Rapid Flexible Docking Using a Stochastic Rotamer Library of Ligands

Feng Ding, Shuangye Yin, and Nikolay V. Dokholyan*

Department of Biochemistry and Biophysics, University of North Carolina at Chapel Hill, School of Medicine, Chapel Hill, North Carolina 27599

Received June 3, 2010

Existing flexible docking approaches model the ligand and receptor flexibility either separately or in a loosely coupled manner, which captures the conformational changes inefficiently. Here, we propose a flexible docking approach, MedusaDock, which models both ligand and receptor flexibility simultaneously with sets of discrete rotamers. We developed an algorithm to build the ligand rotamer library "on-the-fly" during docking simulations. MedusaDock benchmarks demonstrate a rapid sampling efficiency and high prediction accuracy in both self- (to the cocrystallized state) and cross-docking (to a state cocrystallized with a different ligand), the latter of which mimics the virtual screening procedure in computational drug discovery. We also perform a virtual screening test of four flexible kinase targets, including cyclin-dependent kinase 2, vascular endothelial growth factor receptor 2, HIV reverse transcriptase, and HIV protease. We find significant improvements of virtual screening enrichments when compared to rigid-receptor methods. The predictive power of MedusaDock in cross-docking and preliminary virtual-screening benchmarks highlights the importance to model both ligand and receptor flexibility simultaneously in computational docking.

INTRODUCTION

Specific interactions between small molecules and protein receptors are of crucial physiological and pharmacological importance. The ability to predict atomic interactions between ligand and receptor is extremely useful for understanding biological processes and for rational drug design. The major challenge in computational prediction of protein−ligand interactions is the large number of degrees of freedom, including protein backbone and side chain flexibilities, ligand conformational flexibility, and ligand rigid-body motion. Of particular interest is the receptor flexibility, which is essential to capture the receptor conformational changes upon ligand binding, i.e., the induced fit effect.[1−5] However, incorporating receptor flexibility is also computationally challenging and has been one of the foci of recent protein−ligand docking studies.[3−10] For example, the generation of an ensemble of multiple predetermined conformations has been proposed to model the receptor flexibility. The receptor conformation ensemble can be obtained experimentally by X-ray crystallography under different conditions and NMR spectroscopy[11−14] or computationally by molecular dynamics simulations,[6,15−18] comparative modeling,[19] and normal-mode analysis.[10,20] These ensemble approaches also have limitations. For example, multiple experimental conformations are not available for most proteins. Additionally, the predetermined receptor conformational ensemble may not capture receptor conformational changes upon binding to novel ligands. Such a limitation could be avoided if we model the flexibility of both ligand and receptor simultaneously.

Recently, several approaches have been proposed for the simultaneous sampling of the receptor and ligand flexibility during docking.[7−9,21−23] For example, protein side chain rotamer libraries, where the continuous protein side chain conformational space is modeled by a set of discrete states,[24] have been used to model protein flexibility during docking.[7,9,21−23] It has been shown that including receptor side chain rotamers of a few key residues for docking[9,22] or several highly probable rotamers of the binding pocking residues[7] can significantly improve the ability to find near-native poses in cross-docking studies when compared to the rigid-receptor docking. Specifically, Baker and co-workers[8,23] have recently developed the RosettaLigand method to extensively sample the receptor side chain conformations near the binding pocket and also found that the incorporation of receptor flexibility increases the probability of finding near-native poses with low energies. In RosettaLigand, the flexibility of a ligand is modeled by docking with a set of diverse ligand conformations,[8] which can be generated either by other methods[23] or by the recent extension of RosettaLigand.[25] However, during each RosettaLigand docking, the ligand conformation is predetermined, and the ligand conformational flexibility is sampled with limited dihedral angle perturbation, the search for the native pose is constrained by the availability of the native-like ligand conformations in the input ensemble of ligand conformations or rotamers.[26] Therefore, it is important to sample ligand conformations sufficiently during the flexible ligand and receptor docking.

To simultaneously and efficiently sample the receptor and ligand conformational flexibilities, we propose to model the ligand conformations with a set of discrete rotamers, in a similar way to protein side chains. Protein side chain rotamer libraries[24] can be built for 20 amino acids from high-resolution protein structures. However, there is not enough experimental data to build such a rotamer library for a large variety of ligands. Additionally, many small molecule ligands are very flexible, with multiple rotatable bonds, and it is practically impossible to enumerate all possible conformations. Therefore, instead of using a predetermined set of

* Corresponding author. E-mail: dokh@med.unc.edu.

10.1021/ci100218t © 2010 American Chemical Society
Published on Web 08/16/2010

ligand rotamers, we build the ligand rotamer library in a stochastic manner for each ligand during docking simulations. Using the stochastic rotamer library of ligands (STROLL), we are able to model the protein side chain[27] and the ligand flexibilities in a unified and simultaneous manner in MedusaDock. We use the recently developed MedusaScore[28] to guide the docking sampling and to rank the ligand poses. We benchmark MedusaDock with self- and cross-docking studies on a set of proteins cocrystallized with different ligands. We find that MedusaDock is able to sample near-native poses in all self-docking and in 95% of cross-docking cases. The near-native poses are top-ranked in 80 and 56% of self- and cross-docking cases, respectively, and are ranked within the top 10 in 95 and 72% of self- and cross-docking cases, respectively, which is among the best reported performances of flexible-ligand and -receptor docking algorithms. Interestingly, although the success rate in terms of placing the near-native poses as the best ranked is reduced for cross-docking as compared to self-docking, we find that the predicted binding energies—the MedusaScore of the best-ranked poses—between cross- and self-docking are close to each other. Assuming self-docking recapitulates the actual binding, which is often the case with near-native poses best ranked, we can approximate the binding energy with predicted values from cross-docking. This feature makes MedusaDock useful in virtual-screening (VS), where predicted binding energies from cross-docking are used to select the true binding ligands from decoys. Indeed, we find that compared to the rigid docking method, MedusaDock significantly improves the virtual-screening enrichment of cyclin-dependent kinase 2 (CDK2), vascular endothelial growth factor receptor 2 (VEGFr2), HIV reverse transcriptase (HIVRT), and HIV protease (HIVPR), four kinase targets which are known to be very flexible.[7,29,30]

RESULTS

Our docking method is comprised of STROLL generation, clustering of the generated ligand rotamers for initial coarse- and fine-docking for pose minimization (Methods). In order to evaluate whether the randomly generated ligand rotamer library sufficiently samples the conformational space, we test whether the STROLL library contains conformations similar to the X-ray crystallographic structure (Table 1). For each ligand, we generate 100 sets of STROLL using different random seeds. We align two conformations using the Kabsch[31] algorithm and compute the root-mean-square deviation (kRMSD) to determine the similarity between these two ligand conformations. We find that for all cases we are able to find a significant number of native-like rotamer (kRMSD < 2.0 Å and/or < 1.0 Å) conformations. Therefore, the stochastic rotamer library of ligands has sufficient sampling of the ligand conformational space.

Medusa uses an amino acid rotamer library to model protein side chain conformations.[27,32,33] In principle, we can model ligand conformational flexibility in the same way as amino acids using the ligand rotamers of STROLL. However, due to the additional translational and rotational degrees of freedom of the ligand, which are different from those in protein side chains, rigid-body minimization is necessary after each ligand rotamer change. If the number of available ligand rotamers is large, then direct modeling of the ligand rotamer change and the associated rigid-body motion is computationally prohibitive. To increase the ligand pose sampling efficiency, we propose a two-step docking approach, as illustrated in Figure 1. Briefly, each docking run starts with the generation of the STROLL, followed by clustering of the rotamers in STROLL with a kRMSD cutoff of 2.0 Å. For each of the N_C cluster centroids, we perform coarse docking to search for the best-fit poses within the docking boundary box. For the N_C coarse-docked poses, we sort and group similar poses and choose the top N_F (~10% N_C) poses for further fine docking. The coarse-docking step is designed to rapidly sample the rigid-body motion of the ligand for a set of representative ligand conformations. During coarse docking, the ligand rigid-body motion and the receptor side chain rotamers are iteratively sampled. In fine-docking, both the ligand and the receptor side chain rotamers are sampled simultaneously. At the end of a MedusaDock run, we have N_F minimized poses.

The computational time for the STROLL generation and the ligand rotamer clustering is usually smaller than 1 s CPU time on an Intel 2.33G Hz Xeon processor. Each coarse-docking run takes ~2 s, and each fine-docking run takes ~1

Table 1. Properties of the STROLL Library[a]

PDB ID	number of rotamers	P(1.0 Å)	P(2.0 Å)	average N_C
1A4Q	1000	0.023	0.35	8.5
1AQ1	7	0.72	0.71	3.0
1BMK	311	0.18	1.00	2.0
1C1C	540	0.09	0.78	2.4
1CX2	384	0.13	1.00	1.1
1DBJ	1	1.00	1.00	1.0
1DM2	2	0.50	1.00	1.0
1DI9	600	0.09	0.97	2.7
1DWC	1000	0.00	0.05	36.8
1DWD	1000	0.00	0.04	42.2
1ERR	923	0.04	0.45	8.2
1FM9	1000	0.00	0.07	84.1
1KI4	350	0.23	0.92	2.0
1KIM	36	0.42	0.97	2.0
1KSN	1000	0.02	0.24	14.6
1NSC	420	0.23	1.00	1.1
1P8D	27	0.15	0.78	3.0
1PMN	1000	0.12	0.76	3.2
1PMV	1	1.00	1.00	1.0
1PPC	1000	0.00	0.03	56.2
1PPH	1000	0.001	0.09	23.0
1PQ6	1000	0.00	0.03	109.0
1PQC	1000	0.002	0.25	7.6
1Q4L	274	0.06	0.59	5.8
1RTH	10	1.00	1.00	1.0
1STC	4	0.50	0.50	3.0
1XKA	1000	0.003	0.24	17.1
1YDS	420	0.15	0.74	2.6
2DBL	648	0.12	0.97	2.2
2PRG	787	0.004	0.41	13.5
3ERT	1000	0.007	0.50	6.6
3PGH	216	0.24	1.00	1.2
4TIM	53	0.00	1.00	1.0
6TIM	18	0.17	1.00	1.0

[a] Ligands are taken from the corresponding PDB files. For each ligand, 100 independent runs of STROLL generation are performed to compute the statistics. P(1.0 Å) and P(2.0 Å) correspond to the probability of finding ligand rotamers within 1.0 Å or 2.0 Å kRMSD from the native ligand conformation in the PDB. The average N_C corresponds to the average number of clusters using a kRMSD cutoff of 2 Å. Note that we limit the maximum number of rotamers to 1000 (Methods).

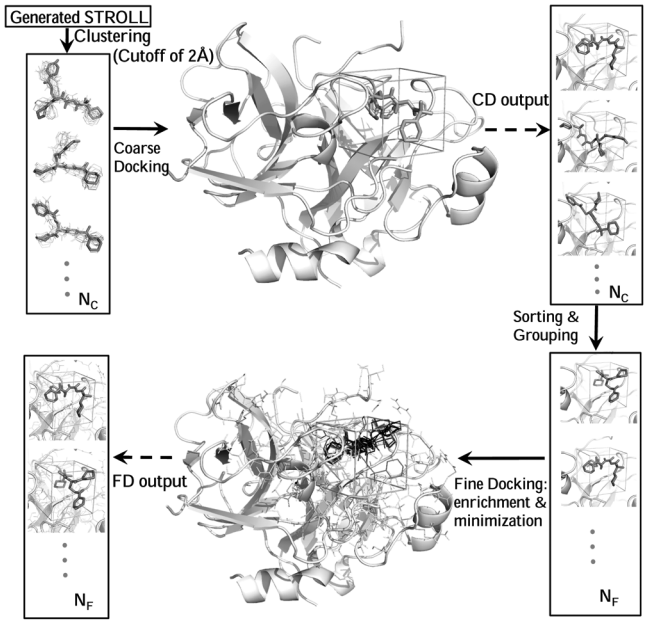

Figure 1. Schema for the MedusaDock protocol. Each docking simulation starts with the generation of the STROLL, followed by the clustering of ligand rotamers (Methods) with a kRMSD cutoff of 2 Å. The N_C cluster centroids are used as the representative ligand conformations for the initial coarse-docking, where the ligand is kept rigid and the ligand binding pocket is rapidly sampled (Methods). Since only a small number of the ligand rotamers is used and the ligand is kept rigid during coarse-docking, the near-native pose is not necessarily the lowest energy pose but usually is among the low-energy poses. We sort the N_C poses according to binding energies. We group similar poses together if the RMSD of two ligand poses is smaller than their kRMSD plus 2 Å. The grouping of similar poses will reduce the necessary number of the more expensive calculations of fine docking. We select the top N_F (~10% N_C) groups of lowest energy poses for the next round of fine docking. For each group of coarse-docked poses, we perform fine docking to enrich the ligand rotamers and to minimize the total energy (Methods). During fine docking, we add the ligand rotamers within 2.0 Å kRMSD to the initial centroid rotamer in order to enrich the ligand conformation. For each of the members in a group, we perform the enrichment step. We only select the lowest energy poses for further minimization (Methods). Therefore, we will have N_F poses for each MedusaDock run.

min. Therefore, for a ligand with 30 clusters, the computational time for each MedusaDock run is approximately 4−5 min.

Self- and Cross-Docking Benchmarks. We compile a set of 18 pairs of receptor−ligand structures from literature,[8,10,34] where each case has two different receptor conformations cocrystallized with different ligands. These proteins with their PDB[35] codes are listed in Table 2. For each docking, we perform 100 independent runs of MedusaDock simulations and rank the poses. In principle, the poses should be ranked according to the total binding energy. However, in some cases a low binding energy pose can result from unfavorable receptor side chain conformations. Therefore, it is important to take the energy of the whole system into account when searching for the near-native ligand pose (Methods).

We perform self-docking simulations, where each ligand is docked to the corresponding cocrystallized receptor conformation state. To evaluate the predictions, we refer the RMSD of the ligands in the two ligand−receptor poses as RMSD, where only the receptors are aligned. For all cases, MedusaDock is able to find near-native ligand poses with RMSD smaller than 2 Å from the native state. In 28 out of 35 cases, the near-native poses are top ranked, and in 31 out of 35, the near-native poses are ranked in the top 2. Only in 2 cases the near-native poses are ranked below the top 10. The results are summarized in Table 2. In Figure 2, we present two examples where the near-native poses are correctly predicted as best-ranked, one is PPARγ (PDB code: 2PRG; Figure 2A,B,C), and the other is LXR β ligand binding domain (PDB code: 1PQC; Figure 2D,E,F). For 2PRG, the MedusaScore alone cannot distinguish the near-native poses from the decoys (Figure 2B). For example, the pose with the lowest MedusaScore has an RMSD larger than 10 Å. Using the ranking energy (Methods), we are able to separate near-native poses from decoys (Figure 2C). In the case of 1PQC, the MedusaScore is able to distinguish the near-native pose from the decoys, although there is one decoy with a MedusaScore very similar to the lowest MedusaScore (Figure 2E). By ranking the poses using the ranking energy, the near-native poses can be unambiguously selected. There are two cases, estrogen receptor (PDB code: 3ERT) and P38 MAP kinase (PDB code: 1BMK), where the near-native poses are ranked below the top 10 (Table 2 and Figure 2G,H,I). For 3ERT, we find that the best-ranked ligand pose and the native pose are partially symmetric, with most of the heavy atoms overlapping (Figure 2G). In the case of 1BMK, the pose with the higher number of contacts to the receptor is selected as the best-ranked pose over the native pose (Figure 2H,I), which is probably due to the inaccuracy of the force field. Therefore, although MedusaDock includes a large number of degrees of freedom in the modeling, the ability to find near-native poses for all self-docking cases and correctly rank the near-native pose as the top one in 80% of cases and the top two in ~90% of cases highlights the docking efficiency and accuracy of MedusaDock.

In cross-docking simulations, the ligand is docked to a receptor conformational state cocrystallized with a different ligand. For the 18 pairs of ligand−receptor complexes, we dock each of the ligands to the receptor structure cocrystallized with the other ligand. In all 36 except two cases, LXR β ligand binding domain (1PQC ligand docked to 1PQ6 receptor) and JNK3 (1PMN ligand docked to 1PMV receptor), the near-native ligand poses are sampled by MedusaDock. In the case of 1PMN ligand docked to 1PMV receptor, the backbone of the 1PMV receptor has severe clashes with the native ligand pose of 1PMN (Figure 3A), which prevents the sampling of the near-native pose. In the case of 1PQC ligand docked to 1PQ6 receptor, the lowest RMSD pose has a RMSD of only 2.21 Å (Figure 3B). In 20 out of 36 cases, the near-native poses are identified as the best-ranked poses, 24 out of 36 cases within the first 5, and 26 out of 36 cases within the first 10. The success rate in terms of ranking the near-native poses as best-ranked is reduced for cross-docking as compared to that of self-docking, as is commonly observed with other docking approaches.[9,10,23] Due to the stochastic nature of the docking approach, it is usually not straightforward to compare the ranking of the near-native poses of specific targets between different methods. For example, there are a few cross-docking cases where the ranking of the near-native poses are not as good as those by other methods, such as 1FM9 versus 2PRG by RosettaLigand[23] or 1BMK versus 1DI9 by ICM.[10] In some

Table 2. Summary for Self- And Cross-Docking Benchmarks[a]

protein	PDB	pairwise RMSD (Å)	self-docking		cross-docking	
			BRP RMSD (Å)	Best NN Rank	BRP RMSD (Å)	Best NN Rank
CDK2	1AQ1	0.39	0.50	1	0.64	1
	1DM2		0.28	1	*4.42*	95
antibody	1DBJ	0.47	0.27	1	0.96	1
	2DBL		1.01	1	0.91	1
thrombin	1DWC	0.13	1.99	1	1.44	1
	1DWD		1.32	1	*2.97*	16
PPARγ	1FM9	0.49	*8.67*	2	*10.58*	11
	2PRG		*1.32*	1	*10.02*	64
LXR β LBD	1P8D	0.65	*4.92*	2	1.99	1
	1PQ6		1.88	1	0.85	1
LXR β LBD	1PQ6	0.62	1.88	1	*3.37*	–
	1PQC		1.27	1	1.53	1
trypsin	1PPC	0.11	1.84	1	4.73	5
	1PPH		1.52	1	1.94	1
isomerase	4TIM	0.21	1.43	1	1.42	1
	6TIM		1.54	1	1.52	1
COX-2	1CX2	0.37	*6.07*	5	1.46	1
	3PGH		1.99	1	0.82	1
estrogen receptor	1ERR	0.34	0.66	1	*7.31*	5
	3ERT		*7.43*	19	0.52	1
factor Xa	1KSN	0.34	0.95	1	1.65	1
	1XKA		1.31	1	1.21	1
GSK-3 β	1Q4L	0.57	*2.59*	2	1.11	1
	1UV5		0.37	1	*2.80*	297
HIV-1 RT	1C1C	0.81	0.86	1	*6.07*	5
	1RTH		0.58	1	1.41	1
JNK 3	1PMN	0.78	0.68	1	*4.68*	10
	1PMV		0.29	1	*7.00*	–
neuramini-dase	1A4Q	0.21	1.60	1	1.06	1
	1NSC		0.96	1	1.57	1
P38 MAP kinase	1BMK	0.49	*8.32*	66	*9.78*	78
	1DI9		*5.34*	10	*2.40*	6
PKA	1STC	0.72	0.43	1	*2.84*	5
	1YDS		1.23	1	*2.29*	147
thymidine kinase	1KI4	0.25	0.79	1	0.61	1
	1KIM		0.54	1	*6.12*	231

[a] We compute the RMSD for the best-ranked pose (BRP) and also list the best rank of the near-native (NN) poses (within 2 Å RMSD from the native state). The solid line (–) corresponds to cases where the near-native poses are not sampled in MedusaDock simulations.

other cases such as 1P8D versus 1PQ6, MedusaDock has better predictions than the others. Overall, the percentages of near-native poses top-ranked and generated for MedusaDock are 56 (20/36) and 94% (34/36), respectively. The same percentages for RosettaLigand[23] are 50 (10/20) and 90% (18/20), and for ICM[10] are 57 (16/28) and 89% (25/28). Here, the bracket indicates the number of cases that have been studied. Therefore, the performance of cross-docking by MedusaDock is comparable to or slightly better than the performances of other flexible docking programs.

One important application of a docking program is virtual screening, where a library of ligands is docked to the apo- or holo-structure of a target receptor in order to find binding ligands, a process mimicked by the cross-docking exercise. The key for the success of virtual screening is to accurately predict the binding energy for a ligand. Although the success rate of cross-docking in terms of identifying the near-native poses as best ranked is not as high as that of self-docking, a docking program will still be useful in virtual-screening if the predicted binding energy of cross-docking—the MedusaScore of the best-ranked pose—is close to the value of that for self-docking. Here, we assume that the predicted binding energy of self-docking is sufficiently accurate to approximate the actual binding affinity,[28] as in most cases of self-docking the near-native poses are best ranked. We plot the binding energy of cross-docking versus that of self-docking (Figure 3C) and find that the two sets of binding energies are very close to each other, with a correlation coefficient of 0.98 and a regression slope of 0.99. Why is the predicted binding energy of cross-docking close to that of self-docking even when the near-native pose is not best ranked in cross-docking? Since MedusaScore is a physically based scoring function, near-native poses in both cross- and self-docking will feature similar physical interactions, and in turn, similar MedusaScores. The above question can be answered if the MedusaScore of the best-ranked pose is close to that of the best-ranked near-native pose. Indeed, we find that the MedusaScore of the best-ranked poses in cross-docking is similar to that of the corresponding best-ranked native-like poses (Figure 3D). For the cases where the best-ranked pose is near-native, the two values will be the same, but for cases where the near-native poses are not top ranked, their energies are also close. Therefore, the benchmark of self- and cross-docking suggests that MedusaDock might be useful in virtual screening. Next, we perform a preliminary virtual-screening test of MedusaDock on a set of flexible kinase targets, including CDK2, VEGFr2, HIVRT, and HIVPR.

Virtual Screening. We test MedusaDock on a VS benchmark set taken from the Directory of Useful Decoys

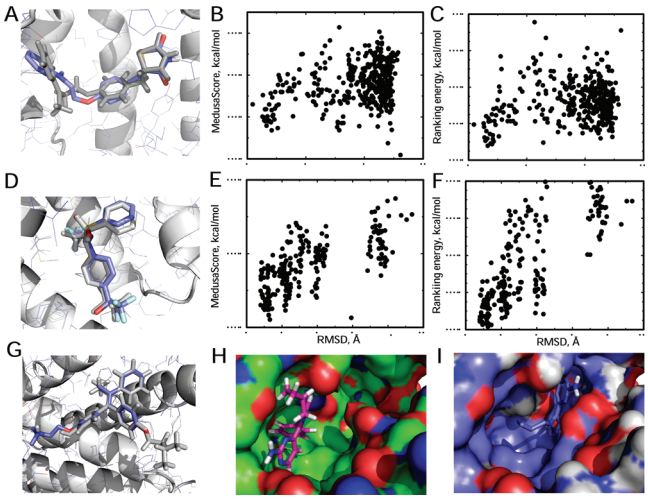

Figure 2. Self-docking results. (A) The best-ranked pose for PPARγ (PDB code: 2PRG) is presented, where the native and predicted ligand poses are in gray and in color, respectively. The MedusaScore (B) and ranking energy (C) are plotted against the corresponding RMSD from the native pose. The MedusaScore alone cannot distinguish the near-native poses from the decoys. For example, the pose with the lowest MedusaScore has an RMSD larger than 10 Å. Using the ranking energy (Methods), we are able to separate near-native poses from decoys. For the LXR β ligand binding domain (PDB code: 1PQC), the predicted pose (D), the MedusaScore (E), and the ranking energy (F) are shown. By ranking the poses using the ranking energy, the near-native poses can be unambiguously selected. (G) The best-ranked (colored) and native (gray) poses of estrogen receptor (PDB code: 3ERT) are partially symmetric. For P38 MAP kinase (PDB code: 1BMK), the native ligand (H) and predicted best-ranked (I) poses show distinct interactions between ligand and receptor.

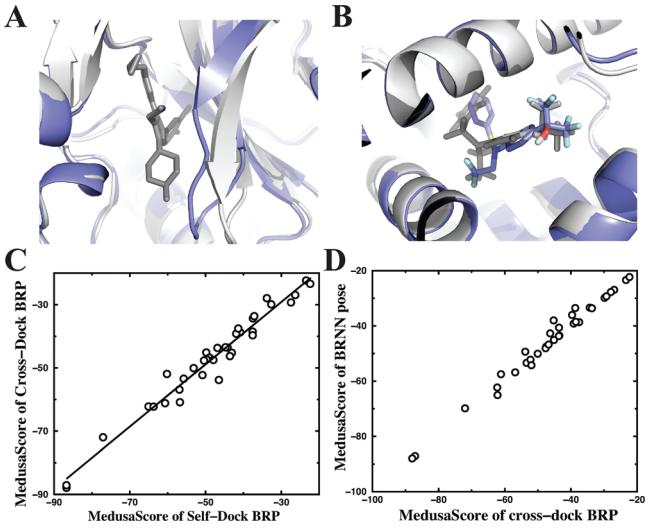

Figure 3. Cross-docking results. For 1PMN ligand docked to 1PMV receptor, the backbone has severe clashes with the native pose (A). For 1PQC ligand docked to 1PQ6 receptor (B), the lowest-RMSD pose (colored) has a RMSD of only 2.2 Å. (C) The MedusaScore of the best-ranked pose (BRP) in crossing-docking is compared to that of self-docking. (D) In cross-docking, the MedusaScore of the best-ranked near-native pose (BRNN) is also close to that of the best-ranked poses.

(DUD) test set.[36] For each target in DUD, there is a set of known binding ligands and corresponding decoy ligands with similar physiochemical properties, making it challenging for VS study. Similar to the original study, we measure the performance of VS by an enrichment plot, which shows the percentage of known binders (true positives) recovered as a function of the percentage of the total library screened. This ratio is the enrichment factor (EF), which is expected to be 1 if the compounds are ranked randomly and more than 1 if a virtual-screening method can rank the compounds so that there are more true binders in the top-ranking compounds.

Among the 40 DUD targets, we select CDK2, VEGFr2, HIVRT, and HIVPR for testing because they are known to be flexible, which possibly leads to the poor enrichment found in the original VS experiments using a rigid-receptor docking protocol.[36] In the original study, the EF_1 and EF_{10} (EF at 1% and 10% library screened) are only 8.0 and 2.4 for CDK2, 3.0 and 1.1 for VEGFr2, 4.9 and 2.5 for HIVRT, and 1.7 and 0.7 for HIVPR.

We apply MedusaDock for VS and find significant enrichments for all four targets (Figure 4). The EF_1 and EF_{10} obtained from MedusaDock are 14 and 3.6 for CDK2, 31 and 4.2 for VEGFr2, 7.5 and 3.0 for HIVRT, and 15.1 and 3.2 for HIVPR. All results are significantly higher than those obtained using Dock3.5.[36] We contribute the improvements to the inclusion of receptor flexibility using MedusaDock. More interestingly, we find that the true positive rates are high among the top 10 ligands for all the 4 targets. For VEGFr2, all the 15 top-ranked ligands are known VEGFr2 binding ligands. For the other 3 targets, there are 7, 6, and 3 true positives among the top 10 ligands for CDK2, HIVPR, and HIVRT, respectively. Such a low false-positive rate makes MedusaDock appropriate for VS because, in reality, only a limited number of ligands can be tested experimentally, constrained by the time and expense of the biological essays.

DISCUSSION

We adopt a rotameric approach to model the side chain conformational flexibility of proteins. Due to the large physiochemical space of ligands and also the large number of degrees of conformational freedom, it is impossible to build an enumerative rotamer library in the same way as for proteins.[24] Therefore, we construct a rotamer library of ligands in a stochastic manner during each MedusaDock simulation. OMEGA by OpenEye Scientific Software (http://www.eyesopen.com/) is often used to precompute the ligand conformations. Although OMEGA uses a different ligand conformation generating method, the depth-first algorithm, we find that our simple method features similar performance in terms of finding the bioactive rotamers (kRMSD < 1 Å) or the computational time (~1 s for each ligand).[37] Although the near-native ligand conformations with kRMSD smaller than 2.0 Å, or even 1.0 Å, can consistently be sampled by STROLL (Table 2), the native ligand rotamer will not always be included in the rotamer library due to the stochastic nature of library generation, which in turn might affect the efficiency of finding and identifying the near-native binding pose. To test the efficiency and accuracy of MedusaDock, we compare the self- and cross-docking benchmark results with and without manually including the native ligand rotamers in STROLL (Figure 5). As expected, manual inclusion of the native ligand rotamer allows MedusaDock to sample the near-native poses more efficiently, with more near-native poses sampled (Figure 5A). However, we find that the inclusion of the native poses does not significantly affect the docking prediction accuracy in terms of ranking the near-

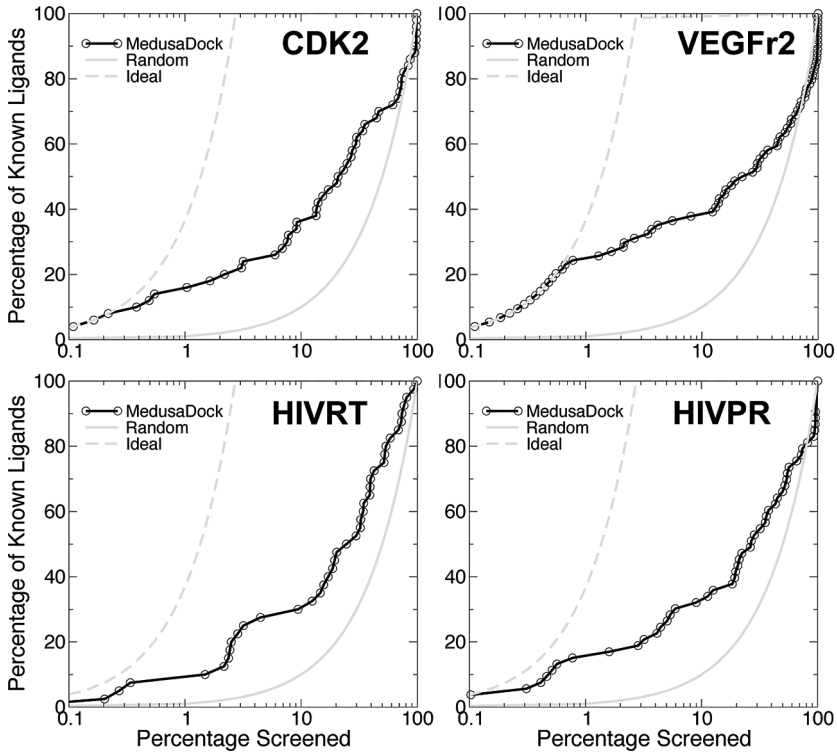

Figure 4. Virtual-screening test on a DUD benchmark set of CDK2, VEGFr2, HIVRT, and HIVPR. The percentage of recovered known binding ligands is plotted as a function of percentage of library screened using various VS strategies. For reference, we also plot the enrichment curves for ideal VS, where all ligands are ranked ahead of decoys, and for random VS, where all molecules are ranked randomly. For VEGFr2, we find all the 15 top-raked ligands are known VEGFr2 binders (ideal VS performance). For the other targets, there are 7, 6, and 3 known binders among the top 10 ligands for CDK2, HIVPR, and HIVRT, respectively.

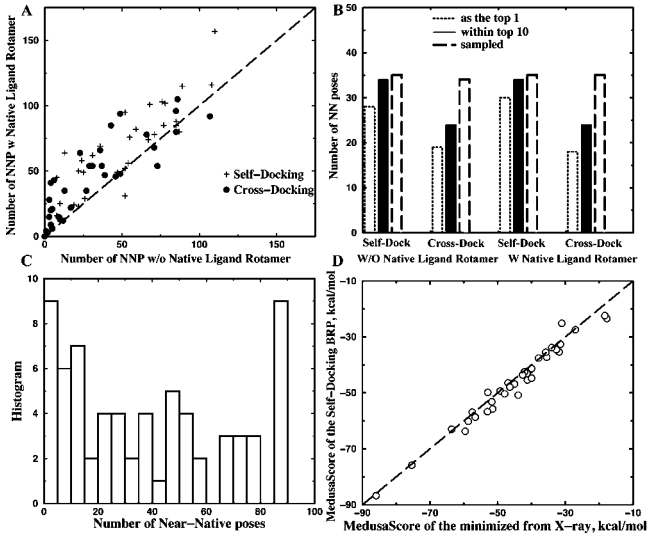

Figure 5. Effect of inclusion of the native ligand rotamer in docking calculations. (A) Inclusion of the native rotamer in STROLL increases the number of near-native poses sampled by MeduasDock. (B) The prediction accuracy in terms of ranking near-native poses as the top one and 10 and sampled does not depend on whether the native ligand rotamers are included. (C) The histogram of the number of near-native poses sampled. (D) The MedusaScore of the best-ranked poses from self-docking simulations is close to that of the minimized poses around the input X-ray crystallographic structure.

native poses (Figure 5B). This observation suggests that our ligand sampling by STROLL is sufficient for accurate docking.

To increase the computational efficiency, we devise a two-step docking protocol, including the initial coarse-docking with representative ligand rotamers followed by fine-docking to enrich the ligand rotamers and to minimize the docking poses. Clustering of the ligand rotamers in STROLL helps to group similar rotamers together, and coarse-docking using the corresponding representative conformations helps to avoid the repetitive calculation of rigid-body motions of similar ligand conformations. Since fine-docking only perturbs translation and rotation around the input coarse-docked pose, it is crucial for the coarse-docking procedure to place the ligand in the proximity of the native position and orientation. As pointed out by previous flexible docking studies,[23,26] one of the challenges of docking is to fit the ligand simultaneously to multiple deep pockets, i.e., a rugged energy landscape (Figure 6A). We smooth the energy landscape by turning off the van der Waals repulsion between the ligand and the receptor side chains (Figures 6B and 7). Therefore, in the coarse-docking step, we adopt an iterative rigid-body docking and receptor side chain packing approach with slowly increasing van der Waals (vdW) repulsion between the ligand and the side chains to facilitate the search of the rugged energy landscape (Methods). During fine docking, the ligand and receptor side chain rotamers are simultaneously modeled in a strongly coupled manner. The self- and cross-docking benchmark of MedusaDock suggests that the proposed docking protocol is quite efficient in the sampling of the near-native poses (Figure 5C); in many cases (~80%), the program identifies more than 10 near-native poses from 100 runs of MedusaDock simulations.

MedusaDock uses MedusaScore to guide docking and to rank the docking poses. MedusaScore uses a physical-based force field to describe the physical interactions between the ligand and the receptor. By correctly docking the ligand to

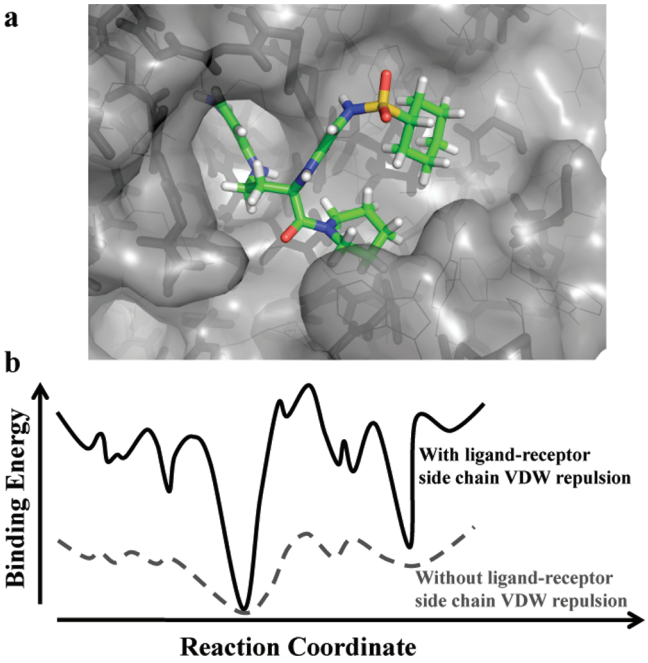

Figure 6. The rugged energy landscape of docking. (A) The ligand binding surface (PDB ID: 1UVS) features several deep subpockets. Successful prediction requires simultaneous fitting of the ligand to all subpockets. (B) The schematic energy landscape with and without vdW repulsion between the ligand and the receptor side chains. By turning off the vdW repulsion between ligand and the receptor side chains, the rugged energy landscape is smoothed.

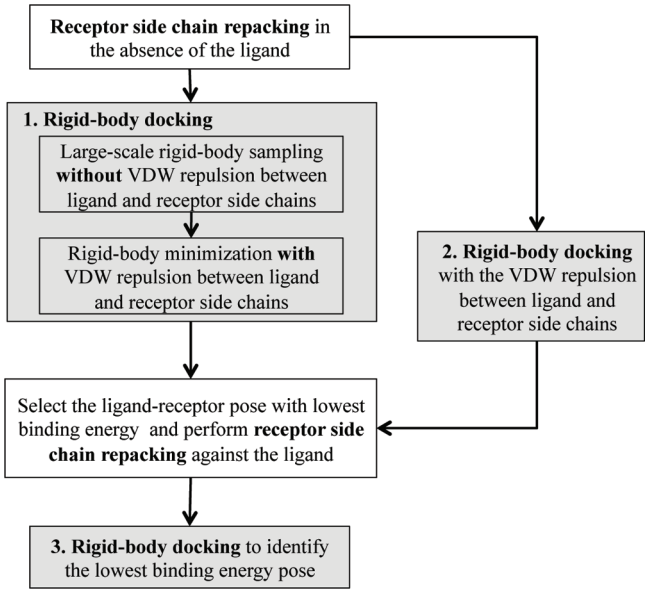

Figure 7. Flowchart of coarse-docking. During coarse-docking, the ligand is kept fixed, with only rigid-body motion allowed. Iterative receptor side chain repacking and rigid-body docking are performed to identify the lowest binding energy pose.

the receptor, MedusaDock can capture the same important binding interactions as in the native pose. For example, the MedusaScore of poses minimized from the X-ray crystallographic structures is close to that of the best-ranked self-docking poses, which are mostly near-native (Figure 5D). Despite the fact that in many cases the near-native poses are not ranked as the top poses in cross-docking, the ranks of these near-native poses are quite high, mostly within the top 10. Hence, the MedusaScore of the best-ranked pose in cross-docking—the predicted binding energy—is similar to that of the best-ranked near-native pose, featuring the same binding interactions as self-docking (Figure 3). Therefore, the MedusaScore from cross-docking simulations can be reliably used to estimate the binding energy for virtual-screening, which has been validated by our preliminary virtual-screening tests. Further testing of the application of MedusaDock in virtual-screen by benchmarking on the whole DUD data set and by comparing with other docking programs is required in a future study.

In this work, we do not model the receptor backbone flexibility. As a result, there are two cases in cross-docking where the program fails to identify the near-native poses due to severe backbone clashes near the ligand binding pocket (Table 2 and Figure 3A). Several approaches have been used to model protein backbone flexibility, including ensemble docking with multiple backbone conformations[10] and backbone relaxation.[7,23] Interestingly, our method is able to capture the near-native poses of several challenging cases that other flexible receptor-backbone methods missed, such as thrombin (1dwc, 1dwd pair; Table 1 in ref 26) and PPARγ (1FM9, 2PRG pair; Table 3 in ref 10). We believe that the ability of MedusaDock to find near-native poses is a result of the efficient modeling of ligand conformations and the protein side chains using rotameric approaches. However, we believe that further modeling of the receptor backbone flexibility will help improve the predictive power.[38]

Large-scale virtual screening requires high computational efficiency, since a large library with millions of ligands will be docked to a target protein. The docking program can be highly parallelized since the calculations are independent of each other. Additionally, since the docking program is able to sample many near-native poses with high efficiency, we can reduce the total number of MedusaDock runs in a future virtual-screening study. The preliminary virtual-screening benchmark results for CDK2, VEGFr2, HIVRT, and HIVPR clearly demonstrate the capabilities of MedusaDock in large-scale virtual screening, even for flexible protein targets.

METHODS

Ligand Conformational Flexibility and STROLL Generation. We model the conformational flexibility of a small molecule ligand by allowing each rotatable bond to rotate freely and the corresponding dihedral angle to adopt preferred values according to the hybridization of the two atoms forming the bond. For each rotatable bond, we use the a similar approach as Miller and Baker[8] to define the torsional degree of freedom. Briefly, if the two atoms are of sp^3 hybridization, we assign a three-fold symmetry to the dihedral angle. If the two atoms are of sp^2 hybridization, we assume a two-fold symmetry. Otherwise, if one atom is of sp^2 and the other is of sp^3, we will adopt a 12-fold symmetry for the dihedral angle. In some cases of sp^2-sp^2 hybridizations where there are persistent clashes, we extend the 2-fold symmetry to 12-fold. In principle, we are able to enumerate all possible conformations of the ligand. However, the total number of rotamers of a ligand can be large. For example, if a ligand has 10 rotatable bonds and the average number of degrees of symmetry is 3, the total possible number of rotamers will be $3^{10} = 59\,049$. Sampling the large conformational space of the ligand along with its rigid-body motion during docking is too computationally expensive. We propose

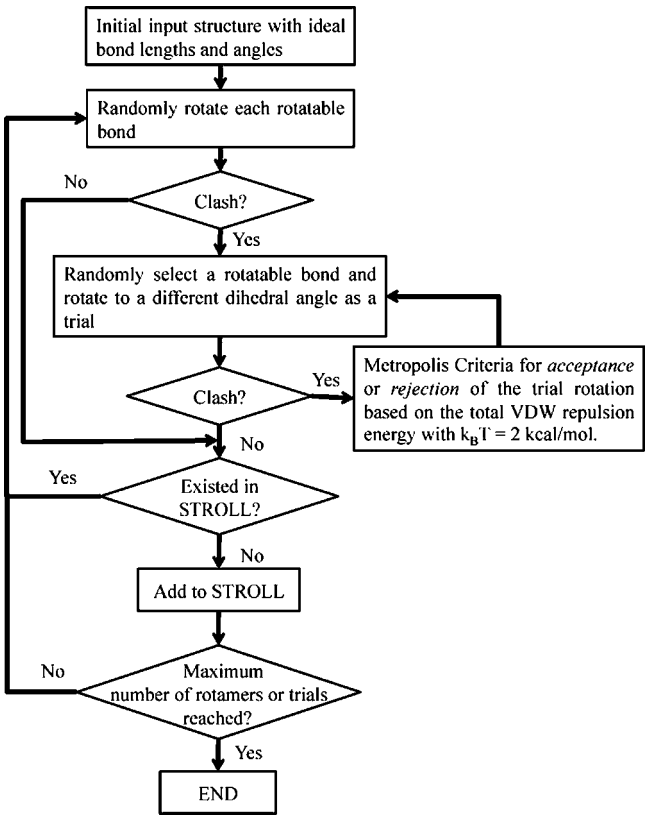

Figure 8. Flowchart of STROLL generation. The input ligand conformation has ideal bond lengths and angles, and only dihedral angles of rotatable bonds are changed during the procedure. To evaluate whether a ligand conformation has clashes, we only compute the vdW repulsions between atom pairs whose pairwise distances are determined by at least one rotatable bond. Two atoms are denoted as nonlocal if their pairwise distance is determined by more than one rotatable bond and as local if their distance is governed by only one rotatable bond. For each local atom pair, we determine the minimal vdW repulsion energy by enumerating the corresponding dihedral angle. We consider a nonlocal atom pair clashing if the vdW repulsion energy is larger than 0.6 kcal/mol and a local atom pair clashing if the vdW repulsion energy is at least 2.4 kcal/mol more than its minimum vdW repulsion energy. A carbon−carbon pair (both local and nonlocal) is considered clashing if their distance is bellow 3 Å. A ligand conformation is determined as clashing if one or more atom pairs are clashing. For ligands with small degrees of freedom, we enumerate all possible nonclashing conformations. For ligands with large degrees of freedom, we stop the rotamer generation once, either the maximum number of rotamers (set to 1000) is generated or the total number of trials (set to 10^6) is reached.

a stochastic rotamer library of ligands to model the ligand flexibility, where we generate a set of ligand rotamers by randomly generating nonclashing ligand conformations using a Monte Carlo-based algorithm (Figure 8). To evaluate whether a ligand conformation has clashes, we compute the vdW repulsions between atom pairs whose pairwise distances are determined by at least one rotatable bond. Two atoms are denoted as nonlocal if their pairwise distance is determined by more than one rotatable bond and as local if their distance is governed by only one rotatable bond. For each local atom pair, we determine the minimal vdW repulsion energy by enumerating the corresponding dihedral angle. We consider a nonlocal atom pair clashing if the vdW repulsion energy is larger than 0.6 kcal/mol and a local atom pair clashing if the vdW repulsion energy is at least 2.4 kcal/mol more than its minimum vdW repulsion energy. A carbon−carbon pair (both local and nonlocal) is considered clashing only if their distance is bellow 3 Å. A ligand conformation is determined as clashing if one or more local or nonlocal atom pairs are clashing. For ligands with small degrees of freedom, we enumerate all possible nonclashing conformations. For ligands with large degrees of freedom, we will stop the rotamer generation once, either the maximum number of rotamers (set to 1000) is generated or the total number of trials (set to 10^6) is reached.

A Simple Clustering Algorithm to Group Similar Ligand Rotamers in STROLL. We iteratively group similar rotamers in STROLL. Initially, the first rotamer is selected from STROLL to be the only element of the first cluster and also the cluster representative. For each newly selected rotamer, we will first decide whether it belongs to any previously constructed cluster by comparing it with the cluster representative. If the smallest kRMSD is less than an input cutoff value, we assign the rotamer to the corresponding cluster. Otherwise, we assign this rotamer to a new cluster as the cluster representative. This iterative procedure ends once all rotamers are assigned. This clustering algorithm does not require the kRMSD calculation between all pairs of rotamers in STROLL as does the commonly used hierarchical clustering methods.

Coarse Docking. During coarse docking, we perform multiple rounds of Monte Carlo (MC)-based rigid-body docking with varying receptor side chain packing (Figure 7). During rigid-body docking, the receptor conformation is fixed, and the ligand rigid-body motion is sampled in two stages. First, the ligand is randomly rotated and translated with the center of mass confined inside a $10 \times 10 \times 10$ Å^3 cubic box around the input pocket center. The docking boundary can be arbitrarily defined. A Metropolis criterion is used to decide whether a new ligand rotamer is accepted or rejected. We usually perform sampling at a high temperature, such as 10 kcal/mol$\cdot k_B$, to avoid trapping in local minima of the energy landscape. After 600 MC steps, we select the lowest energy pose as well as poses with dwell time more than 20 MC steps as the candidates for the second stage of rigid-body minimization. During the minimization, the ligand is randomly perturbed by translation and rotation in small steps with Gaussian distributions. The average length and angle steps are 0.2 Å and 2°, respectively. The maximum minimization step is 100, and the temperature is 0.25 kcal/mol$\cdot k_B$. The pose with the lowest binding energy will be selected for further study.

In order to sample the rugged energy landscape, we devise a three-step coarse-docking protocol using multiple rounds of rigid-body docking with varying receptor side chain packing (Figure 7):

(1) We repack the protein side chains in the absence of the ligand.[27] We exclude residues with any atoms 10 Å away from the docking boundary. The side chain repacking is done by a three-round simulated annealing with temperatures of 10, 3, and 2 kcal/mol$\cdot k_B$, respectively. For each round of repacking, the total number of rotamer trials is twice the number of total available side chain rotamers. Since we only repack the side chains around the ligand binding pocket, the packing converges rapidly. Then, we perform the rigid-body docking to sample the ligand binding pocket defined by the receptor backbone. We turn off the vdW repulsion between the ligand and

the receptor side chains during the large-scale random translation and rotation in order to smooth the energy landscape (see Discussion and Figure 6). For the minimization step, involving small perturbations of translation and rotation, we turn on the vdW repulsion between the ligand and the receptor side chain. The pose with lowest binding energy is selected.

(2) Next, we perform a new round of MC rigid-body docking by turning on the vdW repulsion between the ligand and the receptor side chains. If the binding energy identified during this step is lower than that in the previous step, then we will take the identified ligand pose for the next step. Otherwise, the ligand pose from step 1 will be used.

(3) Last, we repack the protein side chains against the identified pose in steps 1 and 2. We only perform one round of side chain repacking with a MC temperature of 2 kcal/mol$\cdot k_B$. Then, we perform the third round of MC rigid-body docking and identify the lowest binding energy pose as the result of the coarse-docking.

Fine Docking. The purpose of fine-docking is to sample ligand−receptor conformations in the vicinity of the input pose and minimize the total binding energy. The fine-docking is composed of the following two steps:

(1) Ligand Conformation Enrichment: Based on the input ligand rotamer conformation, we select a subset of rotamers from STROLL with kRMSD smaller than 2.0 Å. If the number of rotamers is smaller than 200, then we will regenerate similar rotamers with kRMSD < 2.0 Å until either enough rotamers are generated or the maximum rotamer generation trial is reached. Then, we perform simulated annealing simulations, where receptor side chains and ligand rotamers are randomly sampled and 10 rounds of rotamer searches are performed with the initial and final temperatures of 10.0 and 0.1 kcal/mol$\cdot k_B$, respectively. We do *not* perform ligand rigid-body minimization during the ligand rotamer changes. As the simulation temperature decreases, the ligand rotamer acceptance rate decreases. Once the ligand rotamer can no longer be changed favorably, we perform MC-based rigid-body minimization with small rotation and translation perturbations for each round of rotamer searches.

(2) Pose Minimization: We minimize the ligand poses by simulated annealing with protein side-chain packing, ligand subrotamer vibration, and rigid-body minimization. We sample the ligand subrotamer search by rotating dihedral angles within the allowed dihedral angle variation. We use an angle variation of ±15° for sp^3-sp^3 bonds and ±10° for the remaining bond types. At the end of the simulated annealing, a quenching procedure is applied to find the local energy minima. During the quench, the conjugated gradient method is used for the protein side chains minimization, and a MC-based rigid-body minimization at zero temperature is used for ligand rigid-body minimization.

Scoring Function. We use an extended MedusaScore[28] as the scoring function to dock and rank the ligand poses. The addition to the original MedusaScore is the inclusion of electrostatic interactions. We assign integer charges to various charged chemical moieties and use a distance-dependent dielectric constant, $\sim r$, to model the screening effect. We do not use any distance cutoff. To model the environmental dependence of the electrostatic interactions, we rescale the interaction potential by the extent of solvent exclusion (burriness) of the corresponding residue i, B_i.[39] We define the burriness of a residue based on the number of its contacts N, $B = \min[N/N\max, 1]$. Here, N is the number of contacts computed based on the C_β atoms within a distance cutoff 8.5 Å, and $N\max = 14.5$ is maximal number of contacts. For two charges of the receptors, the coefficient is $(B_i + B_j)/2$. If the interaction is between the ligand and the receptor, then the coefficient is B. By introducing the burriness-dependent factor B, we ensure that the buried charges will have stronger electrostatic interactions than the solvent exposed ones.

The exclusion of the vdW repulsion between the ligand and the receptor in MedusaScore fits better to the experimental binding affinity measurements.[28] Therefore, we include the ligand−receptor vdW repulsion for docking simulations, but we use MedusaScore without the ligand−receptor vdW repulsion as the binding energy between ligand and receptor.

Ligand Pose Ranking. We include the total energy of the ligand−receptor complex in addition to the binding energy in the ranking of ligand poses, $E_{rank} = E_{bind} + c\Theta(E_{total} - E_{cutoff})$. Here, the correction function $\Theta(x)$ is equal to x if $x > 0$ and 0 otherwise, E_{cutoff} is chosen so that the top 5% of poses with the lowest total energy have no correction to the binding energy, and c is the weighting coefficient. The cutoff energy is introduced to account for the large fluctuations in total energy.[23] For our benchmark cases, we find that optimal prediction is not sensitive to the value of the coefficient c. For example, in our test cases, we find that the coefficient c can range approximately from 0.2 to 4 to give the optimal prediction rate. For simplicity, we set $c = 1$.

Virtual Screening. We use a ligand binding crystal structures for DUD. The cognate ligand is removed, and its center is used as the pocket center for restraining the ligand position during docking. For each molecule, we perform 200 runs of MedusaDock simulations, which generate 712 docking poses on average. We then compute the ranking energy based on the binding and total energies for each pose, as described above. Finally, all molecules are ranked according to their lowest E_{rank} of all docking poses.

ACKNOWLEDGMENT

We thank Dr. Adrian Serohijos for insightful discussions and Elizabeth A. Proctor for critical reading of the manuscript. The work is supported by the National Institute of Health Grant R01GM080742 and the ARRA supplement 3R01GM080742-03S1 (to N.V.D.) and by the UNC Research Council (to F.D.). The calculation is performed on the topsail high-performance computing cluster at University of North Carolina.

REFERENCES AND NOTES

(1) Leach, A. R.; Shoichet, B. K.; Peishoff, C. E. Prediction of protein-ligand interactions. Docking and scoring: successes and gaps. *J. Med. Chem.* **2006**, *49*, 5851–5855.

(2) Sousa, S. F.; Fernandes, P. A.; Ramos, M. J. Protein-ligand docking: current status and future challenges. *Proteins* **2006**, *65*, 15–26.

(3) Carlson, H. A.; McCammon, J. A. Accommodating protein flexibility in computational drug design. *Mol. Pharmacol.* **2000**, *57*, 213–218.

(4) Teague, S. J. Implications of protein flexibility for drug discovery. *Nat. Rev. Drug Discovery* **2003**, *2*, 527–541.

(5) Teodoro, M. L.; Kavraki, L. E. Conformational flexibility models for the receptor in structure based drug design. *Curr. Pharm. Des.* **2003**, *9*, 1635–1648.
(6) Koska, J.; Spassov, V. Z.; Maynard, A. J.; Yan, L.; Austin, N.; Flook, P. K.; Venkatachalam, C. M. Fully automated molecular mechanics based induced fit protein-ligand docking method. *J. Chem. Inf. Model.* **2008**, *48*, 1965–1973.
(7) May, A.; Zacharias, M. Protein-ligand docking accounting for receptor side chain and global flexibility in normal modes: evaluation on kinase inhibitor cross docking. *J. Med. Chem.* **2008**, *51*, 3499–3506.
(8) Meiler, J.; Baker, D. ROSETTALIGAND: protein-small molecule docking with full side-chain flexibility. *Proteins* **2006**, *65*, 538–548.
(9) Nabuurs, S. B.; Wagener, M.; de Vlieg, J. A flexible approach to induced fit docking. *J. Med. Chem.* **2007**, *50*, 6507–6518.
(10) Rueda, M.; Bottegoni, G.; Abagyan, R. Consistent improvement of cross-docking results using binding site ensembles generated with elastic network normal modes. *J. Chem. Inf. Model.* **2009**, *49*, 716–725.
(11) Barril, X.; Morley, S. D. Unveiling the full potential of flexible receptor docking using multiple crystallographic structures. *J. Med. Chem.* **2005**, *48*, 4432–4443.
(12) Sheridan, R. P.; McGaughey, G. B.; Cornell, W. D. Multiple protein structures and multiple ligands: effects on the apparent goodness of virtual screening results. *J. Comput.-Aided. Mol. Des.* **2008**, *22*, 257–265.
(13) Knegtel, R. M.; Kuntz, I. D.; Oshiro, C. M. Molecular docking to ensembles of protein structures. *J. Mol. Biol.* **1997**, *266*, 424–440.
(14) Damm, K. L.; Carlson, H. A. Exploring experimental sources of multiple protein conformations in structure-based drug design. *J. Am. Chem. Soc.* **2007**, *129*, 8225–8235.
(15) Cheng, L. S.; Amaro, R. E.; Xu, D.; Li, W. W.; Arzberger, P. W.; McCammon, J. A. Ensemble-based virtual screening reveals potential novel antiviral compounds for avian influenza neuraminidase. *J. Med. Chem.* **2008**, *51*, 3878–3894.
(16) Soliva, R.; Gelpi, J. L.; Almansa, C.; Virgili, M.; Orozco, M. Dissection of the recognition properties of p38 MAP kinase. Determination of the binding mode of a new pyridinyl-heterocycle inhibitor family. *J. Med. Chem.* **2007**, *50*, 283–293.
(17) Karplus, M. Molecular dynamics of biological macromolecules: a brief history and perspective. *Biopolymers* **2003**, *68*, 350–358.
(18) Karplus, M.; Kuriyan, J. Molecular dynamics and protein function. *Proc. Natl. Acad. Sci. U.S.A.* **2005**, *102*, 6679–6685.
(19) Fan, H.; Irwin, J. J.; Webb, B. M.; Klebe, G.; Shoichet, B. K.; Sali, A. Molecular docking screens using comparative models of proteins. *J. Chem. Inf. Model.* **2009**, *49*, 2512–2527.
(20) Keseru, G. M.; Kolossvary, I. Fully flexible low-mode docking: application to induced fit in HIV integrase. *J. Am. Chem. Soc.* **2001**, *123*, 12708–12709.
(21) Kairys, V.; Gilson, M. K. Enhanced docking with the mining minima optimizer: acceleration and side-chain flexibility. *J. Comput. Chem.* **2002**, *23*, 1656–1670.
(22) Anderson, A. C.; O'Neil, R. H.; Surti, T. S.; Stroud, R. M. Approaches to solving the rigid receptor problem by identifying a minimal set of flexible residues during ligand docking. *Chem. Biol.* **2001**, *8*, 445–457.
(23) Davis, I. W.; Baker, D. RosettaLigand docking with full ligand and receptor flexibility. *J. Mol. Biol.* **2009**, *385*, 381–392.
(24) Dunbrack, R. L., Jr.; Karplus, M. Backbone-dependent rotamer library for proteins. Application to side-chain prediction. *J. Mol. Biol.* **1993**, *230*, 543–574.
(25) Kaufmann, K. W.; Glab, K.; Mueller, R.; Meiler, J. Small Molecule Rotamers Enable Simultaneous Optimization of Small Molecule and Protein Degrees of Freedom in ROSETTALIGAND Docking. In *German Conference on Bioinformatics*; Proceedings of the German Conference on Bioinformatics, September 9−12, 2008, Dresden, Germany; Beyer, A., Schroeder, M., Eds. GI: Dresden, 2008; pp 148−157.
(26) Davis, I. W.; Raha, K.; Head, M. S.; Baker, D. Blind docking of pharmaceutically relevant compounds using RosettaLigand. *Protein Sci.* **2009**, *18*, 1998–2002.
(27) Ding, F.; Dokholyan, N. V. Emergence of protein fold families through rational design. *PLoS Comput. Biol.* **2006**, *2*, e85.
(28) Yin, S.; Biedermannova, L.; Vondrasek, J.; Dokholyan, N. V. MedusaScore: an accurate force field-based scoring function for virtual drug screening. *J. Chem. Inf. Model.* **2008**, *48*, 1656–1662.
(29) Bahar, I.; Erman, B.; Jernigan, R. L.; Atilgan, A. R.; Covell, D. G. Collective motions in HIV-1 reverse transcriptase: examination of flexibility and enzyme function. *J. Mol. Biol.* **1999**, *285*, 1023–1037.
(30) Wang, Y. X.; Freedberg, D. I.; Yamazaki, T.; Wingfield, P. T.; Stahl, S. J.; Kaufman, J. D.; Kiso, Y.; Torchia, D. A. Solution NMR evidence that the HIV-1 protease catalytic aspartyl groups have different ionization states in the complex formed with the asymmetric drug KNI-272. *Biochemistry* **1996**, *35*, 9945–9950.
(31) Kabsch, W. Solution for Best Rotation to Relate 2 Sets of Vectors. *Acta Crystallogr., Sect. A: Found. Crystallogr.* **1976**, *32*, 922–923.
(32) Yin, S.; Ding, F.; Dokholyan, N. V. Modeling backbone flexibility improves protein stability estimation. *Structure* **2007**, *15*, 1567–1576.
(33) Yin, S.; Ding, F.; Dokholyan, N. V. Eris: an automated estimator of protein stability. *Nat. Methods* **2007**, *4*, 466–467.
(34) Sherman, W.; Day, T.; Jacobson, M. P.; Friesner, R. A.; Farid, R. Novel procedure for modeling ligand/receptor induced fit effects. *J. Med. Chem.* **2006**, *49*, 534–553.
(35) Berman, H. M.; Westbrook, J.; Feng, Z.; Gilliland, G.; Bhat, T. N.; Weissig, H.; Shindyalov, I. N.; Bourne, P. E. The Protein Data Bank. *Nucleic Acids Res.* **2000**, *28*, 235–242.
(36) Huang, N.; Shoichet, B. K.; Irwin, J. J. Benchmarking sets for molecular docking. *J. Med. Chem.* **2006**, *49*, 6789–6801.
(37) Bostrom, J.; Greenwood, J. R.; Gottfries, J. Assessing the performance of OMEGA with respect to retrieving bioactive conformations. *J. Mol. Graph. Modell.* **2003**, *21*, 449–462.
(38) Friedland, G. D.; Linares, A. J.; Smith, C. A.; Kortemme, T. A simple model of backbone flexibility improves modeling of side-chain conformational variability. *J. Mol. Biol.* **2008**, *380*, 757–774.
(39) Matthew, J. B.; Gurd, F. R. Calculation of electrostatic interactions in proteins. *Methods Enzymol.* **1986**, *130*, 413–436.

CI100218T