Middle-way flexible docking: Pose prediction using mixed-resolution Monte Carlo in estrogen receptor α

Justin Spiriti[1], Sundar Raman Subramanian[2], Rohith Palli[3], Maria Wu[4], Daniel M. Zuckerman[1*]

1 Department of Biomedical Engineering, Oregon Health and Science University, Portland, OR 97239
2 Address unknown
3 Medical Scientist Training Program and Biophysics, Structural & Computational Biology Program, University of Rochester, 601 Elmwood Avenue, Box 350, Rochester, NY 14642
4 1521 Bloomin Spring Ct, Hebron KY 41048

* zuckermd@ohsu.edu

Abstract

There is a vast gulf between the two primary strategies for simulating protein-ligand interactions. Docking methods significantly limit or eliminate protein flexibility to gain great speed at the price of uncontrolled inaccuracy, whereas fully flexible atomistic molecular dynamics simulations are expensive and often suffer from limited sampling. We have developed a flexible docking approach geared especially for highly flexible or poorly resolved targets based on mixed-resolution Monte Carlo (MRMC), which is intended to offer a balance among speed, protein flexibility, and sampling power. The binding region of the protein is treated with a standard atomistic force field, while the remainder of the protein is modeled at the residue level with a Gō model that permits protein flexibility while saving computational cost. Implicit solvation is used. Here we assess three facets of the MRMC approach with implications for other docking studies: (i) the role of receptor flexibility in cross-docking pose prediction; (ii) the use of non-equilibrium candidate Monte Carlo (NCMC) and (iii) the use of pose-clustering in scoring. We examine 61 co-crystallized ligands of estrogen receptor α, an important cancer target known for its flexibility. Adding protein flexibility, not surprisingly, leads to significantly lower total energies and stronger interactions between protein and ligand, but notably we document the important role of backbone flexibility in the improvement. Somewhat unexpectedly, our implementation of NCMC leads to only modestly improved sampling of ligand poses. Overall, the addition of protein flexibility improves the performance of docking, as measured by energy-ranked poses, but we do not find significant improvements based on cluster information or the use of NCMC.

Introduction

Computational structure-based drug design can play an important role in drug development, as exemplified in the development of inhibitors of HIV protease, which have a major impact on treatment for AIDS patients. [1, 2] Because of their potential to reduce the cost and time associated with drug development, a multitude of methods have been developed to screen potential drug candidates virtually and prioritize

possible structures for synthesis. [3–6] Some of the most popular docking methods represent the potential energy due to the receptor using grids and optimize the ligand conformation with respect to this potential. These include Autodock Vina, [7] and the related smina [8], Schrodinger Glide, [9–11] CDOCKER, [12] and others. [13] Once a grid is constructed, however, the protein conformation represented by that grid is fixed, which is a serious approximation, since a number of structural studies have shown that "hidden" protein conformations and protein flexibility play important roles in protein-ligand binding; methods that consider protein flexibility and multiple structures improve docking performance compared to those that make use of only one structure. [14–21]

A number of approaches have been developed to incorporate flexibility into docking. [22] One approach is to try to incorporate protein flexibility into grid-based approaches by collecting multiple confomations of the receptor and docking to all of them, a practice commonly known as ensemble docking. [14, 15, 18, 21, 23, 24] Another grid-based strategy involves leaving certain amino acid side chains out of the grid and optimizing their conformation alongside that of the ligand during the docking procedure. CDOCKER has been modified in this way for example, [25] and Autodock Vina and smina are also capable of this. [7, 8] Other examples of grid-based force fields, some of which are aimed at the related protein-protein docking problem, include the ATTRACT force field [26, 27] and LightDock. [13] However, the amount of flexibility that can be incorporated by both these methods is limited. In particular, allowing only a few side chains to be flexible means these few side chains must be carefully chosen and that important protein motions involving the backbone are not represented. Likewise, ensemble docking requires a careful choice of conformations to be used and only allows for a limited degree of backbone flexibility. Since these conformations come from simulations or structures of the protein without the corresponding ligand, ensemble docking cannot take into account the mutual induced fit that may occur when a ligand binds to a protein.

The RosettaLigand approach to ligand docking [28–30] makes use of the Rosetta knowledge-based force field and, in principle, allows for full receptor flexibility. Like other knowledge-based force fields, it relies on the assumption that the system under discussion is similar to known protein-ligand complexes. Furthermore, the RosettaLigand docking approach used a complex protocol involving several rounds of minimization, which loses information about the relative entropy of the energy minima that are found. Although in theory the protocol should be able to allow full receptor flexibility, in practice it was found that restraints on the α carbons were needed to improve the discrimination between native and nonnative poses. RosettaLigand is also relatively expensive computationally, requiring approximately 80 CPU hours per ligand. [28]

Another way in which full flexibility of the protein can be allowed, at even greater computational cost, is by using molecular dynamics or Monte Carlo simulations. Alchemical free energy methods have received much attention recently. [31–33] These methods require conducting multiple simulations at different values of a coupling parameter λ that serves to scale protein-ligand interactions. In principle, they are exact according to the laws of statistical mechanics and therefore should give perfectly accurate results given an accurate force field potential and a simulation of infinite length. In practice, however, both force field errors and inadequate sampling can result in significant errors in the calculated free energies. [34, 35] Alternatively, methods such as MM-PBSA [36, 37] can be used; these are somewhat less expensive because they only require simulation of the endpoints, but also make additional approximations. There is therefore a need for an *in silico* docking technique that incorporates full flexibility of the protein at modest computational cost.

Multiscale simulation techniques show promise for reducing computational time while maintaining full flexibility of the protein and physical accuracy. They seem especially appropriate for the docking problem because the regions of the protein far from the ligand can be simulated with a less physically accurate, but less computationally expensive, coarse-grained model. Several efforts have been made to implement this idea in simulations, although not always in the context of protein-ligand simulations. An early approach that bears some similarities our own involves dividing the system into three regions: an atomistic region, a coarse-grained region using a Gō-like potential, and an intermediate region. [38] Our group has also done some preliminary work on mixed resolution models, combining a residuel level Gō model with the OPLS-AA force field and conducting some preliminary tests on self-docking to the estrogen receptor. [39] The popular MARTINI force field for water has also been combined with an atomistic force field for proteins; [40] however, balancing the strength of electrostatic forces in the two force fields proved difficult. Others have tried to combine an atomistic region with an elastic network model. [41]. Feig and co-workers have also combined their PRIMO force field with the CHARMM36 atomistic force field, [42] and obtained similar results to fully atomistic or fully coarse-grained simulations, but they note issues with weakened hydrophobic packing interactions, and the amount of speedup they obtained is generally modest because PRIMO is relatively detailed for a coarse-grained model. A few workers have also attempted to adapt grid-based docking methods to work in a multiscale manner; some of these approaches are aimed at the related protein-protein docking problem. [13, 26, 27]

Here we extend our previous mixed-resolution Monte Carlo software [39] and use it to systematically study key aspects of docking in a challenging system. In the software, the majority of the protein is modeled using a residue-level coarse-grained model currently based on the Gō model [43–46] while the ligand and binding site are modeled using an atomistic force field. Interactions between the two regions are treated in a fully atomistic manner. Full flexibility of the ligand and receptor is maintained, allowing the modeling of mutual induced fit, including "breathing" motions away from the binding site. [47, 48] Importantly, the coordinates of all atoms are tracked throughout the simulation, which increases computational cost but enables maintaining proper backbone geometry via standard non-bonded terms even in the coarse region. The computational cost of the method can be adjusted by changing the size of the atomistic region; for the region used here, the mixed-resolution model reduces the computational cost by a factor of 2-4 compared to a fully atomistic treatment of the protein using the same implementation. In addition, coarse-grained models produce potential energy surfaces that are smoother than corresponding atomistic surfaces, because of averaging over omitted degrees of freedom. [49] While we use a Gō-like model for the coarse-grained region in the protein in this work, the mixed-resolution concept can be extended to incorporate other coarse-grained models. The use of Monte Carlo simulation with full flexibility in principle also allows for the inclusion of entropy effects.

In this work, we have replaced the discontinous Gō model functional form used in previous work [45, 46] with a continous Lennard-Jones functional form. We have also replaced the OPLS-AA force field used in our previous work with the AMBER 99SB force field [50] so that ligands can be automatically parameterized using the compatible GAFF forcefield and antechamber tool. [51, 52] To save computer time, we have replaced the Generalized Born solvent model with a simpler solvent exposure-dependent distance-dependent dielectric model. [53] We no longer use precalculated libraries of amino acid conformations in our Monte Carlo moves—a technique worked well for peptides but proved to be less advantageous for Monte Carlo

simulations of dense protein systems because of poor acceptance rates.

Monte Carlo simulation readily allows for the systematic testing of the role of flexibilty in docking, as has been noted. [28, 29] Flexible degrees of freedom and interaction type (coarse-grained vs. all-atom) are readily adjusted. Here, we specifically examine the effects of common choices in conventional docking: rigid side chains and rigid backbones, both of which prove detrimental. We are unaware of a prior study examining these cases together.

In addition to the mixed-resolution model, we examine the recently developed nonequilibrium candidate Monte Carlo (NCMC) method to enhance the sampling of ligand binding modes. [54] In this method, the potential energy is perturbed systematically in a short nonequilibrium simulation over the course of 10^2-10^3 MC trial moves, and the entire sequence of moves is then either accepted or rejected based on the nonequilibrium work done during the simulation. In principle, the perturbation of potential energy can be designed to allow large configurational changes that would have a very low acceptance probability in a standard MC simulation, while the short nonequilibrium simulations allow time for the system to relax after each such configurational change. Gill et al. have applied this method to the binding of toluene to the L99A mutant of T4 lysozyme and found that the NCMC method produced a much higher rate of transitions between distinct binding modes of toluene compared to standard MD simulations or simulations using MD combined with MC. [55]

We examine the effects of flexibility and the efficacy of NCMC in the context of docking known ligands to the ligand binding domain of estrogen receptor α (ER α). ER α is known to undergo a conformational change upon binding by estradiol and other ER agonists in which helix 12 closes over the binding site as shown in Fig 1. [56, 57] In contrast, in the bound structures with tamoxifen and other ER α antagonists, helix 12 does not make this conformational change. [58] The development of drugs to modulate ER activity is of considerable interest because aberrant ER signaling has long been known to be a key player in promoting proliferation of several types of cancer, and multiple ER modulating drugs are currently in clinical use. [59–61] However ER has also been recently observed to evolve drug-resistance mutations during metastatic progression in breast cancer, which limits the ability of current therapeutic agents to affect the progression of secondary tumors. [62–64] We have previously validated a similar MRMC approach for ER in a very limited way using a simple self-docking test. [39] Here we attack the cross-docking problem, in which the co-crystal structure for the tested compound is not used.

The first section of the paper describes the mixed resolution potential and the Monte Carlo and NCMC protocols that were used with it for docking. For this study, ligands with known bound crystal structures were used, so we can compare the docked conformations to experimental crystal structures. We study the docking protocol with and without NCMC and with different levels of protein flexibility, and also tested different methods of ranking the poses, including trying to cluster the poses by structural similarity. We find that having full protein flexibility (including backbone flexibility) results in stronger interactions between protein and ligand, and top-ranked poses that are closer to the corresponding crystal structures. Study of the acceptance rates of large ligand moves and of variations in ligand RMSD during each simulation provide some evidence that NCMC is improving sampling. However, the use of NCMC does not improve the docking results beyond what is obtained with simulations with a fully flexible protein without NCMC. Also, the use of clustering as a part of pose ranking does not appear to improve the overall docking results. Finally, we discuss possible improvements to the mixed resolution potential and to sampling, and the ability of our protocol to sample multiple docked poses for each ligand.

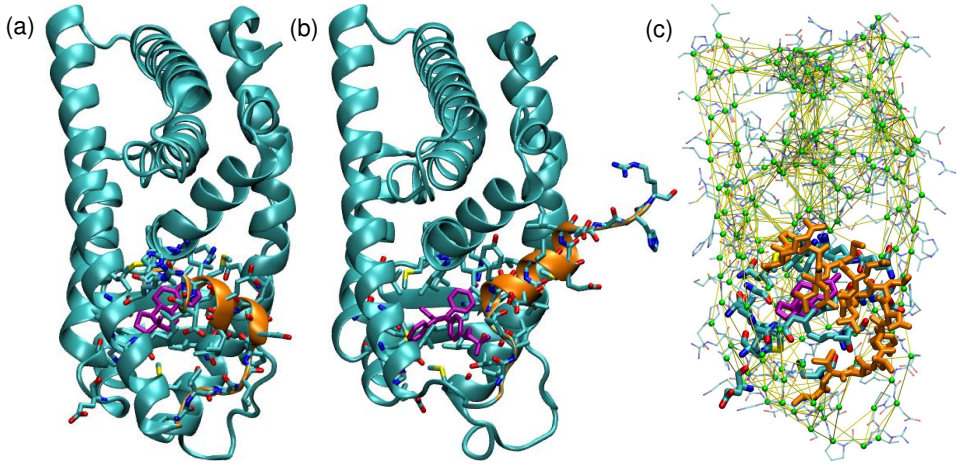

Fig 1. Structure of ER α.
(a) Active and (b) inactive conformation of ER α. (c) Illustration of the mixed-resolution model used in this paper. The ligand (green) and the binding site (heavy structure) constitute the atomistic region and are treated using an all-atom force field. The remainder of the protein is treated as the coarse-grained region is represented by particles located at each α carbon (purple spheres) with native attractions between them (yellow lines). Rigid structures of each amino acid (thin structure) and moved along with the coarse-grained particles and used to calculate interactions between the coarse-grained and atomistic regions. Helix 12 is indicated in orange in all three panels.

Materials and methods

Mixed-Resolution Potential

Simulations were done with a mixed-resolution potential, in which the majority of the protein is treated with a Gō model [43–46, 65] and the atomistic region around the binding site is treated using the AMBER 99SB force field. [50] The overall potential has the following form:

$$U = U_{\mathrm{CG}} + U_{\mathrm{AA}} + U_{\mathrm{CG/AA}} \tag{1}$$

The coarse-grained portion of the potential U_{CG} is in turn given by

$$U_{\mathrm{CG}} = U_{\mathrm{bb}} + U_{\mathrm{G\bar{o}}} \tag{2}$$

$$U_{\mathrm{G\bar{o}}} = \sum_{i,j} \begin{cases} \varepsilon \left[5\left(\frac{r^0_{ij}}{r_{ij}}\right)^{12} - 6\left(\frac{r^0_{ij}}{r_{ij}}\right)^{10} \right] & \text{(native contacts)} \\ 5\varepsilon \left(\frac{r_{\mathrm{HC}}}{r_{ij}}\right)^{12} & \text{(nonnative contacts)} \end{cases} \tag{3}$$

where the sum is taken over all pairs of residues in which both residues belong to the coarse-grained region. The coarse-grained potential also includes a backbone component, which uses the bond, angle, and dihedral terms from the AMBER 99SB force field for the backbone atoms in the coarse-grained region (1-4 van der Waals terms are not included). The 12-10 form used here replaces a square well potential used in previous work [45, 46] and has shown good performance in other settings. [65–67] The Gō model well depth ε and hard core radius r_{HC} are listed in table 1. r^0_{ij} is the native distance between atoms i and j, determined from the original structure.

The all-atom potential U_{AA} is the potential energy of the all-atom region, according to the AMBER 99SB force field. [50] The $U_{\mathrm{CG/AA}}$ term represents the interaction between the two regions, which is also computed atomistically using the AMBER 99SB force field. The all-atom region was defined to contain the ligand and all residues with at least one non-hydrogen atom within 3 Å of the ligand in any of the 66 crystal structures of ligand-ER complexes that were used as references; in addition, residues 533-548, which comprise helix 12 and the neighboring loop, were also included in the all-atom region. Ligands were parameterized with the antechamber tool [52] using the GAFF force field. [51]

In order to incorporate solvation effects in a computationally efficient manner, the electrostatic term of the AMBER 99SB was modified to use a solvent-exposure dependent distance dependent dielectric originally developed by Garden and Zhorov, which was previously found to work well in docking simulations with an AMBER force field. [53] In this model, the electrostatic interaction between atoms i and j is given by $K_{\mathrm{coul}} q_i q_j / \varepsilon_{ij} r_{ij}$ where

$$\varepsilon_{ij} = r_{ij}[\varepsilon_0 + (1 - s_{kl})(\varepsilon_1 - \varepsilon_0)] \quad (4)$$

$$s_{kl} = c(v_k + v_l) \quad (5)$$

where ε_0 and ε_1 are low and high dielectric constants, respectively, and s_{kl} is the overlap of hydration shell volumes for the groups k and l that include atoms i and j. (For all pairs of groups k and l, $0 \leq s_{kl} < 1$.) The hydration volumes v_k and v_l for these groups are calculated using a formula originally used for the EEF1 implicit solvent model. [68] The values for ε_0, ε_1, and c shown in table 1 are those found to optimize docking results for a training set in ref. 53.

Table 1. Parameters of the simulation.

Symbol	Description	Value
ε	Gō model well depth	3.0 kcal/mol
r_{HC}	Hard core radius	1.7 Å
	Cutoff for defining native interactions in Gō model	8 Å
ε_0	Low-dielectric constant	2
ε_1	High dielectric constant	8
c	Constant for determining hydration overlaps	0.625
	Cutoff for atomistic nonbonded interactions	10 Å
T	MC simulation temperature	300 K
	Method of selecting atomistic region	residues 533-548 plus all residues with at least one non-hydrogen atom within 3 Å of the ligand in any of 66 reference structures
	Number of residues in atomistic region	39

Standard Monte Carlo

Sampling was performed using MC simulation employing a standard variety of local and global moves as described in table 2. To investigate the impact of protein flexiblity, in addition to simulations in which the protein was fully flexible, simulations were also conducted in which the protein was kept rigid or only the sidechains were allowed to move. This was done by leaving out the corresponding MC moves and increasing the fraction of others in the move mix, also as shown in table 2. In all simulations, both the translational and rotational degrees of freedom of the ligand with respect to the protein and its internal degrees of freedom were sampled.

Table 2. Monte Carlo moves used in the simulation.

Type	Description	Max. size	Fraction		
			fully flexible	fixed protein	sidechain only
Backbone rotation	Rotation about a randomly selected rotatable bond in protein backbone	2°	0.1	0	0
Sidechain rotation	Rotation about a randomly selected rotatable bond in amino acid side chain (atomistic region only)	180°	0.2	0	0.25
Backrub rotation [69]	Rotation of the part of the protein between two randomly selected amino acids about the axis joining their alpha-carbons	2°	0.1	0	0
Ligand bond rotation	Rotation about a randomly selected rotatable bond in the ligand	180°	0.2	0.33	0.25
Ligand translation	Translation of the ligand by a randomly selected vector	1 Å	0.2	0.33	0.25
Ligand rotation	Random rotation of the ligand about its center of mass	180°	0.2	0.33	0.25

Nonequilibrium candidate Monte Carlo

In order to enhance the sampling of ligand poses, docking runs were also undertaken using a modified version of the nonequilibrium candidate Monte Carlo (NCMC) algorithm. [54, 55] In this method, the system is subjected to NCMC "moves" which are in fact short nonequilibrium simulations (using standard MC) during which the potential energy function is occasionally perturbed such that sampling for particular degrees of freedom may be enhanced.

In the method used here, the potential energy is modified by scaling the van der Waals and electrostatic components of the protein-ligand interaction energy by the factors λ_{VDW} and λ_{elec}:

$$U = U_{\mathrm{protein}} + U_{\mathrm{ligand}} + \lambda_{\mathrm{VDW}} U_{\mathrm{VDW}} + \lambda_{\mathrm{elec}} U_{\mathrm{elec}} \tag{6}$$

Each NCMC move consisted of 800 individual Monte Carlo moves, which were divided into four 100-move phases in which the coupling parameters λ_{VDW} and λ_{elec} were changed according to the schedule shown in Fig. 2. First, the charges on the ligand were removed by systematically driving λ_{elec} towards 0. Second, λ_{VDW} was also driven towards 0, such that the ligand was completely uncoupled from the protein and free to rotate, translate, or change conformation without interference. Then, in the second two phases, first λ_{VDW} and then λ_{elec} weree gradually transitioned back to 1, so that the system relaxed and steric clashes might be resolved. In all four phases, λ_{VDW} and λ_{elec} were changed by 0.05 every 5 trial moves. During all four phases, the individual trial moves were preliminarily accepted or rejected with the standard Metropolis criterion,

$$p_{\mathrm{preliminary}} = \min[1, \exp -\beta(U_{\mathrm{new}} - U_{\mathrm{old}})] \tag{7}$$

where U is the scaled potential given above. This leads to canonical sampling suitable for the given set of λ values.

At the end of each NCMC move, the nonequilibrium work w performed on the system (which accounts for the changes in λ values) was calculated using

$$w = \sum \left[\left(\lambda_{\mathrm{VDW}}^{(i+1)} - \lambda_{\mathrm{VDW}}^{(i)} \right) U_{\mathrm{VDW}}(\mathbf{r}^{(i)}) + \left(\lambda_{\mathrm{elec}}^{(i+1)} - \lambda_{\mathrm{elec}}^{(i)} \right) U_{\mathrm{elec}}(\mathbf{r}^{(i)}) \right] \tag{8}$$

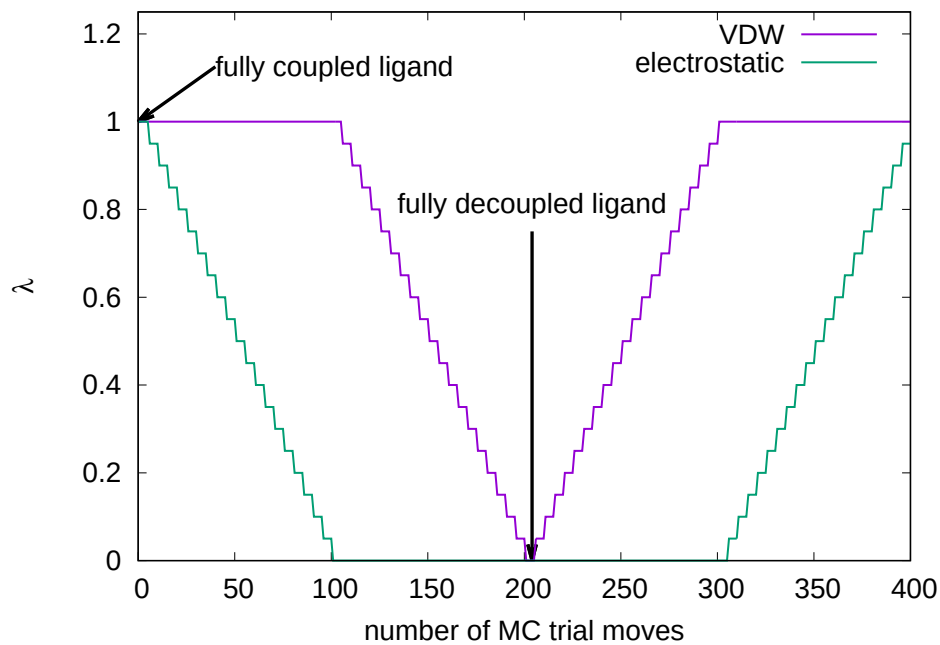

Fig 2. Scaling of λ_{VDW} and λ_{elec} over each NCMC trial-move cycle. The cycle comprises four phases, each containing 100 individual MC trial moves. First, the charges on the ligand are removed by gradually reducing λ_{elec} to 0. Second, the ligand is fully uncoupled from the protein by reducing λ_{VDW} to 0. Third, the van der Waals interactions of the ligand with the protein are restored by increasing λ_{VDW} back to 1. Finally, the charges on the ligand are restored by increasing λ_{elec} back to 1.

In this equation, the index i enumerates values of λ_{VDW} or λ_{elec} that are used during the move (horizontal segments of the graph in Fig. 2) and there is a term in the sum for each change in the λ parameters. The full sequence of MC moves that had taken place during the MC trial was then finally accepted or rejected with probability given by

$$p_{\mathrm{final}} = \min[1, \exp(-\beta w)] \tag{9}$$

If the NCMC move was rejected, the conformation of the system prior to the sequence of MC steps making up the given NCMC move was restored. Because of this, and the low acceptance rate of NCMC moves (about 2%), in pure NCMC simulations the relaxation of the system toward low energy conformations was extremely slow. Consequently, in order to promote more rapid relaxation, the NCMC moves were alternated with 400 moves of regular MC.

Docking protocol

A total of 61 ligands were used in this work, 37 agonists and 24 antagonists. Each ligand had a corresponding reference crystal structure, drawn from the PDB, showing the experimentally determined bound structure. Lists of the ligands and their corresponding reference structures are found in tables S1 Table and S2 Table. Figure 3 shows example ligand structures, where the larger size and flexibility of antagonists is notable.

Depending on whether the compound in question was an agonist or antagonist, it

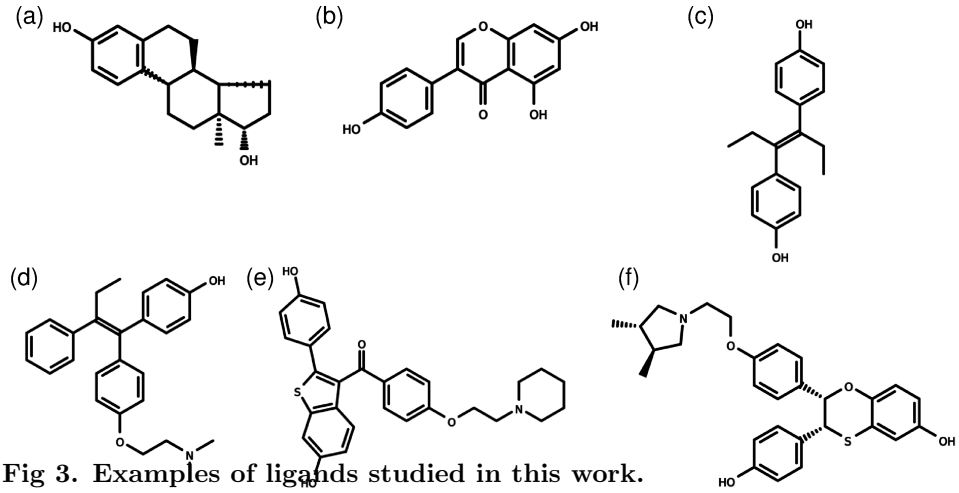

Fig 3. Examples of ligands studied in this work.
(a) estradiol; (b) genistein; (c) diethylstilbestrol; (d) 4-hydroxytamoxifen; (e) raloxifene; (f) the drug AIU. (a)-(c) are agonists, whereas (d)-(f) are antagonists.

was cross-docked against either the active or inactive conformation of ER. (Although in principle the MRMC method is capable of simulating the transition between the active and inactive conformations of ER, such a simulation would likely be longer than the docking runs used here, so docking of agonists against the inactive conformation or of antagonists against the active conformation was not attempted.) The active conformation was taken from the crystal structure of ER in complex with estradiol (PDB code 1QKU) [57] whereas the inactive conformation was taken from the crystal structure in complex with 4-hydroxytamoxifen (PDB code 3ERT). [58] Hydrogen atoms were added using the tleap tool in AMBER, and ionization states for the twelve histidine residues were chosen based on a combination of pKa calculations made using H++ [70, 71] and visual inspection of the crystal structures. Histidine residues 356, 373, 398, 476, 488, 501, 513, and 516 were chosen to be neutral, while histidine residues 377, 474, and 547 were chosen to be ionized. In the active conformation, His 524 makes a hydrogen bond to the hydroxyl group on C17 of estradiol, and it was calcuated to have a pKa of 5.67 using H++, implying a neutral state. In contrast, in the inactive conformation, His 524 makes a salt bridge with Glu 419 and was found to have a pKa of 7.97, implying an ionized state. Given the ambiguous nature of His 524's ionization state and its potential importance for the accuracy of docking results, it was decided to conduct half of the docking runs with an ionized His 524 and half with a neutral His 524.

Structure data files for each ligand were downloaded in SDF format from the Protein Data Bank and used to generate initial coordinates for each ligand. As a part of this process, ionization states for each ligand at pH 7 were determined using OpenBabel [72] and hydrogen atoms were added accordingly.

Fig. 4 gives an overview of the docking and analysis procedure. Each docking run consisted of a search for initial low-energy poses followed by refinement using Monte Carlo simulation. For the initial search, 1000 random poses were generated by placing the ligand in a random position and orientation within 4 A of the center of mass of estradiol (for agonists) or 4-hydroxytamoxifen (for antagonists) in the crystal structures and rotating every rotatable bond in the ligand through a random angle. The energy of each of these random poses was computed and a Monte Carlo

simulation lasting either 40000 trial moves (for regular MC) or 80000 trial moves (for mixed NCMC/MC) was started from the lowest energy pose. A total of 60 docking runs were performed with each His 524 ionization state as described above, for a total of 120 docking runs per drug overall. Each docking run took 2-4 hours on a single CPU, for a total of approximately 240 to 480 CPU-hours per drug.

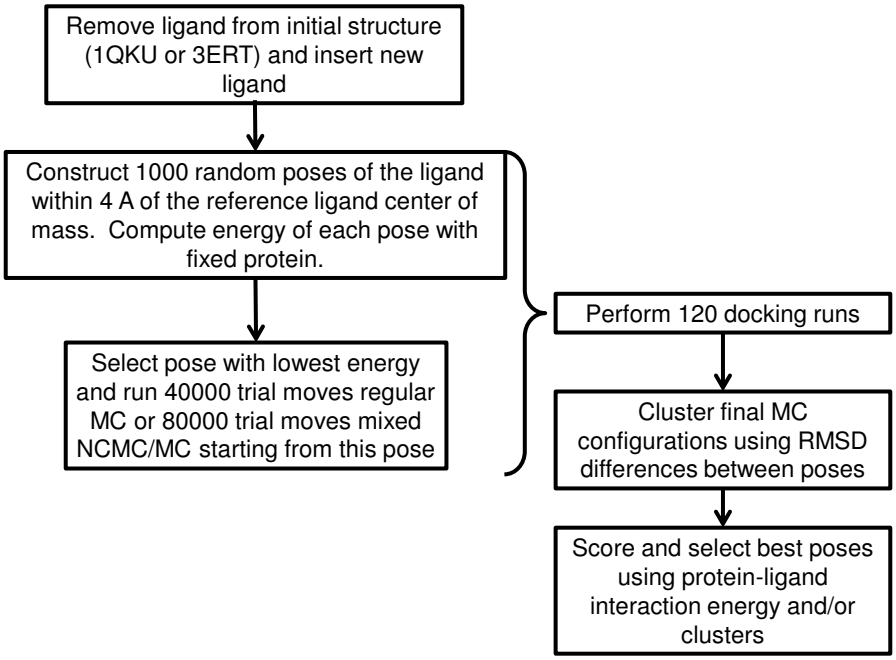

Fig 4. Flowchart showing overview of docking, clustering and scoring procedures used in this paper.

Structural Analysis and Clustering

An important measure of the performance of a docking method is the heavy-atom RMSD of the ligand in the docked pose to a known crystal pose. These measurements could be made for all the ligands tested in this work, since corresponding crystal structures were available. Ligand RMSDs were measured with the protein backbone aligned to the reference structure, and taking into account all possible mappings of chemically equivalent groups on the ligand. The measurements were made using VMD. [73]

Clustering of structures was accomplished by first aligning the structures for each ligand according to the protein backbone and then constructing a pairwise distance matrix among them using ligand heavy atom RMSD as the metric. Complete linkage hierarchical clustering [74] was then applied to this distance matrix to construct a "phylogenetic tree" using ligand heavy atom RMSD taking into account chemically equivalent groups. The clusters were then defined from this tree, using a cutoff defined as the 10th percentile of all distances between structures for each drug. This ensured that the size of the clusters in configuration space would be appropriately scaled to the overall distribution of the poses for each drug.

Results

Impact of Protein Flexibility

We first examined the impact of flexibility on pose generation. Fig. 5 shows scatterplots of ligand RMSD versus interaction energy for two representative agonists (estradiol and genistein) and two representative antagonists (4-hydroxytamoxifen and raloxifene). Since the starting structures used were those of ER α bound to estradiol or 4-hydroxytamoxifen, the simulations with these ligands represent redocking, whereas those with all other ligands represent cross-docking. In fig. 5, simulations in which the protein is fully flexible are compared with simulations in which the protein is rigid, or in which only the sidechains are allowed to move. Full protein flexibility results in interaction energies that are lower than those obtained with flexible side chains, which are in turn lower than those obtained with a completely fixed protein. In some cases – e.g. Panel (d) – the lack of flexibility prevents discovery of low-RMSD poses. This demonstrates that the use of protein flexibility, and particularly backbone flexibility, results in final configurations with stronger interactions between protein and ligand; this point is quantified further below.

The distribution of RMSDs suggests that the final configurations can be divided into clusters in some cases. This is particularly true for the agonists estradiol and genistein, which are relatively flat, rigid drugs and can fit into the binding site in multiple orientations. In these cases, the clusters correspond to these distinct orientations of the drugs. The ER antagonists 4-hydroxytamoxifen and raloxifene are more flexible; consequently, the clusters due to binding in multiple orientations are less clear.

Fig. 5 also shows a comparison between regular Monte Carlo simulations and those conducted with the mixed NCMC/MC protocol. The mixed NCMC/MC simulations had slightly higher interaction energies than the regular MC simulations.

The effect of protein flexibility is also demonstrated by a comparison of the "best" poses (selected on the basis of protein-ligand interaction energy) generated by our protocol, shown in Fig. 6. Figs. 6a and 6b show the distribution of "best" pose RMSDs across all of the drugs, expressed as a cumulative density function. Figs. 6c and 6d show a different approach to evaluating our protocol. For each drug, the best N poses are selected on the basis of protein-ligand interaction energy, the minimum RMSD pose is selected as a function of these, and the RMSD is averaged over all drugs. This average RMSD is then plotted as a function of N. It is clear that a much greater proportion of these best poses are close to the corresponding crystal structures when the protein is allowed to be fully flexible than when the protein is rigid or only sidechains move, and that the average best RMSD achieved is lower as well when more flexibility is allowed. This is especially true for the antagonists, which are generally more flexible than the agonists and whose docking is consequently more challenging. This difference is probably because the inactive structure of ER α is more open and consequently the protein as a whole, particularly helix 12, is more flexible. Consequently, the additional flexibility provided by the MRMC method may be more important for the inactive state.

In principle, allowing protein flexibility should enable us also to predict the changes in protein structure that occur upon docking to different ligands. We constructed Ramachandran and Janin plots [75] for the amino acid residues in the atomistic region for the final conformations from our docking runs. Fig. 7 shows these plots for two representative compounds, genistein and raloxifene. Although the overall conformation of ER around the active site is very similar for all bound agonists (and likewise for antagonists) there are slight differences in the backbone conformation for different ligands. The Janin plots (Figs. 7b and 7d) show relatively thorough sampling

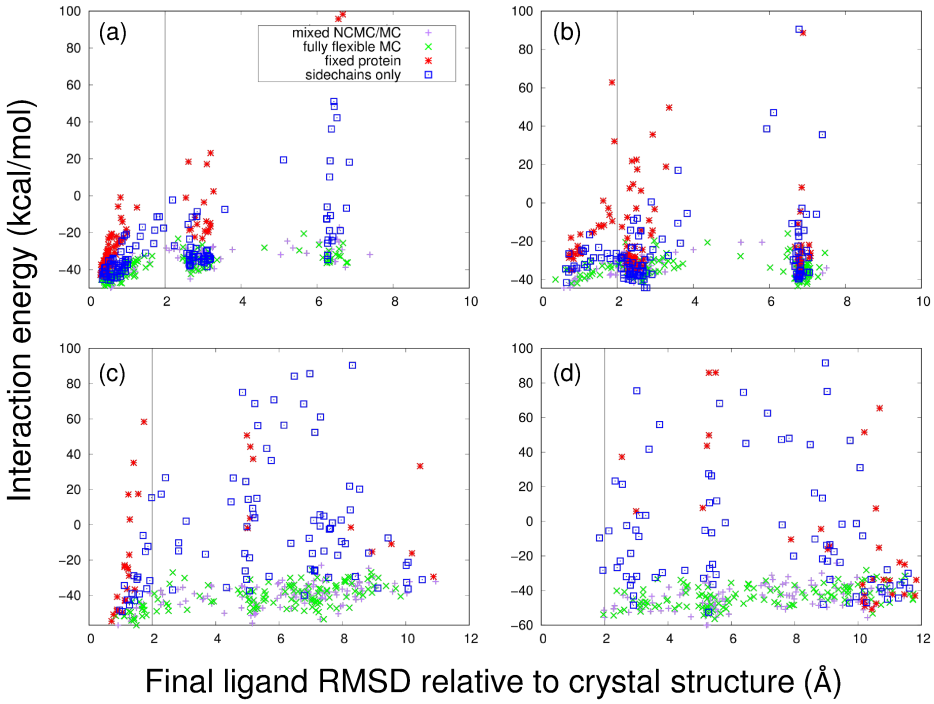

Fig 5. Example ensemble redocking and cross-docking runs for simulations incorporating different levels of flexibility.
Plot of final ligand RMSD relative to crystal structure vs. interaction energy for (a) estradiol (redocking); (b) genistein (cross-docking); (c) 4-hydroxytamoxifen (redocking); (d) raloxifene (cross-docking). In each plot, results for a fully flexible protein are in green (for MC only) or purple (for the mixed NCMC/MC simulations), whereas results for docking simulations in which the entire protein or just its backbone are fixed (MC only simulations) are in red or blue respectively. In order to make the results for the fully flexible protein visible, the vertical axis for each plot is cut off at 100 kcal/mol; as a result, some of the docking runs for the fixed protein are not shown due to their high energies.

of free energy basins in the χ_1-χ_2 plane. The Ramachandran plots (Figs. (Figs. 7a and 7c), on the other hand, show that the backbone sampling is very limited (at least in the binding region) and that final conformations are remaining close to the initial structures rather than adapting to the different ligands. Nevertheless, the overall backbone flexibility which includes the CG region evidently has a significant effect on the generated poses, as shown in fig. 6.

Assessment of NCMC

Fig. 6 also compares the performance of our docking protocol with and without NCMC. The use of NCMC shows at best a modest improvement in the overall docking performance, indicating that the NCMC is not enhancing sampling as much as was expected. In order to investigate this further, we measured both the distribution of overall ligand move sizes generated as well as the average acceptance probability as a function of move size. In the case of the NCMC simulations the size of the overall NCMC move depended on the individual MC moves that were performed within that move. To assess these factors, for each NCMC move, the overall rotation and

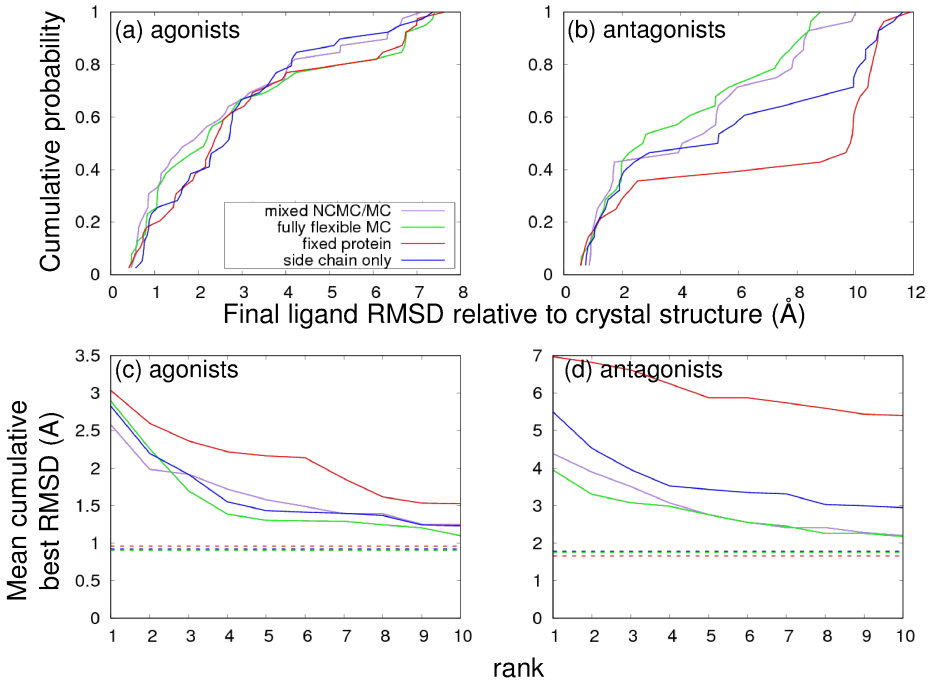

Fig 6. Performance of docking protocols.
(a) and (b) show the cumulative probability distribution of RMSD values for the docking runs with the lowest interaction energies for each drug, for (a) agonists or (b) antagonists. In (c) and (d), for each drug the docking runs with the N lowest interaction energies are chosen and the best RMSD from among these is averaged across (c) agonists or (d) antagonists. This average RMSD is plotted against N; dashed horizontal lines indicate the average best RMSD overall, without regard for interaction energy.

translation of the ligand in the frame of reference of the protein was determined by first performing an RMSD alignment of the protein backbone of the final configuration relative to the initial configuration, then measuring the overall translation and rotation of the ligands needed to minimize the RMSD of the ligand heavy atoms in the two configurations. The cumulative density function of the overall magnitude of the displacement or the overall angle of rotation, and the average final acceptance probability (given by eq. 9) as a function of the overall displacement or rotation angle, were both computed and plotted as shown in Fig. 8. The corresponding distribution of move sizes and average acceptance probability were also plotted for comparison.

The plots show that, compared to regular MC, the distribution of moves generated by NCMC favors smaller moves in both translation and rotation. In addition, while acceptance rates for small moves are comparable for NCMC and MC, the acceptance rates for larger moves are many orders of magnitude smaller for NCMC. This is particularly (and unexpectedly) true for ligand rotations, where acceptance rates for all but the smallest rotations are much smaller for NCMC than for MC. The combined effect of these two trends is that in NCMC, a greater number of small translations and rotations are generated and accepted, compared to MC. The reason for these results seems to be that large ligand translations and rotations typically place the ligand in positions that clash sterically with the protein. Relaxing away such clashes evidently requires large nonequilibrium work and consequently leads to rejections. We note that these results are specific to our implementation, as described below in the Discussion.

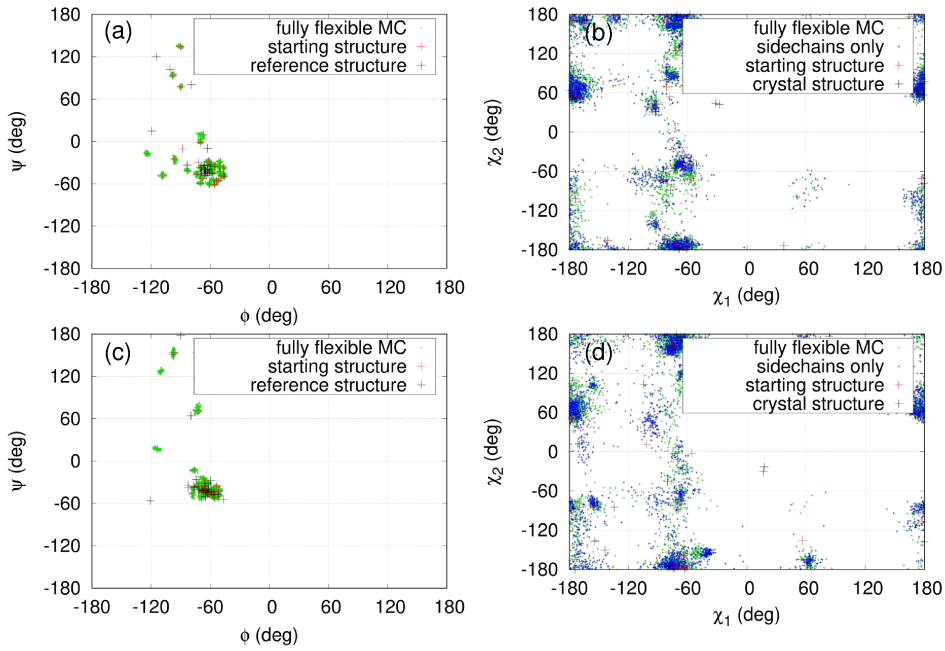

Fig 7. Examination of backbone and sidechain flexibility for two compounds.
Ramachandran and Janin (χ_1 vs. χ_2) plots are shown for for final conformations from docking runs for (a)-(b) genistein and (c)-(d) raloxifene. Plots include only amino acids in the atomistic region. The reference structure is the corresponding crystal structure for each cross-docked compound.

Fig. 9 shows time series of the ligand RMSD for a number of individual docking simulations in both regular MC and NCMC. In many of the simulations, large jumps in RMSD can be seen; these represent large ligand moves that have been accepted. While the docking runs with the mixed NCMC/MC protocol are twice as long as those with regular MC, more than twice as many RMSD jumps can be seen in trajectories using the mixed NCMC/MC protocol. To reconcile this result with the acceptance rate data shown in Fig. 8, note that some of the NCMC transitions occur via a large number of small moves. It should also be noted that NCMC moves also include moving the protein, whereas individual MC moves that involve ligand translation and rotation do not, so Fig. 9 is not a perfect comparison.

Assessment of Clustering in Pose Scoring

The fact that some of the docking simulations show substantial changes in the orientation of the ligand relative to the protein suggests that the ensemble of final conformations generated by the MRMC protocol contains information on multiple binding poses which could make a useful contribution to the docking and scoring process. To study this possibility, the ensemble of structures resulting from the docking runs on each drug were divided into clusters based on the protocol described in Sec. . Three options for choosing the best conformation from the ensemble were then tested:

1. Based on the observation that, for most drugs, the final interaction energy between protein and ligand correlated with ligand RMSD (Fig. 5), the simplest

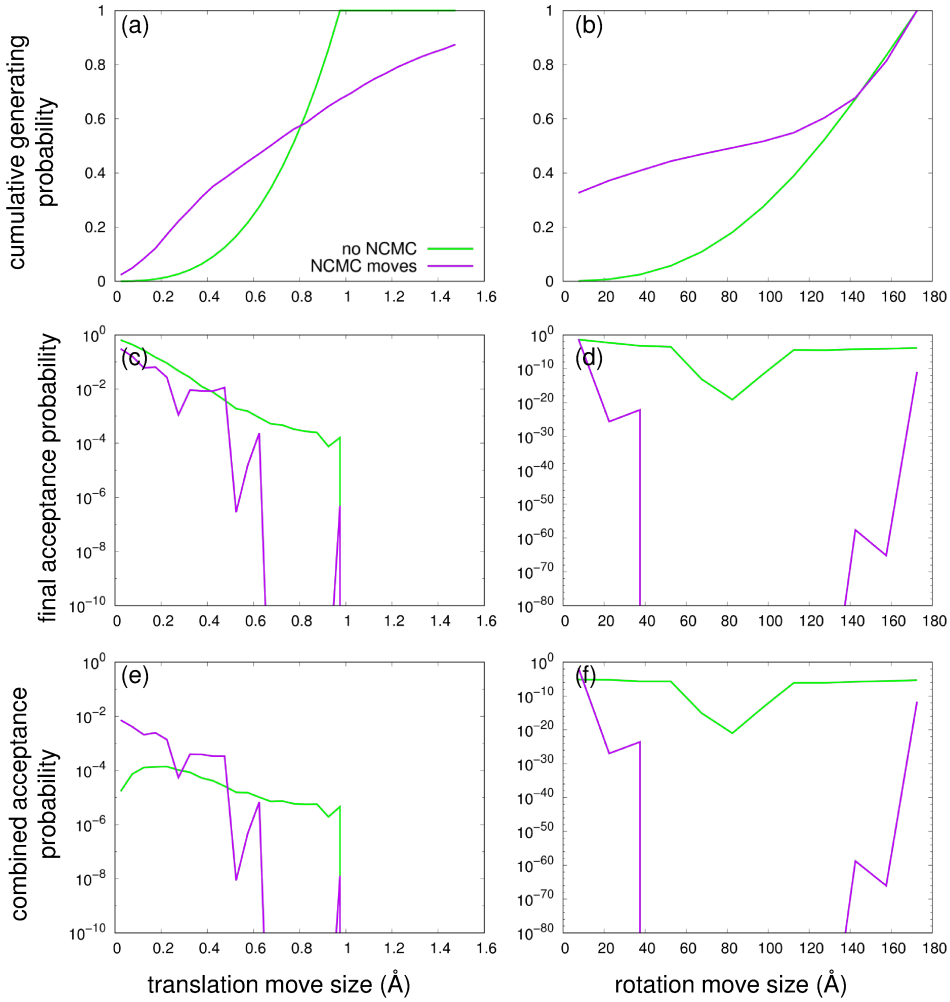

Fig 8. Analysis of NCMC and MC acceptance probabilities in docking simulations of estradiol to the active conformation of ER α as a function of move size.
(a), (c), and (e) are for translations; (b), (d), and (f) are for rotations. (a)-(b) Cumulative probability density function of generated move sizes. (c)-(d) Average final acceptance probability of NCMC and MC moves as a function of move size. (e)-(f) Combined acceptance probability (the overall proportion of all NCMC or MC moves that were accepted and of the given size) which is the product of the generating probability.

approach is simply to choose the conformation with the lowest interaction energy, without regard to clusters.

2. Since the largest cluster represents a binding pose that has the highest entropy relative to the other clusters, the conformation within this cluster with the lowest interaction energy could be chosen.

3. Finally, the cluster with the lowest average energy could be chosen, and then the individual conformation with the lowest interaction energy could be chosen from this cluster.

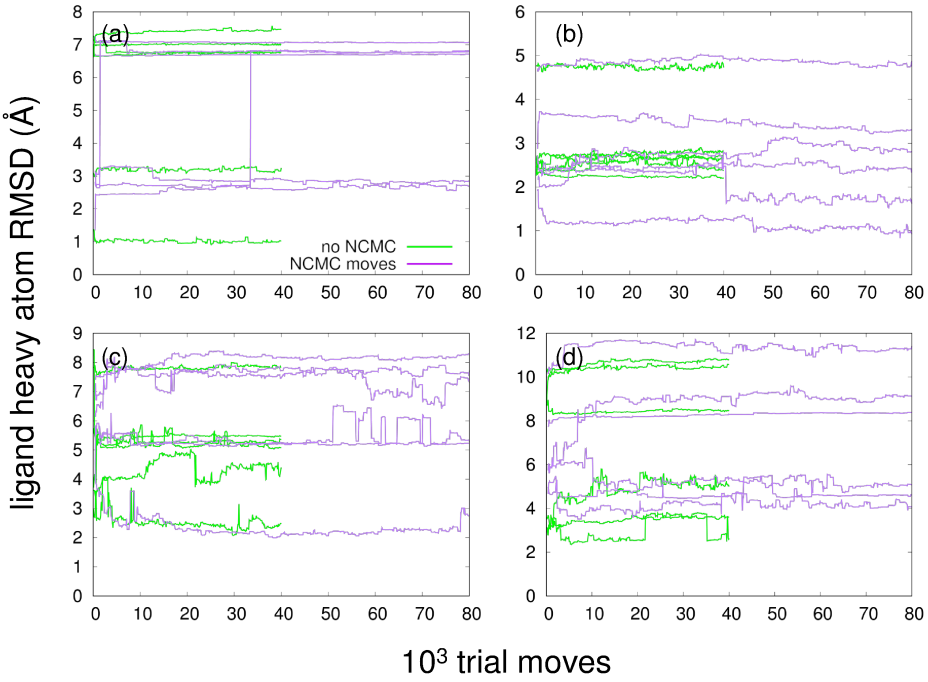

Fig 9. Plots of ligand RMSD relative to crystal structure vs. number of trial moves in docking simulations.
(a) genistein; (b) diethylstilbestrol; (c) raloxifene; (d) the drug AIU. Six out of 120 docking runs are shown for each drug.

Fig. 10 shows the cumulative distribution of the RMSD of the best conformation for all three of these methods. The differences between them are small, but it appears that method 3 above performs slightly better than the others in that the conformations identified by this method are closer to the crystal structures.

Note that we did not employ a more quantitative entropy estimation process because the sampling did not appear to be sufficient - i.e., there were very few jumps between poses in a given MC run (Fig. 9) implying true Boltzmann sampling was not achieved. Also, using the total energy in place of the protein-ligand interaction energy and combining it with similar clustering approaches gave similar results.

Multiple Poses for WAY-169916

The crystal structure of the ER partial agonist WAY-169916 (PDB code 3OS9, ligand ID KN1) [17] bound to the inactive conformation of ER α shows two separate poses for the ligand, with occupancies of approximately 70% and 30%. This provided an opportunity to test whether our protocol can find multiple bound poses for a ligand. The RMSD to each of these poses was calculated separately for our docking simulations of this drug, and a scatterplot is shown in Fig. 11. The ensemble of docked poses for this drug contains two distinct clusters that correspond to the poses in the crystal structure (the closest poses found are approximately 2 Å away from each crystal pose), along with other poses that are similar to neither. This demonstrates that our protocol is able to find both of these poses.

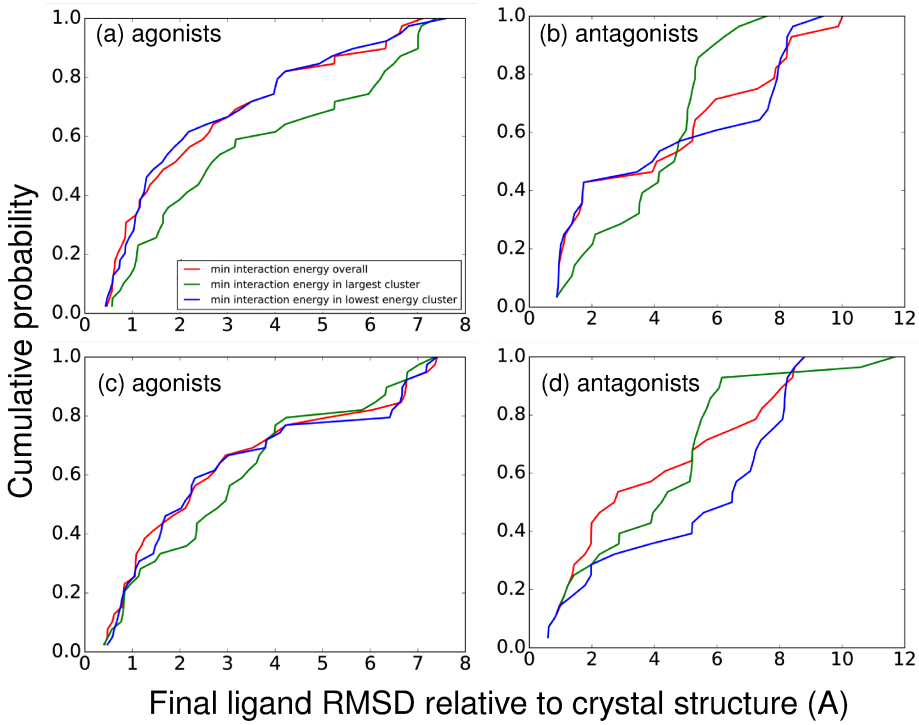

Fig 10. Comparison of methods for ranking ligands for results of simulations with a fully flexible protein.
(a)-(b) using NCMC. (c)-(d) without NCMC. Each graph shows the cumulative distribution of the RMSD of the best conformation (chosen by the indicated method) for (a),(c) agonists and (b),(d) antagonists.

His 524 Ionization States

His 524 plays an important role in the binding of ligands to ER α. For example, the O3 atom of estradiol forms a hydrogen bond with a neutral His 524, whereas when 4-hydroxytamoxifen is bound, His 524 instead is ionized and forms a salt bridge with Glu 419. Because of these changes in the ionization state of His 524, and its importance in ligand binding, half of the docking runs were performed with His 524 in an ionized state and half were performed with His 524 in the neutral state. The results for different ionization states of His 524 were generally similar, with approximately equal numbers of successful docking runs (those with a final RMSD less than 2 Å) coming from runs in which His524 was ionized or neutral. Likewise, simulations with both ionized and neutral His 524 produced docked conformations of WAY-169916 that were close to each of the experimental structures.

Computation time

Table 3 shows a comparison of computation speeds for fully atomistic, mixed-resolution, and coarse-grained representations of ER α. As expected, the Gō model is about 60 times faster than a fully atomistic representation. The mixed resolution model is about four times faster than a fully atomistic representation in vacuum, and about twice as fast when the SEDDD solvation model is used. Thus a mixed-resolution model offers a modest savings in compute time, but there is additional sampling benefit from the landscape-smoothing implicitly provided by coarse-graining. As noted, an MRMC platform also offers significant flexibility in

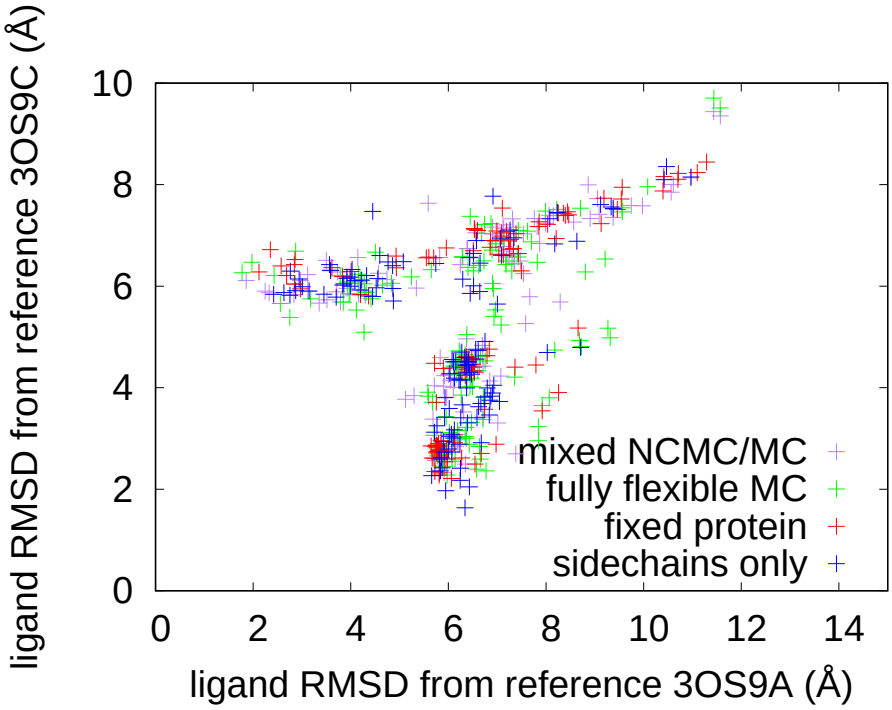

Fig 11. Docking uncovers two crystal poses of the ligand WAY-169916. We show a scatter plot of ligand RMSD relative to the two reference poses found in its reference crystal structure.

implementing docking protocols. For reference, 10 ns of all-atom explicit solvent MD requires about 5 hours with AMBER and one GPU for this system.

Table 3. Comparison of computation speeds (MC trial moves per second) for different resolution representations of ER α in complex with estradiol.

	In vacuum		With SEDDD	
	With NCMC	MC only	With NCMC	MC only
All atom	31.69	32.52	17.70	18.39
Mixed resolution	85.23	88.23	37.33	37.67
Coarse grained only	n/a*	1279.26	n/a*	1279.26

*Coarse grained only results are for uncomplexed ER α; it is not possible to use NCMC without a ligand.

Discussion

In this paper, three separate tactics for improving protein-ligand docking were tested. These included a mixed-resolution potential in which most of the protein is treated using a coarse grained model while the region around the ligand was treated atomistically; a nonequilibrium Monte Carlo method, which is intended to improve sampling by systematically varying the coupling between protein and ligand; and the use of clustering to identify free energy basins corresponding to multiple binding poses, and scoring the poses based on this information. The docking results were evaluated by comparing the final structures to known crystal structures; the sampling

was also evaluated by studying the relationship between acceptance rate and move size for NCMC moves. We found that allowing for full protein flexibility using the mixed-resolution potential significantly improved the docking results, and the use of NCMC produces a further modest improvement. However, clustering did not appear to offer any significant advantages over simply ranking the poses by protein-ligand interaction energy. Overall, we obtained a correct pose within 2 Å for about half the ligands, so there is significant room for improvement. Below, we discuss ways that several aspects of the approach could be improved. We note that systematic investigation of these different aspects of the docking problem is facilitated by having a highly flexible/adjustable Monte Carlo platform.

Improving the MRMC Potential

The mixed-resolution model used here is motivated by the concept that the most important approximations in protein-ligand binding will be those between the ligand and the closest amino acid residues within the protein. Therefore, it makes the most sense to model the closest amino acids at the fully atomistic level, while saving computation time by modeling the remainder of the protein using a more approximate coarse-grained model. That said, it is also reasonable to examine every part of the mixed-resolution potential to see if they can be made more physically accurate, while continuing to save computer time over fully atomistic approaches. These include the coarse-grained force field, the atomistic force field, the coupling between them, and the choice of atomistic region.

In this work, a relatively simple Gō model was used to represent the coarse-grained region of ER α. This is effective in saving computer time. On the other hand, its reliance on native interactions gives it a strong bias toward the native state, allowing only limited conformational flexibility outside the atomistic region. In the case of ER α, the only significant conformational change is the motion of helix 12, so any necessary conformational flexibility could be included by ensuring that helix 12 and the loop connecting it to the rest of the protein were part of the atomistic region. However, with other target proteins this may not be adequate. Additional flexibility could be incorporated by making use of a double well Gō potential [45] or by replacing the Gō potential by another potential that is less dependent on native interactions. The popular MARTINI force field [76, 77] has been used in a mixed resolution configuration [40] but frequently leads to distorted structures for soluble proteins unless reinforced by an elastic network model. [78] Other potential force fields that could in principle be used include OPEP [79] and UNRES. [80] We have also developed a tunable coarse-grained force field based on constructing interaction energy tables and applying variable amounts of smoothing to them [81]; one of the goals for this force field was to use it in a mixed resolution setting, but substantial additional effort would be required to implement that combination.

Compared to the coarse-grained potential, the atomistic potential used here would seem to offer less room for improvement. The AMBER 99SB force field is a well-tested, commonly used force field, although it could conceivably be replaced with another, newer force field. A more significant area for improvment concerns the treatment of solvation in the simulations. The SEDDD method is based on a distance dependent dielectric with a dielectric constant that includes some solvent exposure. The linear dependence of the dielectric constant on interatomic distance is not physically correct, however, since the dielectric constant should approach that of water as the distance between two atoms increases. There is also no explicit term representing the hydrophobic effect. A generalized Born model [82] would be more physically realistic but also more computationally expensive. Explicit solvent (using water molecules restrained around the atomistic region) would be even better, but

would require even more computational cost, and steric clashes between the ligand and water molecules would make it more difficult to sample ligand configurations via Monte Carlo. Another relatively inexpensive solvation model is the Sheffield solvation model, [83] which might provide improved results.

The choice of atomistic region also plays an important role in establishing the tradeoff between physical accuracy and computation time. We selected the atomistic region used here to be as small as possible while including all those residues in direct contact with the ligand in any of the reference structures. We also included helix 12 and its loop to allow for the possibility of transitions between the active and inactive conformations, although we did not observe any such transitions due to the short duration of the simulations. A larger atomistic region would trade computation speed for greater physical accuracy.

Improved Sampling

While crystal structures of protein-ligand complexes frequently show only one bound pose for the ligand, there is experimental evidence that some ligands can bind to proteins in multiple configurations. The two poses for WAY-169916 are a case in point. Likewise, differences have also been found in ligand binding to aldose reductase depending on the crystallization conditions [16]. In principle, with a sufficient number and length of runs, a docking algorithm based on Monte Carlo or molecular dynamics simulation should be able to find all relevant bound poses for the ligand with the proportion that would be expected based on their relative free energy. Although our MRMC runs indeed found multiple poses for ligands (including both experimental poses for WAY-169916), jumps between poses in a single run were extremely rare.

With our implementation of NCMC, the improvement over standard MC is marginal, in contrast to the success reported recently by the Mobley group for the T4 lysozyme/toluene system using MD. [55] This is likely because the ER α ligands studied here are larger than toluene, and shaped in such a way that the barriers separating different ligand orientations are higher than those separating distinct orientations of toluene bound to T4 lysozyme. In addition, our ligands are also more flexible, and we used a larger range of moves and attempted to sample a greater number of degrees of freedom. In an investigation of NCMC applied to side chain rotations of amino acids in explicit solvent, Kurut and coworkers found that NCMC did not enhance sampling for methionine as much as valine, because enhancing the sampling of one dihedral degree of freedom did not improve the sampling of other dihedrals. [84] We may be observing the same effect here, where couplings between the ligand position and orientation and the side chain degrees of freedom of nearby amino acids reduce the efficiency of NCMC. Alternatively, as suggested by Chodera [personal communication] our NCMC protocol may require better targeted ligand MC moves such as "smart darting," [85] although our initial tests of this idea did not show a substantial improvement.

There are also straightforward means to improve sampling. Two simple ways are to run longer docking simulations or perform more docking runs for each drug. The docking runs used here are fairly short (40000-80000 trial moves) which inherently limits the amount of sampling possible in a single docking run. Of course, both increasing the length and number of docking runs would increase the amount of CPU time needed to dock a drug. In addition, using a longer or otherwise redesigned schedule for λ_{VDW} might improve the NCMC acceptance rate, which was relatively low (approximately 2%) in the work reported here, and thereby improve the sampling. Another way to improve the NCMC acceptance rate might be to use a soft core potential for the van der Waals interactions. This might reduce the potential energy changes associated with steric clashes between the ligand and receptor as λ_{VDW} is

increased from 0 to 1, which contribute to large values of the nonequilibrium work and consequently to poor acceptance rates.

Use of Clustering Information

If a sufficient degree of sampling can be obtained, it should in principle be possible to identify basins in the free energy surface corresponding to different possible ligand or protein conformations. Each basin will correspond to a cluster of similar conformations obtained in the ensemble. In principle, once all of the basins are identified, if true Boltzmann sampling has been achieved the ensemble of conformations should also give information on the relative free energies of different binding conformations, which can then be used to calculate binding free energies. [86–88] Motivated by this reasoning, we sought to apply clustering algorithms to the ensemble of configurations we obtained from our docking simulations and use this information to aid in identifying the most representative structures. We found, however, that clustering information did not improve the ranking of poses in practice.

The main reason why the clustering was not useful may simply have been that the protocol used here did not allow for adequate sampling, as described above. Another flaw may have been the choice of clustering algorithm or metric used. The complete linkage clustering algorithm used here is relatively crude, being primarily intended for the construction of phylogenetic trees using distances between protein or DNA sequences. [74] Despite this, it was selected because many other algorithms, such as the K-means algorithm, rely on averaging coordinates from distinct configurations, an operation of unclear physical meaning. In addition, complete linkage clustering guarantees that any two configurations classified in the same cluster will have an RMSD less than the selected cutoff. The Cheatham group has tested a number of clustering algorithms on MD trajectories; [89] while they recommend average-linkage hierarchical clustering for circumstances in which the number of clusters is not known in advance (as here) they also point out that the performance of a clustering algorithm is influenced by the choice of atoms used for pairwise comparison and that hierarchical clustering is sensitive to outliers.

Conclusion

We used a highly adjustable mixed-resolution Monte Carlo (MRMC) platform to examine several aspects of docking protocols in a systematic way. Most importantly, we examined the effects of rigidifying both side chains and the protein backbone. The detrimental results are not completely surprising, but the systematic comparison underscores the importance of backbone flexibility, which is absent from almost all grid-based docking studies. We further examined the sampling improvement afforded by non-equilibrium candidate Monte Carlo, finding only modest improvement in our implementation. Our test case was the flexible ligand binding domain of the estrogen receptor alpha, which is an important cancer target and also a model for other nuclear hormone receptors.

As computing power increases, a ‘middle way’ of docking between grid-based approaches and all-atom free energy calculations may prove useful in drug-design pipelines. This study is a step toward developing such a highly adaptable platform. We recognize that further improvements to sampling and entropy-based pose evaluation will be necessary to make a middle-way tool more valuable for the drug-design enterprise.

Supporting information

S1 Table. Agonists used in this work

S2 Table. Antagonists used in this work.

Acknowledgments

We thank David Mobley and his group for the suggestion to use the nonequilibrium Monte Carlo method and for valuable discussions. We also acknowledge Andy Stern, John Chodera, David Koes, Jocelyn Sunseri, Ernesto Suarez, Barmak Mostofian and Apoorva Shrivastava for helpful discussions. We also thank the Department of Computational and Systems Biology and the Center for Research Computing at the University of Pittsburgh for computer time. This work was supported by NIH grant no. P41-GM103712, NSF grant nos. MCB-1119091 and CNS-1229064, and a Commonwealth Universal Research Enhancement Program grant from the Commonwealth of Pennsylvania Department of Health (SAP 4100062224).

References

1. Navia MA, Fitzgerald PMD, McKeever BM, Leu CT, Heimbach JC, Herber WK, et al. Three-dimensional structure of aspartyl protease from human immunodeficiency virus HIV-1. Nature. 1989;337:615–620.

2. Wlodawer A, Vondrasek J. Inhibitors of HIV-1 Protease: A Major Success of Structure-Assisted Drug Design. Annu Rev Biophys Biomol Struct. 1998;27:249–284.

3. Shoichet BK, McGovern SL, Wei BQ, Irwin JJ. Lead discovery using molecular docking. Curr Opin Struc Biol. 2002;6(4):439–446.

4. Kitchen DB, Decornez H, Furr JR, Bajorath J. Docking and Scoring in Virtual Screening for Drug Discovery: Methods and Applications. Nature Reviews. 2004;3:935–949.

5. Klebe G. Virtual ligand screening: Strategies, perspectives and limitations. Drug Discov Today. 2006;11:580–594.

6. Jorgensen WL. Efficient Drug Lead Discovery and Optimization. Acc Chem Res. 2009;42(6):724–733.

7. Trott O, Olson AJ. AutoDock Vina: Improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading. J Comput Chem. 2010;31(2):455–461.

8. Koes DR, Baumgartner MP, Camacho CJ. Lessons Learned in Empirical Scoring with smina from the CSAR 2011 Benchmarking Exercise. J Chem Inf Model. 2013;53(8):1893–1904.

9. Halgren TA, Murphy RB, Friesner RA, Beard HS, Frye LL, Pollard WT, et al. Glide: A New Approach for Rapid, Accurate Docking and Scoring. 2. Enrichment Factors in Database Screening. J Med Chem. 2004;47(7):1750–1759.

10. Friesner RA, Banks JL, Murphy RB, Halgren TA, Klicic JJ, Mainz DT, et al. Glide: A New Approach for Rapid, Accurate Docking and Scoring. 1. Method and Assessment of Docking Accuracy. J Med Chem. 2004;47(7):1739–1749.

11. Friesner RA, Murphy RB, Repasky MP, Frye LL, Greenwood JR, Halgren TA, et al. Extra Precision Glide: Docking and Scoring Incorporating a Model of Hydrophobic Enclosure for Protein-Ligand Complexes. J Med Chem. 2006;49(21):6177–6196.

12. Wu G, Robertson DH, III CLB, Vieth M. Detailed Analysis of Grid-Based Molecular Docking: A Case Study of CDOCKER–A CHARMm-Based MD Docking Algorithm. J Comput Chem. 2003;24:1549–1562.

13. Jimenez-Garcia B, Roel-Touris J, Romero-Durana M, Vidal M, Jimenez-Gonzalez D, Fernandez-Recio J. LightDock: a new multi-scale approach to protein-protein docking. Bioinformatics. 2018;34(1):49–55.

14. Lorber DM, Shoichet BK. Flexible ligand docking using conformational ensembles. Protein Sci. 1998;7(4):938–950.

15. Osguthorpe DJ, Sherman W, Hagler AT. Exploring Protein Flexibility: Incorporating Structural Ensembles From Crystal Structures and Simulation into Virtual Screening Protocols. J Phys Chem B. 2012;116(23):6952–6959.

16. Steuber H, Zentgraf M, Gerlach C, Sotriffer CA, Heine A, Klebe G. Expect the unexpected or caveat for drug designers: Multiple structure determinations using aldose reductase crystals treated under varying soaking and co-crystallisation conditions. J Mol Biol. 2006;363(1):174–187.

17. Bruning JB, Parent AA, Gil G, Zhao M, Nowak J, Pace MC, et al. Coupling of receptor conformation and ligand orientation determine graded activity. Nat Chem Biol. 2010;6(11):837–843.

18. Korb O, Olsson TSG, Bowden SJ, Hall RJ, Verdonk ML, Liebeschuetz JW, et al. Potential and Limitations of Ensemble Docking. J Chem Inf Model. 2012;52(5):1262–1274.

19. Bowman GR, Geissler PL. Equilibrium fluctuations of a single folded protein reveal a multitude of potential cryptic allosteric sites. Proc Natl Acad Sci U S A. 2012;109(29):11681–11686.

20. Srinivasan S, Nwachukwu JC, Parent AA, Cavett V, Nowak J, Hughes TS, et al. Ligand-binding dynamics rewire cellular signaling via estrogen receptor α. Nat Chem Biol. 2013;9:326.

21. Hart KM, Ho CMW, Dutta S, Gross ML, Bowman GR. Modelling proteins' hidden conformations to predict antibiotic resistance. Nat Commun. 2016;7:10.

22. Antunes DA, Devaurs D, Kavraki LE. Understanding the challenges of protein flexibility in drug design. Expert Opin Drug Discov. 2015;10(12):1301–1313.

23. Amaro RE, Li WW. Emerging Methods for Ensemble-Based Virtual Screening. Curr Top Med Chem. 2010;10:3–13.

24. Tian S, Sun HY, Pan PC, Li D, Zhen XC, Li YY, et al. Assessing an Ensemble Docking-Based Virtual Screening Strategy for Kinase Targets by Considering Protein Flexibility. J Chem Inf Model. 2014;54(10):2664–2679.

25. Gagnon JK, Law SM, III CLB. Flexible CDOCKER: Development and Application of a Pseudo-Explicit Structure-Based Docking Method Within CHARMM. J Comput Chem. 2016;37:753–762.

26. de Vries SJ, Zacharias M. Fast and Accurate Grid Representations for Atom-Based Docking with Partner Flexibility. J Comput Chem. 2017;38(17):1538–1546.

27. Cheron JB, Zacharias M, Antonczak S, Fiorucci S. Update of the ATTRACT force field for the prediction of protein-protein binding affinity. J Comput Chem. 2017;38(21):1887–1890.

28. Meiler J, Baker D. ROSETTALIGAND: Protein-small molecule docking with full side-chain flexibility. Proteins. 2006;65(3):538–548.

29. Davis IW, Baker D. ROSETTALIGAND Docking with Full Ligand and Receptor Flexibility. J Mol Biol. 2009;385:381–392.

30. DeLuca S, Khar K, Meiler J. Fully Flexible Docking of Medium Sized Ligand Libraries with RosettaLigand. PLoS One. 2015;10(7):19.

31. Singh N, Warshel A. Absolute binding free energy calculations: On the accuracy of computational scoring of protein-ligand interactions. Proteins. 2010;78:1705–1723.

32. Mobley DL, Klimovich PV. Perspective: Alchemical free energy calculations for drug discovery. J Chem Phys. 2012;137:230901.

33. Mobley DL, Gilson MK. Predicting Binding Free Energies: Frontiers and Benchmarks. Ann Rev Biophys. 2017;46:531–558.

34. Mobley DL. Let's get honest about sampling. J Comput Aided Mol Des. 2012;26(1):93–95.

35. Hansen N, van Gunsteren WF. Practical Aspects of Free-Energy Calculations: A Review. J Chem Theory Comput. 2014;10(7):2632–2647.

36. Miller BR, McGee TD, Swails JM, Homeyer N, Gohlke H, Roitberg AE. MMPBSA.py: An Efficient Program for End-State Free Energy Calculations. J Chem Theory Comput. 2012;8(9):3314–3321.

37. Genheden S, Ryde U. The MM/PBSA and MM/GBSA methods to estimate ligand-binding affinities. Expert Opin Drug Discov. 2015;10(5):449–461.

38. Neri M, Anselmi C, Cascella M, Maritan A, Carloni P. Coarse-grained model of proteins incorporating atomistic detail of the active site. Phys Rev Lett. 2005;95(21):4.

39. Mamonov AB, Lettieri S, Ding Y, Sarver JL, Palli R, Cunningham TF, et al. Tunable, Mixed-Resolution Modeling Using Library-Based Monte Carlo and Graphics Processing Units. J Chem Theory Comput. 2012;8(8):2921–2929.

40. Wassenaar TA, Ingolfsson HI, Priess M, Marrink SJ, Schaefer LV. Mixing MARTINI: Electrostatic Coupling in Hybrid Atomistic-Coarse-Grained Biomolecular Simulations. J Phys Chem B. 2013;117(13):3516–3530.

41. Fogarty AC, Potestio R, Kremer K. A multi-resolution model to capture both global fluctuations of an enzyme and molecular recognition in the ligand-binding site. Proteins. 2016;84(12):1902–1913.

42. Kar P, Feig M. Hybrid All-Atom/Coarse-Grained Simulations of Proteins by Direct Coupling of CHARMM and PRIMO Force Fields. J Chem Theory Comput. 2017;.

43. Taketomi H, Ueda Y, Gō N. Studies on Protein Folding, Unfolding and Fluctuations by Computer Simulation. 1. Effect of Specific Amino-Acid Sequence Represented by Specific Inter-Unit Interactions. Int J Pept Protein Res. 1975;7(6):445–459.

44. Ueda Y, Taketomi H, Gō N. Studies on Protein Folding, Unfolding and Fluctuations by Computer Simulation. 2. 3-Dimensional Lattice Model of Lysozyme. Biopolymers. 1978;17(6):1531–1548.

45. Zuckerman DM. Simulation of an Ensemble of Conformational Transitions in a United-Residue Model of Calmodulin. J Phys Chem B. 2004;108:5127–5137.

46. Zhang BW, Jasnow D, Zuckerman DM. Efficient and verified simulation of a path ensemble for conformational change in a united-residue model of calmodulin. Proc Natl Acad Sci U S A. 2007;104:18043–18048.

47. Csermely P, Palotai R, Nussinov R. Induced fit, conformational selection and independent dynamic segments: an extended view of binding events. Trends Biochem Sci. 2010;35(10):539–546.

48. Carroll MJ, Gromova AV, Miller KR, Tang H, Wang XS, Tripathy A, et al. Direct Detection of Structurally Resolved Dynamics in a Multiconformation Receptor-Ligand Complex. J Am Chem Soc. 2011;133(16):6422–6428.

49. Takada S. Coarse-grained molecular simulations of large biomolecules. Curr Opin Struc Biol. 2012;22(2):130–137.

50. Hornak V, Abel R, Okur A, Strockbine B, Roitberg A, Simmerling C. Comparison of Multiple Amber Force Fields and Development of Improved Protein Backbone Parameters. Proteins. 2006;65:712–725.

51. Wang J, Wolf RW, Caldwell JW, Kollman PA, Case DA. Development and Testing of a General Amber Force Field. J Comput Chem. 2004;25:1157–1174.

52. Wang JM, Wang W, Kollman PA, Case DA. Automatic atom type and bond type perception in molecular mechanical calculations. J Mol Graph. 2006;25(2):247–260.

53. Garden DP, Zhorov BS. Docking flexible ligands in proteins with a solvent exposure- and distance-dependent dielectric function. J Comput Aided Mol Des. 2010;24(2):91–105.

54. Nilmeier JP, Crooks GE, Minh DDL, Chodera JD. Nonequilibrium candidate Monte Carlo is an efficient tool for equilibrium simulation. Proc Natl Acad Sci U S A. 2011;108(45):E1009–E1018.

55. Gill SC, Lim NM, Grinaway PB, Rustenburg AS, Fass J, Ross GA, et al. Binding Modes of Ligands Using Enhanced Sampling (BLUES): Rapid Decorrelation of Ligand Binding Modes via Nonequilibrium Candidate Monte Carlo. J Phys Chem B. 2018;122(21):5579–5598.

56. Jordan VC. In: Williams DA, Foye WO, Lemke TL, editors. Selective Estrogen Receptor Modulators. Baltimore: Lippincott Williams and Wilkins; 2002. p. 1059–1069.

57. Gangloff M, Ruff M, Eiler S, Duclaud S, Wurtz JM, Moras D. Crystal Structure of a Mutant hERα Ligand-binding Domain Reveals Key Structural Features for the Mechanism of Partial Agonism. J Biol Chem. 2001;276(18):15059–15065.

58. Shiau AK, Barstad D, Loria PM, Cheng L, Kushner PJ, Agard DA, et al. The structural basis of estrogen receptor/coactivator recognition and the antagonism of this interaction by tamoxifen. Cell. 1998;95(7):927–37.

59. Osborne CK, Fuqua SAW. Selective estrogen receptor modulators: Structure, function, and clinical use. J Clin Oncol. 2000;18(17):3172–3186.

60. Jordan VC. Selective estrogen receptor modulation: Concept and consequences in cancer. Cancer Cell. 2004;5(3):207–213.

61. Cuzick J, DeCensi A, Arun B, Brown PH, Castiglione M, Dunn B, et al. Preventive therapy for breast cancer: a consensus statement. Lancet Oncol. 2011;12(5):496–503.

62. Toy W, Shen Y, Won H, Green B, Sakr RA, Will M, et al. ESR1 ligand-binding domain mutations in hormone-resistant breast cancer. Nature Genet. 2013;45(12):1439–U189.

63. Robinson DR, Wu YM, Vats P, Su FY, Lonigro RJ, Cao XH, et al. Activating ESR1 mutations in hormone-resistant metastatic breast cancer. Nature Genet. 2013;45(12):1446–U197.

64. Fanning SW, Mayne CG, Dharmarjan V, Carlson KE, Martin TA, Novick SJ, et al. Estrogen receptor alpha somatic mutations Y537S and D538G confer breast cancer endocrine resistance by stabilizing the activating function-2 binding conformation. eLife. 2016;5:25.

65. Levy Y, Caflisch A, Onuchic JN, Wolynes PG. The folding and dimerization of HIV-1 protease: Evidence for a stable monomer from simulations. J Mol Biol. 2004;340(1):67–79.

66. Sulkowska JI, Cieplak M. Selection of optimal variants of Gō-like models of proteins through studies of stretching. Biophys J. 2008;95(7):3174–3191.

67. Cieplak M, Sulkowska JI. Tests of the Structure-Based Models of Proteins. Acta Phys Pol A. 2009;115(2):441–445.

68. Lazaridis T, Karplus M. Effective energy function for proteins in solution. Proteins. 1999;35:133–152.

69. Betancourt MR. Optimization of Monte Carlo trial moves for protein simulations. J Chem Phys. 2011;134(1):13.

70. Gordon JC, Myers JB, Folta T, Shoja V, Heath LS, Onufriev A. H++: a server for estimating pKas and adding missing hydrogens to macromolecules. Nucleic Acids Res. 2005;33:W368–W371.

71. Anandakrishnan R, Aguilar B, Onufriev AV. H++ 3.0: automating pK prediction and the preparation of biomolecular structures for atomistic molecular modeling and simulations. Nucleic Acids Res. 2012;40:W537–W541.

72. O'Boyle NM, Banck M, James CA, Morley C, Vandermeersch T, Hutchison GR. Open Babel: An open chemical toolbox. J Chemoinformatics. 2011;3(1):33.

73. Humphrey W, Dalke A, Schulten K. VMD: visual molecular dynamics. J Mol Graphics. 1996;14(1):33–38.

74. Press WH, Teukolsky SA, Vetterling WT, Flannery BP. In: Numerical Recipes: The Art of Scientific Computing, Third Ed. New York: Cambridge University Press; 2007. p. 868–883.

75. Janin J, Wodak S, Levitt M, Maigret B. Conformation of amino acid side-chains in proteins. J Mol Biol. 1978;125(3):357–386.

76. Marrink SJ, Risselada HJ, Yefimov S, Tieleman DP, de Vries AH. The MARTINI force field: Coarse grained model for biomolecular simulations. J Phys Chem B. 2007;111(27):7812–7824.

77. Monticelli L, Kandasamy SK, Periole X, Larson RG, Tieleman DP, Marrink SJ. The MARTINI Coarse-Grained Force Field: Extension to Proteins. J Chem Theory Comput. 2008;4(5):819–834.

78. Periole X, Cavalli M, Marrink SJ, Ceruso MA. Combining an Elastic Network With a Coarse-Grained Molecular Force Field: Structure, Dynamics, and Intermolecular Recognition. J Chem Theory Comput. 2009;5(9):2531–2543.

79. Chebaro Y, Pasquali S, Derreumaux P. The Coarse-Grained OPEP Force Field for Non-Amyloid and Amyloid Proteins. J Phys Chem B. 2012;116:8741–8752.

80. Liwo A, He Y, Scheraga HA. Coarse-grained force field: general folding theory. Phys Chem Chem Phys. 2011;13(38):16890–16901.

81. Spiriti J, Zuckerman DM. Tunable Coarse Graining for Monte Carlo Simulations of Proteins via Smoothed Energy tables: Direct and Exchange Simulations. J Chem Theory Comput. 2014;10:5161–77.

82. Bashford D, Case DA. Generalized born models of macromolecular solvation effects. Annu Rev Phys Chem. 2000;51:129–152.

83. Grant JA, Pickup BT, Sykes MJ, Kitchen CA, Nicholls A. A simple formula for dielectric polarisation energies: The Sheffield Solvation Model. Chem Phys Lett. 2007;441(1):163–166.

84. Kurut A, Fonseca R, Boomsma W. Driving Structural Transitions in Molecular Simulations Using the Nonequilibrium Candidate Monte Carlo. J Phys Chem B. 2018;122(3):1195–1204.

85. Andricioaei I, Straub JE, Voter AF. Smart darting Monte Carlo. J Chem Phys. 2001;114(16):6994–7000.

86. Minh DDL. Implicit ligand theory: Rigorous binding free energies and thermodynamic expectations from molecular docking. J Chem Phys. 2012;137(10).

87. Xie B, Nguyen TH, Minh DDL. Absolute Binding Free Energies between T4 Lysozyme and 141 Small Molecules: Calculations Based on Multiple Rigid Receptor Configurations. J Chem Theory Comput. 2017;13(6):2930–2944.

88. Chen W, Gilson MK, Webb SP, Potter MJ. Modeling Protein-Ligand Binding by Mining Minima. J Chem Theory Comput. 2010;6(11):3540–3557.

89. Shao JY, Tanner SW, Thompson N, Cheatham TE. Clustering molecular dynamics trajectories: 1. Characterizing the performance of different clustering algorithms. J Chem Theory Comput. 2007;3(6):2312–2334.