chul Moon, Frederick West, Michael Overduin, and Khaled H. Barakat. Discovery of allosteric shp2 inhibitors through ensemble-based consensus molecular docking, endpoint and absolute binding free energy calculations. *Computers in Biology and Medicine*, 152:106442, 1 2023. ISSN 00104825. doi:10.1016/j.compbiomed.2022.106442.

[26] Yudong Qiu, Daniel G.A. Smith, Simon Boothroyd, Hyesu Jang, David F. Hahn, Jeffrey Wagner, Caitlin C. Bannan, Trevor Gokey, Victoria T. Lim, Chaya D. Stern, Andrea Rizzi, Bryon Tjanaka, Gary Tresadern, Xavier Lucas, Michael R. Shirts, Michael K. Gilson, John D. Chodera, Christopher I. Bayly, David L. Mobley, and Lee Ping Wang. Development and Benchmarking of Open Force Field v1.0.0 - The Parsley Small-Molecule Force Field. *J. Chem. Theory Comput.*, 17(10):6262–6280, 2021. ISSN 15499626. doi:10.1021/acs.jctc.1c00571.

[27] Chao Lu, Chuanjie Wu, Delaram Ghoreishi, Wei Chen, Lingle Wang, Wolfgang Damm, Gregory A. Ross, Markus K. Dahlgren, Ellery Russell, Christopher D. Von Bargen, Robert Abel, Richard A. Friesner, and Edward D. Harder. OPLS4: Improving Force Field Accuracy on Challenging Regimes of Chemical Space. *J. Chem. Theory Comput.*, 2021. ISSN 1549-9618. doi:10.1021/acs.jctc.1c00302.

[28] Lingle Wang, Richard A Friesner, and Berne B. J.. Replica Exchange with Solute Scaling: A more efficient version of Replica Exchange with Solute Tempering (REST2). *J. Phys. Chem. B*, 115(30):9431–9438, 2011. ISSN 15378276. doi:10.1038/jid.2014.371. URL https://pubmed.ncbi.nlm.nih.gov/21714551/.

[29] Hannah E. Bruce Macdonald, Christopher Cave-Ayland, Gregory A. Ross, and Jonathan W. Essex. Ligand Binding Free Energies with Adaptive Water Networks: Two-Dimensional Grand Canonical Alchemical Perturbations. *J. Chem. Theory Comput.*, 14(12):6586–6597, 2018. ISSN 15499626. doi:10.1021/acs.jctc.8b00614.

[30] E. Prabhu Raman, Thomas J. Paul, Ryan L. Hayes, and Charles L. Brooks. Automated, Accurate, and Scalable Relative Protein-Ligand Binding Free-Energy Calculations Using Lambda Dynamics. *J. Chem. Theory Comput.*, 16(12):7895–7914, 2020. ISSN 15499626. doi:10.1021/acs.jctc.0c00830.

[31] Haoyu S. Yu, Yuqing Deng, Yujie Wu, Dan Sindhikara, Amy R. Rask, Takayuki Kimura, Robert Abel, and Lingle Wang. Accurate and Reliable Prediction of the Binding Affinities of Macrocycles to Their Protein Targets. *J. Chem. Theory Comput.*, 13(12):6290–6300, 2017. ISSN 15499626. doi:10.1021/acs.jctc.7b00885.

[32] Lingle Wang, Yuqing Deng, Yujie Wu, Byungchan Kim, David N Lebard, Dan Wandschneider, Mike Beachy, Richard A Friesner, and Robert Abel. Accurate Modeling of Scaffold Hopping Transformations in Drug Discovery. *J. Chem. Theory Comput.*, 13:42–54, 2017. doi:10.1021/acs.jctc.6b00991.

[33] Haoyu S. Yu, Cen Gao, Dmitry Lupyan, Yujie Wu, Takayuki Kimura, Chuanjie Wu, Leif Jacobson, Edward Harder, Robert Abel, and Lingle Wang. Toward Atomistic Modeling of Irreversible Covalent Inhibitor Binding Kinetics. *J. Chem. Inf. Model.*, 59(9):3955–3967, 2019. ISSN 15205142. doi:10.1021/acs.jcim.9b00268.

[34] Gregory A. Ross, Ellery Russell, Yuqing Deng, Chao Lu, Edward D. Harder, Robert Abel, and Lingle Wang. Enhancing Water Sampling in Free Energy Calculations with Grand Canonical Monte Carlo. *J. Chem. Theory Comput.*, 2020. ISSN 1549-9618. doi:10.1021/acs.jctc.0c00660.

[35] Zoe Cournia, Christophe Chipot, Benoît Roux, Darrin M. York, and Woody Sherman. Free energy methods in drug discovery—introduction. In Kira A. Amarcost and David C. Thompson, editors, *Free Energy Methods in Drug Discovery: Current State and Future Directions*, chapter 1, pages 39–66. ACS Publications, 2021.

[36] 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, 2016. ISSN 15205142. doi:10.1021/acs.jcim.6b00220.

[37] David Hahn, Christopher Bayly, Melissa L. Boby, Hannah Bruce Macdonald, John Chodera, Vytautas Gapsys, Antonia Mey, David Mobley, Laura Perez Benito, Christina Schindler, Gary Tresadern, and Gregory Warren. Best practices for constructing, preparing, and evaluating protein-ligand binding affinity benchmarks. *Living Journal of Computational Molecular Science*, (1):1497, Aug 2022. doi:https://doi.org/10.33011/livecoms.4.1.1497.

[38] Simon Boothroyd, Pavan Kumar Behara, Owen C Madin, David F Hahn, Hyesu Jang, Vytautas Gapsys, Jeffrey R Wagner, Joshua T Horton, David L Dotson, Matthew W Thompson, Jessica Maat, Trevor Gokey, Lee-Ping Wang, Daniel J Cole, Michael K Gilson, John D Chodera, Christopher I Bayly, Michael R Shirts, and David L Mobley. Development and benchmarking of open force field 2.0.0 - the sage small molecule force field. *ChemRxiv*, 2023. doi:10.26434/chemrxiv-2022-n2z1c-v2.

[39] Minyi Su, Qifan Yang, Yu Du, Guoqin Feng, Zhihai Liu, Yan Li, and Renxiao Wang. Comparative assessment of scoring functions: The casf-2016 update. *Journal of Chemical Information and Modeling*, 59:895–913, 2 2019. ISSN 1549960X. doi:10.1021/acs.jcim.8b00545.

[40] Peter J. Tonge. Quantifying the Interactions between Biomolecules: Guidelines for Assay Design and Data Analysis. *ACS Infect. Dis.*, 5(6):796–808, 2019. ISSN 23738227. doi:10.1021/acsinfecdis.9b00012.

[41] Cheng Yung-Chi and William H. Prusoff. Relationship between the inhibition constant (KI) and the concentration of inhibitor which causes 50 per cent inhibition (I50) of an enzymatic reaction. *Biochem. Pharmacol.*, 22(23): 3099–3108, 1973. ISSN 00062952. doi:10.1016/0006-2952(73)90196-2.

[42] Michelle J. Cannon, Giuseppe A. Papalia, Iva Navratilova, Robert J. Fisher, Lindsey R. Roberts, Karen M. Worthy, Andrew G. Stephen, Gerardo R. Marchesini, Edward J. Collins, Dave Casper, Huawei Qiu, Daulet Satpaev, Stefano F. Liparoto, Dax A. Rice, Inna I. Gorshkova, Ryan J. Darling, Donald B. Bennett, Michael Sekar, Eric Hommema, Amy M. Liang, Eric S. Day, Jean Inman, Shannon M. Karlicek, Stephen J. Ullrich, Dianne Hodges, Teresa Chu, Eric Sullivan, Jack Simpson, Ashique Rafique, Béatrice Luginbühl, Susanne Nyholm Westin, Magdalena Bynum, Paul Cachia, Yue Jin Li, Daniel Kao, Amy Neurauter, Melanie Wong, Michael Swanson, and David G. Myszka. Comparative analyses of a small molecule/enzyme interaction by multiple users of Biacore technology. *Anal. Biochem.*, 330(1):98–113, 2004. ISSN 00032697. doi:10.1016/j.ab.2004.02.027.

[43] Giuseppe A. Papalia, Stephanie Leavitt, Maggie A. Bynum, Phinikoula S. Katsamba, Rosemarie Wilton, Huawei Qiu, Mieke Steukers, Siming Wang, Lakshman Bindu, Sanjay Phogat, Anthony M. Giannetti, Thomas E. Ryan, Victoria A. Pudlak, Katarzyna Matusiewicz, Klaus M. Michelson, Agnes Nowakowski, Anh Pham-Baginski, Jonathan Brooks, Bryan C. Tieman, Barry D. Bruce, Michael Vaughn, Michael Baksh, Yun Hee Cho, Mieke De Wit, Alexandra Smets, Johan Vandersmissen, Lieve Michiels, and David G. Myszka. Comparative analysis of 10 small molecules binding to carbonic anhydrase II by different investigators using Biacore technology. *Anal. Biochem.*, 359(1):94–105, 2006. ISSN 10960309. doi:10.1016/j.ab.2006.08.021.

[44] Scott P. Brown, Steven W. Muchmore, and Philip J. Hajduk. Healthy skepticism: assessing realistic model performance. *Drug Discov. Today*, 14(7-8):420–427, 2009. ISSN 13596446. doi:10.1016/j.drudis.2009.01.012.

17

[45] D. G. Myszka, Y. N. Abdiche, F. Arisaka, O. Byron, E. Eisenstein, P. Hensley, J. A. Thomson, C. R. Lombardo, F. Schwarz, W. Stafford, and M. L. Doyle. The ABRF-MIRG'02 study: Assembly state, thermodynamic, and kinetic analysis of an enzyme/inhibitor interaction. *J. Biomol. Tech.*, 14(4):247–269, 2003. ISSN 15240215.

[46] Christian Kramer, Tuomo Kalliokoski, Peter Gedeck, and Anna Vulpetti. The experimental uncertainty of heterogeneous public K(i) data. *J. Med. Chem.*, 55(11):5165–5173, jun 2012. ISSN 1520-4804. doi:10.1021/jm300131x. URL http://dx.doi.org/10.1021/jm300131x.

[47] Joel Tellinghuisen and John D Chodera. Systematic errors in isothermal titration calorimetry: Concentrations and baselines. *Anal. Biochem.*, 414(2):297–299, jul 2011. ISSN 00032697. doi:10.1016/j.ab.2011.03.024. URL http://dx.doi.org/10.1016/j.ab.2011.03.024.

[48] Fouad H. Darras and Yuan Ping Pang. On the use of the experimentally determined enzyme inhibition constant as a measure of absolute binding affinity. *Biochem. Biophys. Res. Commun.*, 489(4):451–454, 2017. ISSN 10902104. doi:10.1016/j.bbrc.2017.05.168. URL http://dx.doi.org/10.1016/j.bbrc.2017.05.168.

[49] Vaida Paketurytė, Vaida Linkuvienė, Georg Krainer, Wen Yih Chen, and Daumantas Matulis. Repeatability, precision, and accuracy of the enthalpies and Gibbs energies of a protein–ligand binding reaction measured by isothermal titration calorimetry. *Eur. Biophys. J.*, 48(2):139–152, 2019. ISSN 14321017. doi:10.1007/s00249-018-1341-z. URL https://doi.org/10.1007/s00249-018-1341-z.

[50] Sarah E. Boyce, Joel Tellinghuisen, and John D. Chodera. Avoiding accuracy-limiting pitfalls in the study of protein-ligand interactions with isothermal titration calorimetry. *bioRxiv*, 2015. doi:10.1101/023796. URL https://www.biorxiv.org/content/early/2015/08/03/023796.

[51] Trung Hai Nguyen, Ariën S. Rustenburg, Stefan G. Krimmer, Hexi Zhang, John D. Clark, Paul A. Novick, Kim Branson, Vijay S. Pande, John D. Chodera, and David D.L. Minh. Bayesian analysis of isothermal titration calorimetry for binding thermodynamics. *PLoS One*, 13(9):1–26, 2018. ISSN 19326203. doi:10.1371/journal.pone.0203224.

[52] Tuomo Kalliokoski, Christian Kramer, Anna Vulpetti, and Peter Gedeck. Comparability of mixed IC50 data - a statistical analysis. *PLoS One*, 8(4):e61007+, apr 2013. ISSN 1932-6203. doi:10.1371/journal.pone.0061007. URL http://dx.doi.org/10.1371/journal.pone.0061007.

[53] Zaneta Nikolovska-Coleska, Renxiao Wang, Xueliang Fang, Hongguang Pan, York Tomita, Peng Li, Peter P. Roller, Krzysztof Krajewski, Naoyuki G. Saito, Jeanne A. Stuckey, and Shaomeng Wang. Development and optimization of a binding assay for the XIAP BIR3 domain using fluorescence polarization. *Anal. Biochem.*, 332 (2):261–273, 2004. ISSN 00032697. doi:10.1016/j.ab.2004.05.055.

[54] Wei Chen, Yuqing Deng, Ellery Russell, Yujie Wu, Robert Abel, and Lingle Wang. Accurate Calculation of Relative Binding Free Energies between Ligands with Different Net Charges. *J. Chem. Theory Comput.*, 14(12): 6346–6358, 2018. ISSN 15499626. doi:10.1021/acs.jctc.8b00825.

[55] Thomas B. Steinbrecher, Markus Dahlgren, Daniel Cappel, Teng Lin, Lingle Wang, Goran Krilov, Robert Abel, Richard Friesner, and Woody Sherman. Accurate Binding Free Energy Predictions in Fragment Optimization. *J. Chem. Inf. Model.*, 55(11):2411–2420, 2015. ISSN 15205142. doi:10.1021/acs.jcim.5b00538.

[56] 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, Ariele Viacava Follis, 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, 2020. ISSN 15205142. doi:10.1021/acs.jcim.0c00900.

[57] Francesco Deflorian, Laura Pérez-Benito, Eelke B. Lenselink, Miles Congreve, Herman Van Vlijmen, Jonathan S. Mason, Chris de Graaf, and Gary Tresadern. Accurate Prediction of GPCR Ligand Binding Affinity with Free Energy Perturbation. *J. Chem. Inf. Model.*, 2020. ISSN 1549-9596. doi:10.1021/acs.jcim.0c00449.

[58] Callum J Dickson, Viktor Hornak, and Jose S Duca. Relative Binding Free-Energy Calculations at Lipid-Exposed Sites: Deciphering Hot Spots. *J. Chem. Inf. Model.*, 61(12):5923–5930, 2021. ISSN 15205142. doi:10.1021/acs.jcim.1c01147.

[59] Vincent Wagner, Linda Jantz, Hans Briem, Kai Sommer, Matthias Rarey, and Clara D. Christ. Computational Macrocyclization: From de novo Macrocycle Generation to Binding Affinity Estimation. *ChemMedChem*, 12(22): 1866–1872, 2017. ISSN 18607187. doi:10.1002/cmdc.201700478.

[60] 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, 2017. ISSN 15499626. doi:10.1021/acs.jctc.6b01141.

[61] Daniel Cappel, Steven Jerome, Gerhard Hessler, and Hans Matter. Impact of Different Automated Binding Pose Generation Approaches on Relative Binding Free Energy Simulations. *J. Chem. Inf. Model.*, 60(3):1432–1444, 2020. ISSN 15205142. doi:10.1021/acs.jcim.9b01118.

[62] Michael F.T. Koehler, Philippe Bergeron, Elizabeth M. Blackwood, Krista Bowman, Kevin R. Clark, Ron Firestein, James R. Kiefer, Klaus Maskos, Mark L. McCleland, Linda Orren, Laurent Salphati, Steve Schmidt, Elisabeth V. Schneider, Jiansheng Wu, and Maureen H. Beresini. Development of a potent, specific cdk8 kinase inhibitor which phenocopies cdk8/19 knockout cells. *ACS Medicinal Chemistry Letters*, 7:223–228, 1 2016. ISSN 19485875. doi:10.1021/acsmedchemlett.5b00278.

[63] Kristoffer Peterson. *Molecular basis for galectin-ligand interactions. Design , synthesis and analysis*. PhD thesis, Lund University, 2018.

[64] Christopher R. Smith, Douglas R. Dougan, Mallareddy Komandla, Toufike Kanouni, Beverly Knight, J. David Lawson, Mark Sabat, Ewan R. Taylor, Phong Vu, and Corey Wyrick. Fragment-Based Discovery of a Small Molecule Inhibitor of Bruton's Tyrosine Kinase. *J. Med. Chem.*, 58(14):5437–5444, 2015. ISSN 15204804. doi:10.1021/acs.jmedchem.5b00734.

[65] Jenny K Ekegren, Nina Ginman, Åsa Johansson, Hans Wallberg, Mats Larhed, Bertil Samuelsson, Torsten Unge, and Anders Hallberg. Microwave-Accelerated Synthesis of P1'-Extended HIV-1 Protease Inhibitors Encompassing a Tertiary Alcohol in the Transition-State Mimicking Scaffold. *J. Med. Chem.*, 49:1828–1832, 2006. doi:10.1021/jm051239z CCC:.

[66] Arjun Saha, Amy Y. Shih, Taraneh Mirzadegan, and Mark Seierstad. Predicting the Binding of Fatty Acid Amide Hydrolase Inhibitors by Free Energy Perturbation. *J. Chem. Theory Comput.*, 14(11):5815–5822, 2018. ISSN 15499626. doi:10.1021/acs.jctc.8b00672.

[67] César De Oliveira, Haoyu S. Yu, Wei Chen, Robert Abel, and Lingle Wang. Rigorous Free Energy Perturbation Approach to Estimating Relative Binding Affinities between Ligands with Multiple Protonation and Tautomeric States. *J. Chem. Theory Comput.*, 15(1):424–435, 2019. ISSN 15499626. doi:10.1021/acs.jctc.8b00826.

[68] G. J. Tuckerman, M; Berne, B. J.; Martyna. Reversible multiple time scale molecular dynamics. *J. Phys. Chem.*, 97(3):1990–2001, 1992. ISSN 00223654. doi:10.1021/j100153a002.

[69] Glenn J Martyna, Michael L Klein, and Mark Tuckermana. Nosé-Hoover chains: The canonical ensemble via continuous dynamics. *J. Chem. Phys.*, 97(4):2635–2643, 1992.

[70] Glenn J. Martyna, Douglas J. Tobias, and Michael L. Klein. Constant pressure molecular dynamics algorithms. *The Journal of Chemical Physics*, 101:4177–4189, 9 1994. ISSN 0021-9606. doi:10.1063/1.467468.

[71] Chresten R. Søndergaard, Mats H.M. Olsson, Michał Rostkowski, and Jan H. Jensen. Improved treatment of ligands and coupling effects in empirical calculation and rationalization of p K a values. *J. Chem. Theory Comput.*, 7(7):2284–2295, 2011. ISSN 15499618. doi:10.1021/ct200133y.

[72] Lingle Wang, Yuqing Deng, Jennifer L. Knight, Yujie Wu, Byungchan Kim, Woody Sherman, John C. Shelley, Teng Lin, and Robert Abel. Modeling local structural rearrangements using FEP/REST: Application to relative binding affinity predictions of CDK2 inhibitors. *J. Chem. Theory Comput.*, 9(2):1282–1293, 2013. ISSN 15499618. doi:10.1021/ct300911a.

[73] Huafeng Xu. Optimal measurement network of pairwise differences. *J. Chem. Inf. Model.*, 59:4720–4728, 2019. ISSN 23318422.

[74] Qingyi Yang, Woodrow Burchett, Gregory S. Steeno, Shuai Liu, Mingjun Yang, David L. Mobley, and Xinjun Hou. Optimal designs for pairwise calculation: An application to free energy perturbation in minimizing prediction variability. *J. Comput. Chem.*, 41(3):247–257, 2020. ISSN 1096987X. doi:10.1002/jcc.26095.

[75] Gregory A. Ross and Lingle Wang. The FEP+ benchmark and experimental reproducibility survey. `https://github.com/schrodinger/public_binding_free_energy_benchmark`, 2023.

[76] Lorenzo D'Amore and David Hahn. Follow-up workshop on benchmarking, September 2021. URL `https://doi.org/10.5281/zenodo.5369858`.

[77] Janet L. Paulsen, Haoyu S. Yu, Dan Sindhikara, Lingle Wang, Todd Appleby, Armando G. Villaseñor, Uli Schmitz, and Devleena Shivakumar. Evaluation of Free Energy Calculations for the Prioritization of Macrocycle Synthesis. *J. Chem. Inf. Model.*, 60(7):3489–3498, 2020. ISSN 15205142. doi:10.1021/acs.jcim.0c00132.

[78] The COVID Moonshot Consortium, Hagit Achdout, Anthony Aimon, Elad Bar-David, Haim Barr, Amir Ben-Shmuel, James Bennett, Vitaliy A. Bilenko, Vitaliy A. Bilenko, Melissa L. Boby, Bruce Borden, Gregory R. Bowman, Juliane Brun, Sarma BVNBS, Mark Calmiano, Anna Carbery, Daniel Carney, Emma Cattermole, Edcon Chang, Eugene Chernyshenko, John D. Chodera, Austin Clyde, Joseph E. Coffland, Galit Cohen, Jason Cole, Alessandro Contini, Lisa Cox, Milan Cvitkovic, Alex Dias, Kim Donckers, David L. Dotson, Alice Douangamath, Shirly Duberstein, Tim Dudgeon, Louise Dunnett, Peter K. Eastman, Noam Erez, Charles J. Eyermann, Mike Fairhead, Gwen Fate, Daren Fearon, Oleg Fedorov, Matteo Ferla, Rafaela S. Fernandes, Lori Ferrins, Richard Foster, Holly Foster, Ronen Gabizon, Adolfo Garcia-Sastre, Victor O. Gawriljuk, Paul Gehrtz, Carina Gileadi, Charline Giroud, William G. Glass, Robert Glen, Itai Glinert, Andre S. Godoy, Marian Gorichko, Tyler Gorrie-Stone, Ed J. Griffen, Storm Hassell Hart, Jag Heer, Michael Henry, Michelle Hill, Sam Horrell, Victor D. Huliak, Matthew F.D. Hurley, Tomer Israely, Andrew Jajack, Jitske Jansen, Eric Jnoff, Dirk Jochmans, Tobias John, Steven De Jonghe, Anastassia L. Kantsadi, Peter W. Kenny, J. L. Kiappes, Serhii O. Kinakh, Lizbe Koekemoer, Boris Kovar, Tobias Krojer, Alpha Lee, Bruce A. Lefker, Haim Levy, Ivan G. Logvinenko, Nir London, Petra Lukacik, Hannah Bruce Macdonald, Beth MacLean, Tika R. Malla, Tatiana Matviiuk, Willam McCorkindale, Briana L. McGovern, Sharon Melamed, Kostiantyn P. Melnykov, Oleg Michurin, Halina Mikolajek, Bruce F. Milne, Aaron Morris, Garrett M. Morris, Melody Jane Morwitzer, Demetri Moustakas, Aline M. Nakamura, Jose Brandao Neto, Johan Neyts, Luong Nguyen, Gabriela D. Noske, Vladas Oleinikovas, Glaucius Oliva, Gijs J. Overheul, David Owen, Ruby Pai, Jin Pan, Nir Paran, Benjamin Perry, Maneesh Pingle, Jakir Pinjari, Boaz Politi, Ailsa Powell, Vladimir Psenak, Reut Puni, Victor L. Rangel, Rambabu N. Reddi, St Patrick Reid, Efrat Resnick, Emily Grace Ripka, Matthew C. Robinson, Ralph P. Robinson, Jaime Rodriguez-Guerra, Romel Rosales, Dominic Rufa, Kadi Saar, Kumar Singh Saikatendu, Chris Schofield, Mikhail Shafeev, Aarif Shaikh, Jiye Shi, Khriesto Shurrush, Sukrit Singh, Assa Sittner, Rachael Skyner, Adam Smalley, Bart Smeets, Mihaela D. Smilova, Leonardo J. Solmesky, John Spencer, Claire Strain-Damerell, Vishwanath Swamy, Hadas Tamir, Rachael Tennant, Warren Thompson, Andrew Thompson, Susana Tomasio, Igor S. Tsurupa, Anthony Tumber, Ioannis Vakonakis, Ronald P. van Rij, Laura Vangeel, Finny S. Varghese, Mariana Vaschetto, Einat B. Vitner, Vincent Voelz, Andrea Volkamer, Frank von Delft, Annette von Delft, Martin Walsh, Walter Ward, Charlie Weatherall, Shay Weiss, Kris M. White, Conor Francis Wild, Matthew Wittmann, Nathan Wright, Yfat Yahalom-Ronen, Daniel Zaidmann, Hadeer Zidane, and Nicole Zitzmann. Open science discovery of oral non-covalent sars-cov-2 main protease inhibitor therapeutics. *bioRxiv*, 2022. doi:10.1101/2020.10.29.339317. URL https://www.biorxiv.org/content/early/2022/01/30/2020.10.29.339317.

[79] Steven Dajnowicz, Delaram Ghoreishi, Kalyan Modugula, Wolfgang Damm, Edward D. Harder, Robert Abel, Lingle Wang, and Haoyu S. Yu. Advancing Free-Energy Calculations of Metalloenzymes in Drug Discovery via Implementation of LFMM Potentials. *J. Chem. Theory Comput.*, 16(11):6926–6937, 2020. ISSN 15499626. doi:10.1021/acs.jctc.0c00615.

[80] Sayan Mondal, Gary Tresadern, Jeremy Greenwood, Byungchan Kim, Joe Kaus, Thomas Steinbrecher, Lingle Wang, Craig Masse, Ramy Farid, and Robert Abel. A free energy perturbation approach to estimate the intrinsic solubilities of drug-like small molecules. *ChemRxiv*, 2019. doi:10.26434/chemrxiv.10263077.v1.

[81] Mehtap Işık, Ariën S. Rustenburg, Andrea Rizzi, M. R. Gunner, David L. Mobley, and John D. Chodera. *Overview of the SAMPL6 pK a challenge: evaluating small molecule microscopic and macroscopic pK a predictions*, volume 35. Springer International Publishing, 2021. ISBN 0123456789. doi:10.1007/s10822-020-00362-6. URL https://doi.org/10.1007/s10822-020-00362-6.

[82] Kevin Hauser, Christopher Negron, Steven K. Albanese, Soumya Ray, Thomas Steinbrecher, Robert Abel, John D. Chodera, and Lingle Wang. Predicting resistance of clinical Abl mutations to targeted kinase inhibitors using alchemical free-energy calculations. *Commun. Biol.*, 1(1):1–14, 2018. ISSN 23993642. doi:10.1038/s42003-018-0075-x. URL http://dx.doi.org/10.1038/s42003-018-0075-x.

[83] Dana Katz, Michael A. Dimattia, Dan Sindhikara, Hubert Li, Nikita Abraham, and Abba E. Leffler. Potency-and selectivity-enhancing mutations of conotoxins for nicotinic acetylcholine receptors can be predicted using accurate free-energy calculations. *Mar. Drugs*, 19(7), 2021. ISSN 16603397. doi:10.3390/md19070367.

[84] Joseph W. Kaus, Edward Harder, Teng Lin, Robert Abel, J. Andrew McCammon, and Lingle Wang. How to deal with multiple binding poses in alchemical relative protein-ligand binding free energy calculations. *J. Chem. Theory Comput.*, 11(6):2670–2679, 2015. ISSN 15499626. doi:10.1021/acs.jctc.5b00214.

[85] David L Mobley, John D Chodera, and Ken A Dill. On the use of orientational restraints and symmetry corrections in alchemical free energy calculations. *J. Chem. Phys.*, 125(8), 2006. ISSN 00219606. doi:10.1063/1.2221683.