

Physics-Based Scoring of Protein–Ligand Complexes: Enrichment of Known Inhibitors in Large-Scale Virtual Screening

Niu Huang, Chakrapani Kalyanaraman, John J. Irwin, and Matthew P. Jacobson*

Department of Pharmaceutical Chemistry, University of California San Francisco, Genentech Hall,
600 16th St., Box 2240, San Francisco, California 94143-2240

Received July 13, 2005

We demonstrate that using an all-atom molecular mechanics force field combined with an implicit solvent model for scoring protein–ligand complexes is a promising approach for improving inhibitor enrichment in the virtual screening of large compound databases. The rescoring method is evaluated by the extent to which known binders for nine diverse, therapeutically relevant enzymes are enriched against a background of ~100 000 drug-like decoys. The improvement in enrichment is most robust and dramatic within the top 1% of the ranked database, that is, the first thousand compounds; below the first few percent of the ranked database, there is little overall improvement. The improved early enrichment is likely due to the more realistic treatment of ligand and receptor desolvation in the rescoring procedure. We also present anecdotal but encouraging results assessing the ability of the rescoring method to predict specificity of inhibitors for structurally related proteins.

1. INTRODUCTION

Structure-based virtual screening, also referred to as small-molecule docking, orients and scores small molecules from large chemical databases (typically tens or hundreds of thousands of compounds) for complementarity to a macromolecular binding site. The results from virtual screening can be used to prioritize compounds for experimental testing in a cost-effective fashion;^{1–3} numerous studies have used such an approach to identify novel inhibitors for various protein targets.^{4–13}

Despite these successes, docking remains a challenging field in structure-based drug design. The critical issues include the methods for exploring the conformational space of the flexible ligands (sampling) and the estimation of binding affinities for ligand–receptor complexes (scoring). The scoring function is used both to identify the correct binding orientation and conformation (docking pose) out of enormous numbers of alternative modes for each ligand and to rank different ligands with respect to their estimated binding affinity. Therefore, to dock a large compound library, a scoring function has to be simple, fast, and derived from a physically reasonable equation.

Currently available scoring functions can be divided into three classes: force-field-based (e.g., DOCK and AUTODOCK),^{14,15} empirical (e.g., FlexX and Glide),^{16,17} and knowledge-based (e.g., PMF and SMOG).^{18,19} Force-field-based scoring methods attempt to approximately calculate the atomic interaction energies in the system. Empirical scoring functions obtain parameter coefficients by fitting to many crystal complexes with known binding affinities. Knowledge-based functions are derived from a statistical analysis of the interaction distances among different pairs of atom types in cocrystallized protein–ligand structures. Typically, scoring functions have been evaluated by testing

their ability to reproduce ligand binding poses or affinities; however, none of them is able to predict experimental binding free energies accurately in all situations.^{20–24} Several studies have also evaluated the rate of “enrichment”—the increase in the proportion of active compounds found in selected subsets from docking calculations compared with the proportion expected from random selection—obtained using various scoring functions and docking algorithms.²⁵

There is no fundamental reason that the same scoring function must be used to both select the correct binding pose for a ligand and rank ligands with respect to their estimated binding affinities. It is possible to define two-step strategies that decouple these two processes.^{10,17,24,26–28} The second stage, used to rank the compounds, may then use a more computationally intensive scoring procedure, because it is applied only to single poses of ligands. Molecular-mechanics-based scoring functions have been applied in a second stage of this type to “rescore” docking results, either to improve the ability to reproduce crystallographic binding poses^{24,29} or to rerank ligands in a virtual screening context.^{24,27} However, to date, there has been no large-scale evaluation of the ability of such scoring functions to *enrich* known binders among large numbers of decoys. Enrichment is qualitatively more challenging than simply reproducing crystallographic poses, in terms of both the demands on algorithmic efficiency and the accuracy of the scoring function. In evaluating a scoring function, it is also important to use a diverse set of binding sites, each of which challenges the scoring function in different ways, because of differences in chemical composition and conformation.

Other methods related to the work reported here, in terms of the use of molecular mechanics energy functions, include free-energy perturbation (FEP), thermodynamic integration (TI), one-window free energy grid (OWFEG), mining minima, and molecular mechanics Poisson–Boltzmann/surface area (MM-PB/SA). These methods have been successfully used to estimate the relative, or in a few cases

* Corresponding author phone: 415-514-9811; fax: 415-514-4260; e-mail: matt@cgl.ucsf.edu.

Table 1. Measures of Enrichment of the Known Inhibitors for Nine Enzyme Systems Achieved by Docking Alone (D) and the Rescoring Procedure (R)^a

enzyme	PDB code	number of known inhibitors	% of ranked database needed to find 25% of known inhibitors		maximum enrichment factor achieved		% of ranked database where maximum enrichment factor occurred		number of known inhibitors found in top 0.1% of ranked database		number of known inhibitors found in top 0.5% of ranked database	
			D	R	D	R	D	R	D	R	D	R
DHFR	3dfr ⁶⁴	117	0.3	0.3	111	204	0.1	0.1	15	25	34	32
DHFR (mod)	3dfr ⁶⁴	100	0.3	0.1	110	239	0.1	0.1	13	25	31	32
GART	1c2t ⁶⁵	50	0.9	0.8	46	159	0.3	0.1	0	8	8	12
AR	1ah3 ⁶⁶	722	3.5	4.0	8	12	2.0	0.1	6	9	19	33
PARP	1efy ⁶⁷	45	4.6	2.3	6	11	5.2	3.8	0	0	0	2
PNP	1b8o ⁶⁸	25	1.2	0.1	60	358	0.2	0.1	1	9	3	11
SAHH	1a7a ⁶⁹	37	2.1	1.8	14	19	1.3	2.0	0	0	0	0
thrombin	1ba8 ⁷⁰	243	4.2	0.8	25	49	0.1	0.1	6	13	21	58
AChE	1e66 ⁷¹	554	5.0	5.1	21	25	0.4	0.1	5	13	59	30
TS	2bbq ⁷²	171	1.5	0.5	25	52	0.3	0.1	3	9	19	44

^a Abbreviations: AR, aldose reductase; DHFR, dihydrofolate reductase; GART, glycinamide ribonucleotide transformylase; PARP, poly(ADP-ribose) polymerase; PNP, purine nucleoside phosphorylase; SAHH, *S*-adenosylhomocysteine hydrolase; AChE, acetylcholinesterase; TS, thymidylate synthase. DHFR (mod) refers to excluding 17 inhibitors with known problems (prodrugs, incorrect tautomerization, and parametrization failure), as discussed in the text.

absolute, binding free energy of a series of docked ligands.^{30–34} However, these methods are computationally expensive and have generally been applied only to dozens of compounds. By limiting the sampling performed during the molecular mechanics rescoring stage to simple minimization, we are able to apply it to tens or hundreds of thousands of ligands using a small Linux cluster.

In our previous work, a similar docking and rescoring approach was shown to improve the ability to identify substrates of α - β barrel enzymes by virtual metabolite screening.³⁵ However, those enzymes have small and highly charged binding sites, which is not the case for most drug targets. Here, we demonstrate that an all-atom molecular mechanics force field (OPLS-AA), combined with an implicit solvent model (Generalized Born, GB), can be used to enrich known inhibitors of a diverse set of therapeutically relevant enzymes. Specifically, the maximum enrichment factors observed increased for all nine of the test cases, by up to a factor of 6. The improvement in enrichment is most robust and dramatic within the top 1% of the ranked database, that is, the first thousand compounds. The improved early enrichment is likely due to the more realistic treatment of ligand and, especially, receptor desolvation in the rescoring procedure; the fully flexible minimization of the ligands in the receptor during the rescoring stage may also contribute to the improved enrichment. To our knowledge, this work represents the most extensive test to date of the utility of an all-atom force field/implicit solvent model scoring function in the context of high-throughput virtual screening.

Although the rescoring method improves enrichment the most within the top 1% of the ranked compound database, in all cases, known inhibitors are enriched significantly relative to random selection throughout at least the top 20% of the ranked database. This behavior may be important if the method is to be used in combination with high-throughput experimental screening methods. In four of the nine test cases, the rescoring method robustly improves enrichment, relative to docking alone, well beyond the top 1% of the ranked database. In the other test cases, however, the results of the docking and rescoring methods are roughly comparable beyond the top 1%. By carefully analyzing the results for

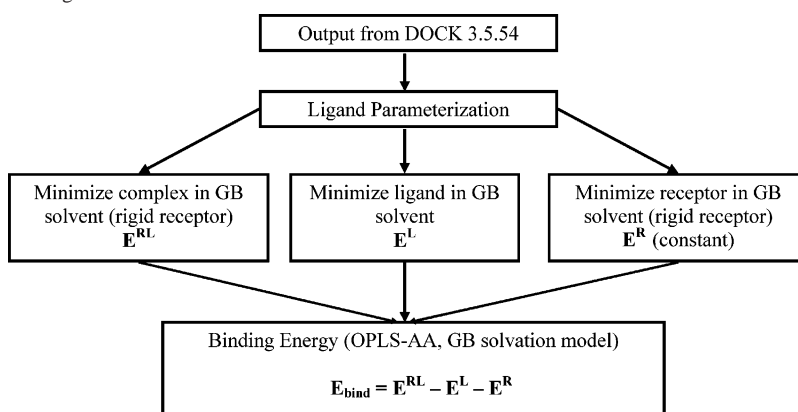
one test case (DHFR), we suggest some reasons for this behavior, which reflects many factors in addition to the quality of the scoring function.

We also present anecdotal but encouraging results assessing the ability of the rescoring method to predict specificity of inhibitors for structurally related proteins.

2. METHODS

Molecular Docking. We use the test set of McGovern and Shoichet, containing nine therapeutically important enzymes (Table 1). Briefly, each protein was prepared for docking in the same manner as previously described.³⁶ When cofactors were present, they were treated as part of the protein. The molecular solvent-accessible surface³⁷ was calculated with the program DMS³⁸ using a probe radius of 1.4 Å. Polar hydrogens were added to the proteins using SYBYL.³⁹ Matching spheres, required for initial placement of the ligand during database screening, were obtained from the position of the crystallographic ligand using the program SPHGEN.¹⁴ A “coloring” scheme⁴⁰ was applied to individually label the matching spheres on the basis of their hydrogen-bond properties and the charge states of nearest neighboring atoms in the protein binding site. Four different types of grids were generated before the docking calculations, including an excluded volume grid obtained from DISTMAP,⁴¹ a united AMBER-based van der Waals potential grid computed by CHEMGRID,⁴¹ an electrostatic potential grid calculated using DelPhi,⁴² and a ligand desolvation grid computed using SOLVMAP (B. K. Shoichet, unpublished results).

The program DOCK 3.5.54 was used to dock the MDL Drug Data Report (MDDR) database into the protein binding site.^{43,44} DOCK 3.5.54 implements an alternative method of whole-molecule-based docking^{14,15,45–47} to sample the ligand conformational space, where ensembles of precalculated conformers from conformationally expanded databases are used to significantly speed up docking calculations.⁴⁴ To sample ligand orientations, the bin size for both receptor and ligand was set to 0.4 Å and the overlap bin size was set to 0.3 Å. A distance tolerance (dislim) of 1.5 Å was applied for matching the ligand to the spheres, and ligand orientations

Scheme 1. Refinement and Rescoring Protocol^a

^a The superscript R refers to the free receptor in solution, L refers to the ligand in solution, and RL refers to the protein-ligand complex in solution. E_{bind} is the predicted ligand binding energy, the free receptor energy in solution (E^{R}) is a constant value, E^{L} is the energy of the free ligand in solution, and E^{RL} is the energy of the ligand-protein complex in solution.

were rejected if the color of a ligand-receptor pair did not match. For each ligand orientation, the conformational ensemble was filtered for steric complementarity using DISTMAP with polar and nonpolar close contact limits of 2.3 and 2.6 Å, respectively. Ligand conformations are scored on the basis of the docking total energy ($E_{\text{tot}} = E_{\text{ele}} + E_{\text{vdw}} - \Delta G_{\text{lig-solv}}$), which is the sum of electrostatic (E_{ele}) and van der Waals (E_{vdw}) interaction energies corrected by the ligand partial desolvation energy ($\Delta G_{\text{lig-solv}}$). This precalculated atomic desolvation penalty was from AMSOL, as previously described,⁴⁸ and the partial desolvation penalty was based on the fraction of surface area buried by the receptor for each ligand atom (B. K. Shoichet, unpublished results). Final energies were computed after 25 steps of rigid-body minimization. Then, a single docking pose with the best total energy score was saved for each docked molecule. For ligands with multiple protonation states, only the best scoring version was retained for further consideration.

Molecular Mechanics Rescoring. The top 25% of the ranked database from the docking was submitted to our rescoring protocol. As discussed in the Results, this threshold was chosen because the docking algorithm typically ranked most known binders within the top 25% of the database.

All energy minimizations were performed using the Protein Local Optimization Program (PLOP)⁴⁹⁻⁵¹ with the all-atom OPLS force field (OPLS-AA)^{52,53} and the Surface Generalized Born (SGB) implicit solvent model.^{54,55} PLOP implements a multiscale truncated-Newton (MSTN) minimization algorithm. This algorithm is described in detail elsewhere;⁵⁶ because it is critical to the success of this work, we briefly summarize the results here. The algorithm is adapted from TNPACK⁵⁷ and optimized by applying multiscale methods, analogous to those used in molecular dynamics (e.g., r-RESPA).⁵⁸ The molecular mechanics forces are divided into short- (bond, angle, torsion, and local nonbonded) and long-range components, with the long-range forces updated only intermittently (never during the inner TN cycles and infrequently during the outer cycles). The speedup of MSTN relative to the unmodified TNPACK algorithm depends on system size and the distance cutoffs used for defining the short- and long-range interactions and the long-range force updating frequency, but it is a factor of 4.0-4.5 faster with the parameters used here. The algorithm is also optimized for minimizations with generalized Born implicit solvents,

using a self-consistent procedure that increases the computational expense, relative to vacuum, by only a small factor (~ 3). The termination criterion for the minimization was determined by the root-mean-squared gradient and the maximum number of truncated Newton (TN) steps with default values of 0.001 kcal/mol/Å and 65, respectively. Cutoffs for the nonbonded interactions are residue-based and depend on the type of side chain (charged or neutral). We employ fixed absolute cutoffs (long-range cutoffs) of 30 Å for charged-charged residue pairs, 20 Å for charged-neutral pairs, and 15 Å for neutral-neutral pairs, with no smoothing. The short-range cutoffs for charged-charged residue pairs are 15 Å, and for all other pairs, they are 10 Å. The long-range interactions are updated during every fifth outer cycle of the TN minimization.

Each protein was prepared for ligand rescoring using the same procedure. The same protein structure file used in the docking was used for rescoring. When cofactors were present, the program IMPACT⁵⁹ was used to generate OPLS force field parameters for it. Hydrogen atoms were added in standard geometries as defined by the OPLS force field using PLOP. The positions of hydrogen atoms on OH and SH groups were determined as the lowest energy state by scanning the hydrogen dihedral angles at 10° intervals using the OPLS force field with GB solvation, followed by energy minimization of all hydrogen atoms. The resulting protein structure was used for generating ligand-protein complexes for the rescoring step. Note that all heavy atoms were held fixed during this receptor preparation procedure.

The rescoring procedure for a single protein-ligand complex is shown in Scheme 1. The first step is to generate OPLS force field parameters for each ligand using IMPACT,⁵⁹ after which the coordinate and parameter files are passed to PLOP. The protein-ligand complex and the free ligand were then submitted to energy minimization in a GB solvent. The binding energy ($E_{\text{bind}} = E^{\text{RL}} - E^{\text{L}} - E^{\text{R}}$) was calculated by subtracting the energies of the free ligand in solution (E^{L}) and the free protein in solution (E^{R}) from the ligand-protein complex's energy in solution (E^{RL}). In this work, the protein was kept rigid during minimization of the ligand-protein complex to reduce the computational expense. However, in other works, portions of the receptor are allowed to relax during minimization, to account for receptor strain (unpublished results).

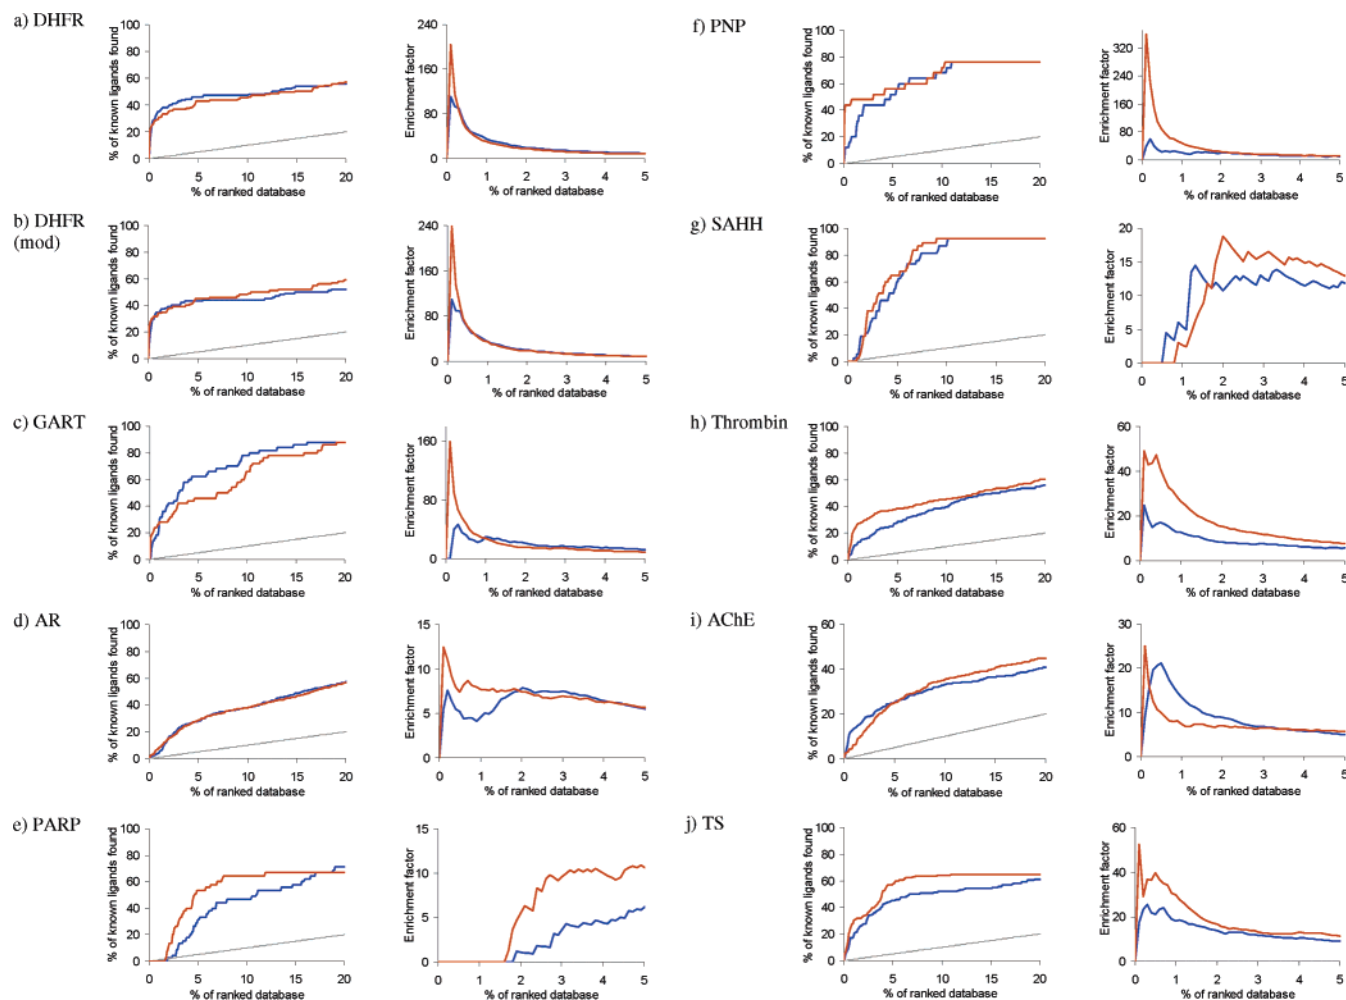


Figure 1. Enrichment plots for nine enzyme systems obtained after docking alone (blue line) and after rescoring (orange line). (Left) The percent of known inhibitors identified in increasingly large subsets of the ranked database. The gray line represents the results expected from a random selection of ligands. (Right) Enrichment factor as a function of the fraction of the ranked database. DHFR (mod) refers to excluding 17 DHFR inhibitors with known problems (prodrugs, incorrect tautomerization, and parametrization failure), as discussed in the text. This additional analysis, resulting from an extensive manual inspection of the results, was not performed for the other cases.

Enrichment Calculations. The MDDR database containing 95 579 unique molecules was used as a background of drug-like decoys for enrichment calculations. We assume that molecules annotated as inhibitors of a given enzyme in the MDDR are true positives and the remaining molecules are true negatives (neither of these assumptions is likely to be entirely correct, as discussed below). The quality of enrichment is measured as the proportion of true binders found in selected subsets from the docking (or rescoring) calculations compared with the proportion expected from random selection. The enrichment factor (EF) is calculated as $EF_{\text{subset}} = \{\text{binders}_{\text{subset}}/N_{\text{subset}}\}/\{\text{binders}_{\text{total}}/N_{\text{total}}\}$.⁴⁸

3. RESULTS

Early Enrichment Improved via Rescoring. The key results are summarized in Figure 1 and Table 1. Two types of graphs are presented in Figure 1.^{36,48} The left panel presents the percent of known inhibitors found (y axis) as a function of the percent of the ranked database (x axis). The line $y = x$, shown in gray, is the curve expected when randomly selecting compounds; thus, the curves obtained from the docking and rescoring should rise above this line by an amount related to the enrichment of the known

inhibitors. The higher the percentage of known inhibitors found at a given percentage of the ranked database, the better the performance of the virtual screening. The right panel of Figure 1 is the enrichment factor (y axis) as a function of the percentage of the ranked database (x axis), which emphasizes the enrichment of known binders in the top 5% of the ranked database. High enrichment factors near the top of the database (i.e., those compounds that would actually be tested in a typical experimental screen) are desired.

Table 1 summarizes the results using five indicators: the percent of the ranked database required to find 25% of the known inhibitors, the maximum enrichment factor, the location of the maximum enrichment factor, and the numbers of known inhibitors found in the top 100 and 500 ranked compounds. A key objective in database screening is to find active compounds as early as possible in the ranked database. In many typical screening efforts (excluding high-throughput screening), only a few dozen to a few hundred compounds from the top of the docking hit list are selected for experimental screening. In Table 1, we report the number of known inhibitors found in the top 100 and 500 ranked compounds ($\sim 0.1\%$ and $\sim 0.5\%$ of the database, respectively). Encouragingly, the rescoring procedure appears to

robustly improve early enrichment, especially within the top few tenths of a percent of the database, by several measures (Figure 1 and Table 1).

As a case study, we carefully examine the results for DHFR. Docking alone gives a maximum enrichment factor of 111, occurring at the top 0.1% of the database, which corresponds to 15 known inhibitors found in top 100 compounds in the ranked list. With the rescoring procedure, the maximum enrichment factor (204) is nearly doubled, with 25 known inhibitors found in the top 0.1% of the database. By contrast, the rescoring actually performs slightly worse than docking alone beyond the top 0.3% of the database (Figure 1a, left panel). At 0.5% of the database, 34 known inhibitors are found by docking while only 32 inhibitors are found by rescoring. The maximum difference along the y axis (8% of known inhibitors) between docking and rescoring occurs at the top 3.8% of the database (Figure 1a, left panel), which corresponds to the net loss of nine DHFR inhibitors after rescoring. Five DHFR inhibitors (with docking ranks ranging from 0.7% to 4%) were skipped in rescoring because of a failure in the parametrization step; however, this cannot explain the difference completely.

A similar scenario is observed in GART, where rescoring triples the enrichment factor and lowers the percentage of the database from 0.3% to 0.1% where the maximum enrichment factor occurs. However, the rescoring starts to perform worse than docking beyond the top 1% of the database (Figure 1c, left panel). The maximum difference along the y axis (20% of known inhibitors) between docking and rescoring occurs at the top 6% of the database, corresponding to the net loss of 10 GART inhibitors by rescoring. Clearly, rescoring improves early enrichment, which we consider a significant benefit, but also penalizes some known binders, further down the hit list, in these two cases. AR arguably belongs in this category as well, with significant improvements in enrichment up through the top 2% of the database and a similar to slightly worse enrichment further down the list (Figure 1d, left panel).

To understand what causes these differences, the atomistic interactions between representative known binders were analyzed in detail. For instance, DHFR inhibitors **1** and **2** (structures shown in Figure 2a) are ranked #1828 and #22 after the initial docking and #4 and #17 020 after rescoring, respectively. The docking poses of both compounds superimpose well with the crystallographic ligand methotrexate. The refinement and rescoring procedure significantly improves the relative rank of compound **1** but dramatically worsens the rank of compound **2**. The key hydrogen bonding interactions between methotrexate and the DHFR binding site residues are illustrated in Figure 2b. Minimization of the ligand inside the protein binding pocket generates relatively small conformational changes but clearly enhances energetically important interactions such as the hydrogen bond between the two charged carboxylate groups of compound **1** and the binding site residues Arg57 and His28 (Figure 2c). On the other hand, rescoring of compound **2** magnifies the unfavorable interactions between the hydrogen-bond acceptor atom O6 of compound **2** and the backbone oxygen atoms of protein residues Leu4 and Ala97 (Figure 2d), where chemical functional groups containing hydrogen-bond donors (pteridines and pyrimidines) are generally presented in known DHFR inhibitors. The problem here is

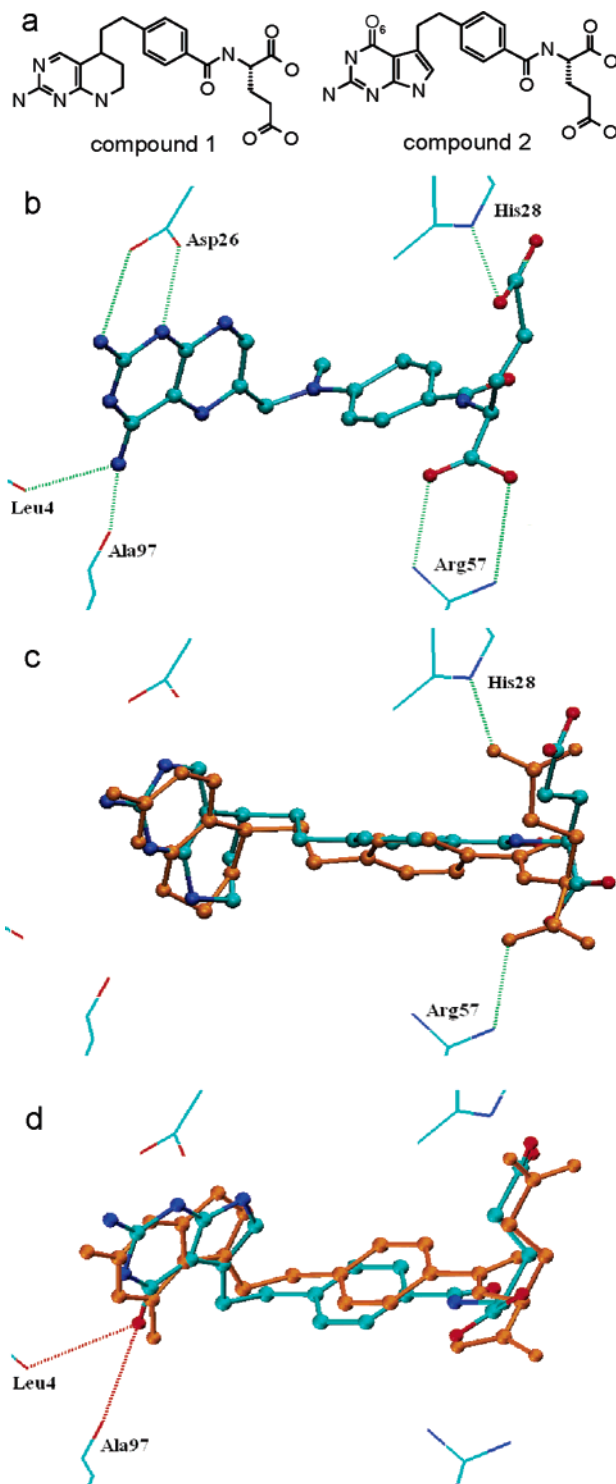


Figure 2. (a) Structures of two DHFR inhibitors, compound **1** and compound **2**, which ranked #1828 and #22 after docking versus #4 and #17 020 after rescoring, respectively. (b) The crystallographic ligand (methotrexate; structure shown in Figure 6a) is represented by a CPK model colored by atom type. The key hydrogen bond interactions between the protein and methotrexate are illustrated with dashed green lines. (c) The docked pose of compound **1** is colored by atom type, while the refined binding pose is in orange. Two hydrogen-bond interactions, identified by dashed green lines, are recovered after minimizing the ligand during the refinement stage. (d) The docked pose of compound **2** is colored by atom type, while the refined binding pose is in orange. The unfavorable interactions between backbone oxygen atoms of protein residue binding site residues and the O6 atom of compound **2** are illustrated with dashed red lines. Molecular images were generated with VMD.⁷³

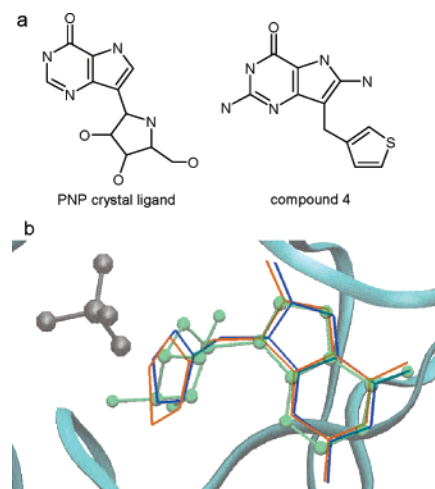


Figure 3. (a) Structures of the PNP cocrystallized ligand and compound **4**, which ranked #674 after docking alone and #1 after rescoring. (b) The docked pose of compound **4** (blue) is superimposed over the minimized binding pose after the rescoring stage (orange). The protein structure is represented as an aqua-colored ribbon, while the cocrystallized ligand is shown in green. The phosphate group cofactor is shown in gray.

simply that the wrong tautomer state for compound **2** was presented in our study because of an incomplete treatment of tautomerization in our ligand database preparation. In this case, enolization of the O6 atom of compound **2** would restore the favorable interactions in the rescoring. In the MDDR database, 10 DHFR inhibitors are derivatives of compound **2**, and all of these are ranked reasonably highly by the docking scoring function (within the top 3% of the database) but poorly by rescoring (outside of the top 10% of the database). Another set of six compounds were “de-enriched” during rescoring because they are ester prodrugs instead of free glutamate acid moieties.

It appears that the simple scoring function employed in the docking method is less sensitive to such errors, while the more physically reasonable molecular mechanics energy employed in the rescoring requires accurate treatment of protonation and charge states to correctly account for the electrostatic properties of ligands. Improvements to the database preparation procedure are being pursued. Simply excluding the problematic DHFR ligands discussed above improves the enrichments obtained upon rescoring, particularly beyond the top 0.3% of the database (Figure 1b and “DHFR (mod)” in Table 1).

In the cases of PARP, PNP, thrombin, and TS (Figure 1), rescoring consistently performs better than docking alone. In PNP, rescoring increases the maximum enrichment factor 6-fold over docking alone. One of the most potent PNP inhibitors (compound **4** in Figure 3a), currently in phase I trials, ranks #1 after rescoring and #674 after docking alone. Figure 3b presents the pose of this inhibitor overlapped with the PNP cocrystallized ligand.

The remaining two test cases (SAHH and AChE) are more complicated. The maximum enrichment factors for docking and rescoring are similar and below 25. SAHH is the only case where the earliest enrichment is worse after rescoring, albeit not by a huge margin. Shown in Figure 4 is a known SAHH inhibitor, which ranks #568 after docking alone and #1705 after rescoring. Clearly, the pentose ring of compound **6** is rotated ~ 180 degrees from the orientation observed in

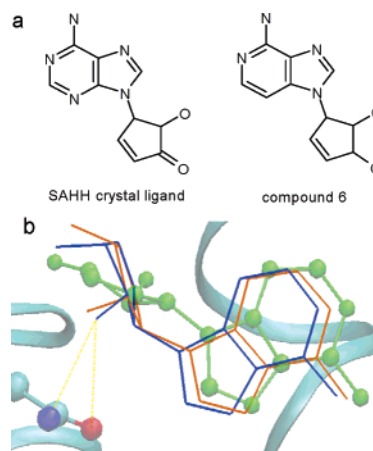


Figure 4. (a) Structures of the SAHH cocrystallized ligand and compound **6**, which ranked #568 after docking and #1705 after rescoring. (b) The docked pose of compound **6** is shown in blue, while the refined binding pose is in orange. The newly formed hydrogen bonds between the sugar moiety of compound **6** and cofactor NAD⁺, after minimization during the rescoring procedure, are marked with yellow lines. The protein structure is represented by aqua ribbon, while the crystallographic ligand is shown in green, and the cofactor is colored by atom type.

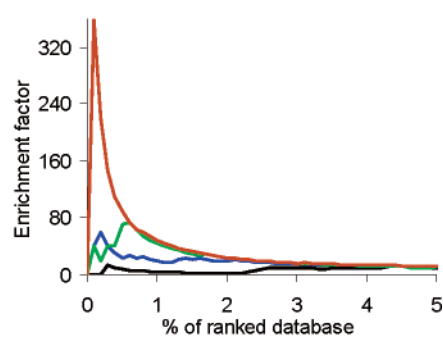
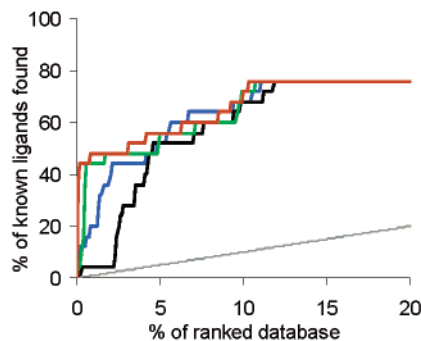
the cocrystallized ligand. We do not know the origin of this error; possibilities include inadequate sampling of the pentose ring conformations, poor parametrization of the cofactor, or simply limitations of the scoring function used to select the poses.

For AChE, the rescoring improves enrichment at the very top of the database but performs significantly worse lower down, even at 1%. We believe many of the ligands that rank worse after rescoring may have incorrect binding poses generated by docking, although we cannot exclude other explanations. The AChE binding cavity is large, and compounds can bind to several different regions within the pocket, making it a particularly difficult test case.⁶⁰

Correct Protonation States Are Critical for the Rescoring. Unsurprisingly, incorrect protonation states on the ligands, receptor, or cofactors significantly affect the electrostatic potential, which in turn strongly affects the rescoring calculations. In this work, we did not attempt any sophisticated prediction of protonation states. Instead, we simply inspected binding sites “by eye”. With respect to the protein, we simply identified His residues that should be positively charged as those that are within hydrogen-bonding distance of a carboxylate group. With respect to cofactors, a phosphate group is present in the PNP binding site and interacts with ligands (shown in Figure 3). The assignment of different protonation states to the phosphate group dramatically changes the enrichments obtained from rescoring (Figure 5a); however, it affects the enrichments obtained from docking very little (Figure S1 in the Supporting Information). Strikingly, the maximum enrichment factor after rescoring increases by factors of 3 and 30 when the formal charge on the phosphate group is changed from -3 to -2 and -1 , respectively. The correct protonation state of this phosphate group is clearly -1 on the basis of physical principles. The pK_a value for $\text{HPO}_3^{2-}/\text{H}_2\text{PO}_3^-$ is ~ 7 , and desolvation of the ion in the binding site will shift the pK_a lower.

In DHFR, the enrichment of known inhibitors was originally negatively impacted by incorrect formal charges

a). PNP



b). DHFR

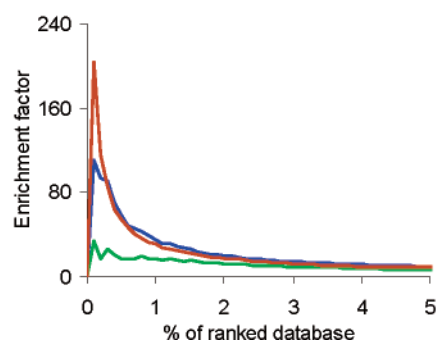
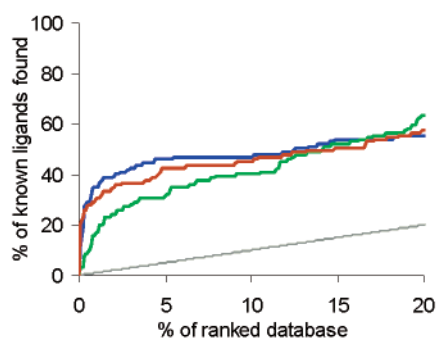


Figure 5. (a) Enrichment plots for PNP obtained from docking (blue line) and rescoring with different charge states assigned to the cofactor phosphate group: H_2PO_3^- (orange line), $\text{H}_1\text{PO}_3^{2-}$ (green line), and PO_3^{3-} (black line). (b) DHFR enrichment results obtained from docking (blue line) and rescoring with incorrectly parametrized ligands (green line) and correctly parametrized ligands (orange line). The incorrectly parametrized ligands have a formal charge of 0 on the pteridine ring.

being assigned to many DHFR inhibitors (Figure 5b), such as methotrexate (Figure 6). The program IMPACT failed to recognize protonated methotrexate analogues and assigned a formal charge of 0 to pteridine rings, which resulted in the rescoring producing enrichment factors almost 3-fold worse than docking alone (Figure 5b). However, when this problem was corrected by assigning partial atomic charges calculated using AMSOL⁴⁸ before rescoring the database compounds, the enrichment was significantly improved.

These results fully agree with previous work demonstrating that improved modeling of protonation states leads to a better prediction of binding affinities when implicit solvent models are used.²⁹ Clearly, the proper treatment of tautomeric forms and protonation states on receptor binding residues and ligands is a key requirement for taking advantage of the more accurate electrostatics calculations carried out in the rescoring method.

The Fraction of the Database Selected for Rescoring.

It is common in virtual screening to subject top-ranking compounds identified by a docking algorithm to some type of rescoring procedure (generally not a force-field-based method as in this work). The fraction of the database subjected to rescoring is generally arbitrary, ranging from 1% to 10% of the entire database.²⁴ This parameter can affect not only efficiency (i.e., rescoring more ligands takes a longer time) but also, potentially, accuracy. That is, if few true inhibitors are ranked below, for example, 10%, then rescoring a larger portion of the database will do little to improve enrichment and could even make it worse if some decoys rank highly because of limitations of the energy function or other technical problems.

Here, we address this issue systematically by performing additional enrichment studies where different fractions (1%, 5%, 10%, 15%, 20%, and 25%) of the docked database were

selected for rescoring. The results are shown in Figure S2 in the Supporting Information. In the systems with relatively poor enrichment (AR, AChE, PARP, and SAHH), the best enrichment is given by rescoring a relatively small portion of the ranked database, approximately the top 5%. For these four cases, rescoring a great fraction of the database increases the “noise” more than the “signal” and degrades the overall enrichment. In the systems with moderate enrichment after docking (thrombin, GART, and TS), enrichments are improved maximally by rescoring $\sim 15\text{--}25\%$ of the database. Finally, for DHFR and PNP, the enrichments are essentially insensitive to the rescoring cutoff.

In general, the systems with the best overall enrichment after docking alone show the least sensitivity to the fraction of the database chosen for rescoring. This arguably makes sense, because the rescoring procedure relies on the docking algorithm to generate accurate poses (inaccurate poses can be due either to inadequate sampling or scoring by the docking algorithm or to induced fit effects that are not accounted for with a rigid receptor). The systems with the best enrichment presumably reflect a higher fraction of ligands being docking with accurate poses, and we postulate that these ligands are generally enriched further upon rescoring. Generally, 5% of docked database seems to be the lower limit for rescoring, while 25% is the upper limit.

Better Selectivity Profiles Achieved by Rescoring.

Selectivity is a major challenge in drug discovery efforts that target proteins in families with many members possessing close structural similarity. Here, we provide a preliminary and admittedly anecdotal assessment of the ability of the rescoring protocol to improve selectivity.

The substrates of DHFR, GART, and TS enzymes are folate-related compounds. Generally, their inhibitors are chemically similar (the cocrystallized ligands are shown in

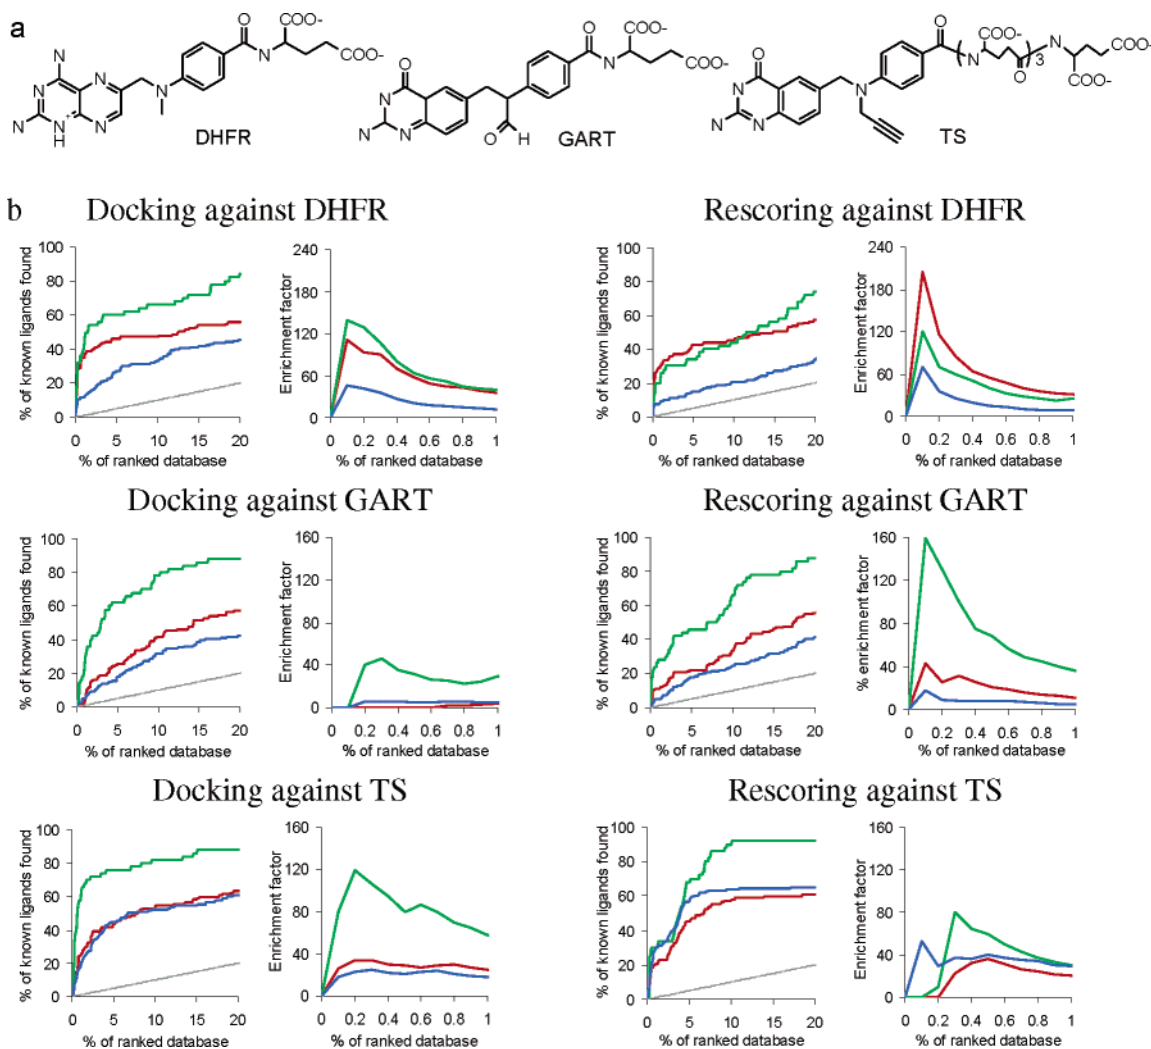


Figure 6. (a) Structures of the crystallographic ligands of DHFR, GART, and TS. (b) Enrichment against three folate enzyme systems: DHFR, GART, and TS. The enrichment of ligands annotated as DHFR inhibitors is shown in green, the enrichment of GART inhibitors is in red, and the enrichment of TS inhibitors is in blue. Results after docking alone are shown on the left, and those after rescoring are shown on the right.

Figure 6a), and there exists evidence that the DHFR inhibitor methotrexate inhibits GART and TS.⁶¹ In the MDDR database, 16 compounds are annotated as inhibitors of all three enzymes (DHFR, GART and TS), while another set of nine compounds are annotated as GART and TS inhibitors. The docking results for DHFR show that GART inhibitors are actually enriched more significantly than the ligands annotated as DHFR inhibitors (Figure 6; a similar result is seen for TS). Encouragingly, the results after rescoring correct this situation. That is, the earliest enrichment (0.1% of the database) shows a strong preference of DHFR for the inhibitors annotated as DHFR selective and less preference for the ligands annotated as TS and GART inhibitors. Similar results are seen for rescoring against TS and GART. These results suggest that our rescoring method is able to discriminate not only between actives and inactives but also between closely related analogues, at least in this case.

4. DISCUSSION AND CONCLUSION

We have developed a two-stage virtual screening protocol, in which a rapid-to-compute, grid-based scoring function (implemented in DOCK 3.5.54) is used to dock large compound databases to a receptor, and a more computationally

intensive molecular-mechanics-based energy function is used to rescore single poses for the top 25% of the ligands from the docking phase. The rescoring procedure uses the OPLS all-atom force field and a generalized Born implicit solvent model and accounts for ligand/receptor desolvation and, to a lesser extent, ligand strain energies in a more physically realistic manner than the docking algorithm (vide infra). The overall computational expense of the rescoring protocol is relatively modest, about 1 min per ligand on recent-generation personal computers, and thus can be applied to the large compound databases commonly employed for docking applications. One critical technical advance making this work possible is the multiscale truncated Newton minimization algorithm for rapidly relaxing the ligands in the protein receptor (with implicit solvation), as described in the Methods.

We evaluated the success of the new rescoring procedure by the extent to which known inhibitors were enriched against a background of drug-like decoys. Encouragingly, for all nine cases, the maximum enrichment factor increased upon rescoring, by up to a factor of 6. After rescoring, the maximum enrichment factor occurred early in the ranked database, around ~0.1% (top 100 compounds) for seven out

Table 2. Comparison of Various Characteristics of the Scoring Methods Used in Docking Alone (with DOCK 3.5.54), Molecular Mechanics Rescoring as Described Here, and More Rigorous Methods for Estimating Relative or Absolute Binding Affinities

	docking	rescoring	free energy methods
force field	united atom	all-atom	all-atom
nonbonded interaction energy	grid-based potential	pairwise potential	pairwise potential
ligand strain	N/A ^a	partial treatment, with flexible ligand minimization	treated via MD or MC
receptor strain	N/A ^a	none in this work (can be included by minimization)	treated via MD or MC
ligand desolvation	partial atomic desolvation energy	implicit solvent	implicit or explicit solvent
receptor desolvation	N/A ^a	implicit solvent	implicit or explicit solvent
computational timing	one second per ligand for sampling millions of docking poses	one minute per ligand for minimizing one docking pose	days per ligand

^a N/A = not applicable.

of nine cases. Enrichment factors beyond the top 1% of the database were generally not affected substantially by rescoring and, in a few cases, were worse than docking alone. Our detailed evaluation of the results for DHFR suggested several reasons for this, including many artifacts unrelated to the scoring function. As discussed below, we believe that the two most significant limitations of the rescoring method in its current form are related to incorrect poses generated by the docking algorithm and the rigid receptor approximation applied in this work.

Although we do not have sufficient data to make a firm conclusion, the rescoring method appears to work best on cases where the docking alone generates good enrichment. For five of the nine test cases, the docking algorithm generated maximum enrichment factors of 25 or greater; in all five of these cases, the rescoring at least doubles the maximum enrichment factor, with the largest improvement for PNP (6-fold). Finally, results on three folate enzymes suggest that our post-docking rescoring process may help to predict the selectivity of ligands toward related proteins, although we acknowledge that these results are anecdotal.

Table 2, which summarizes the differences between the sampling and scoring functions used in the docking and rescoring stages, helps both to rationalize the success of the rescoring at improving enrichment and to point the way to further improvements. The major physical effect treated by the rescoring but not by the docking scoring function is desolvation. DOCK 3.5.54 includes a partial treatment of ligand desolvation but currently does not treat receptor desolvation.^{42,48,62} A complete treatment of ligand and receptor desolvation, even at the implicit solvent level, is incompatible with the grid-based scoring required for the rapid screening of many ligand poses. The rescoring method includes a full implicit solvent treatment of ligand and receptor desolvation. The use of an all-atom force field, as opposed to the united atom force field used in DOCK, also undoubtedly improves the treatment of hydrogen bond and other nonbonded interactions.

Nonetheless, the fully flexible minimization of the ligand in the receptor is critical to the success of the rescoring; results without this minimization show very poor enrichment (results not shown). In other works, energy minimization of docking poses was shown to significantly improve the enrichments in systems with sterically demanding binding pockets.²⁴ Clearly, the minimization performed by the rescoring method cannot rescue grossly misdocked ligands (e.g., Figure 4). However, in the cases of docking poses close

to the native states (Figure 3), energy minimization is capable of locally refining the binding geometries, and the minimized energy is effective at improving enrichment.

In Table 2, we also compare the docking and rescoring methods with free energy methods, by which we refer to more rigorous methods of computing free energies, such as FEP, TI, mining minima, and OWFEG, as well as MM-PB/SA and MM-GB/SA, which involve some simplifying approximations.^{30–34} We view our rescoring method as an intermediary between high-throughput docking methods and more rigorous molecular-mechanics-based methods. It uses the same all-atom force fields typically applied in more rigorous free energy methods but uses a generalized Born implicit solvent and limits sampling to simple minimization (beyond the extensive ligand sampling provided by the docking algorithm). Our rescoring approach is also intermediate between high-throughput docking and free-energy methods in terms of computational expense. It is orders of magnitude slower than the docking algorithm but orders of magnitude faster than more rigorous free energy estimates. Ultimately, we can envision following up the physics-based rescoring with even more computationally intensive (but presumably more accurate) methods for a subset of ligands.

The minimization of a ligand in a rigid protein receptor requires only ~15 s, in implicit solvent. However, PLOP takes ~45 s to load a protein-ligand complex, so that the overall computational expense is ~1 min per complex. A more efficient rescoring procedure may be developed simply by optimizing the data-loading algorithm. In addition, there are several ways to further improve our rescoring protocol without greatly increasing the computational expense.

1. More Robust and Automated Treatment of Ligand and Receptor Protonation States. We have shown several examples where the successful application of the rescoring method requires an accurate assignment of protonation states. In this work, we relied on visual inspection for assigning protonation states to the protein, while ligand protonation states were assigned during the automated ligand library preparation. Both of these processes can be made more robust. With respect to the proteins, we cannot rule out the possibility that some protonation states remain incorrectly assigned, and we are testing automated methods for assigning protonation states. With respect to the ligands, tautomerization was incompletely addressed in the library preparation in this work; in addition, the automated parameter assignment failed for a small fraction of the ligands. Fixes for these technical problems are being pursued.

2. *Improving the Quality of the Molecular Mechanics Energy Function.* The 2003 version of the OPLS-AA force field greatly improves the parametrization of ligands, especially for torsional parameters. We are also testing a polarizable version of the OPLS force field and Poisson–Boltzmann implicit solvent in our rescoring scheme.

3. *Rescoring Multiple Poses.* One major limitation of the current protocol is that it relies entirely on the docking algorithm to identify the correct binding pose. However, the better treatment of desolvation applied at the rescoring stage can, in principle, also help to identify correct poses, as has been shown in other work.^{29,63} A simple extension of the current method is to subject a small number of dissimilar binding poses to minimization in the rescoring step and use the most favorable binding energy for rank-ordering ligands.

4. *Incorporating Receptor Flexibility.* Small amounts of receptor flexibility can be included in the current scheme simply by minimizing the receptor along with the ligand during the rescoring stage. In early tests of this procedure, receptor minimization increases the computational expense by only a small factor (<2), if only the residues contacting the ligands are minimized.

Finally, we have not yet tested the ability of the rescoring method to predict the relative binding affinities of inhibitors. In principle, the improved enrichment shown here, relative to a high-throughput docking program, reflects the improved estimation of relative binding affinities, at least for a subset of the known inhibitors. Nonetheless, it is clear that binding affinity estimation is a much more challenging goal.

ACKNOWLEDGMENT

The California Institute for Quantitative Biomedical Research (QB3) is acknowledged for financial support to N.H. J.J.I. is supported by NIH Grant GM71896. The work of C.K. and M.P.J. is supported by NIH Grants GM071790, AI035707, and GM56531 and startup funds provided by HHMI Biomedical Research Support Program Grant #5300246 to the UCSF School of Medicine. The Center of Computational Proteomics Research (CCPR) in UCSF is thanked for computational support, MDL Inc. is thanked for providing the MDDR database and ISIS software (to B. Shoichet), the Shoichet lab is thanked for making the computers on which this software was installed available for this work, and Schrödinger Inc. is thanked for use of the program IMPACT. M.P.J. is a member of the Scientific Advisory Board of Schrödinger Inc.

Supporting Information Available: Figure S1 presents the enrichment plots for PNP obtained from docking with different charge states assigned to the cofactor phosphate group. Figure S2 presents the enrichment plots for the enzymes as a function of the percent of the database that is subjected to rescoring. This material is available free of charge via the Internet at <http://pubs.acs.org>.

REFERENCES AND NOTES

- Shoichet, B. K.; McGovern, S. L.; Wei, B.; Irwin, J. J. Lead discovery using molecular docking. *Curr. Opin. Chem. Biol.* **2002**, *6*, 439–446.
- Walters, W. P.; Stahl, M. T.; Murcko, M. A. Virtual screening – an overview. *Drug Discovery Today* **1998**, *3*, 160–178.
- Shoichet, B. K. Virtual screening of chemical libraries. *Nature* **2004**, *432*, 862–865.
- Shoichet, B. K.; Stroud, R. M.; Santi, D. V.; Kuntz, I. D.; Perry, K. M. Structure-based discovery of inhibitors of thymidylate synthase. *Science* **1993**, *259*, 1445–1450.
- Chen, I.-J.; Neamati, N.; Nicklaus, M. C.; Orr, A.; Anderson, L.; Barchi, J. J. J.; Kelley, J. A.; Pommier, Y.; MacKerell, A. D. J. Identification of HIV-1 integrase inhibitors via three-dimensional database searching using ASV and HIV-1 integrases as targets. *Bioorg. Med. Chem.* **2000**, *8*, 2385–2398.
- Deborah, A. L.; Willian, V. M.; Linda, K. J. Application of Virtual screening tools to a protein–protein interaction: database mining studies on the growth hormone. *Med. Chem. Res.* **1999**, *9*, 579–591.
- Doman, T. N.; McGovern, S. L.; Witherbee, B. J.; Kasten, T. P.; Kurumbail, R.; Stallings, W. C.; Connolly, D. T.; Shoichet, B. K. Molecular Docking and High-Throughput Screening for Novel Inhibitors of Protein Tyrosine Phosphatase-1B. *J. Med. Chem.* **2002**, *45*, 2213–2221.
- Powers, R. A.; Morandi, F.; Shoichet, B. K. Structure-based discovery of a novel, noncovalent inhibitor of AmpC beta-lactamase. *Structure* **2002**, *10*, 1013–1023.
- Peng, H.; Huang, N.; Qi, J.; Xie, P.; Xu, C.; Wang, J.; Yang, C. Identification of novel inhibitors of BCR-ABL tyrosine kinase via virtual screening. *Bioorg. Med. Chem. Lett.* **2003**, *13*, 3693–3699.
- Schapiro, M.; Abagyan, R.; Totrov, M. Nuclear hormone receptor targeted virtual screening. *J. Med. Chem.* **2003**, *46*, 3045–3059.
- Huang, N.; Nagarsekar, A.; Xia, G.; Hayashi, J.; MacKerell, A. D., Jr. Identification of Non-Phosphate-Containing Small Molecular Weight Inhibitors of the Tyrosine Kinase p56 Lck SH2 Domain via in Silico Screening against the pY + 3 Binding Site. *J. Med. Chem.* **2004**, *47*, 3502–3511.
- Alvarez, J. C. High-throughput docking as a source of novel drug leads. *Curr. Opin. Chem. Biol.* **2004**, *8*, 1–6.
- Oprea, T. I.; Matter, H. Integrating virtual screening in lead discovery. *Curr. Opin. Chem. Biol.* **2004**, *8*, 1–10.
- Kuntz, I. D.; Blaney, J. M.; Oatley, S. J.; Langridge, R.; Ferrin, T. E. A geometric approach to macromolecule–ligand interactions. *J. Mol. Biol.* **1982**, *161*, 269–288.
- Jones, G.; Willett, P.; Glen, R. C.; Leach, A. R.; Taylor, R. Development and validation of a genetic algorithm for flexible docking. *J. Mol. Biol.* **1997**, *267*, 727–748.
- Rarey, M.; Kramer, B.; Lengauer, T.; Klebe, G. A fast flexible docking method using an incremental construction algorithm. *J. Mol. Biol.* **1996**, *261*, 470–489.
- Friesner, R. A.; Banks, J. L.; Murphy, R. B.; Halgren, T. A.; Klicic, J. J.; Mainz, D. T.; Repasky, M. P.; Knoll, E. H.; Shelley, M.; Perry, J. K.; Shaw, D. E.; Francis, P.; Shenkin, P. S. Glide: a new approach for rapid, accurate docking and scoring. 1. Method and assessment of docking accuracy. *J. Med. Chem.* **2004**, *47*, 1739–1749.
- Muegge, I.; Martin, Y. C. A general and fast scoring function for protein–ligand interactions: a simplified potential approach. *J. Med. Chem.* **1999**, *42*, 791–804.
- DeWitte, R.; Shakhnovich, E. SMOG: de novo design method based on simple, fast and accurate free energy estimates. 1. Methodology and supporting evidence. *J. Am. Chem. Soc.* **1996**, *118*, 11733–11744.
- Bissantz, C.; Folkers, G.; Rognan, D. Protein-based virtual screening of chemical databases. 1. Evaluation of different docking/scoring combinations. *J. Med. Chem.* **2000**, *43*, 4759–4767.
- Charifson, P. S.; Corkery, J. J.; Murcko, M. A.; Walters, W. P. Consensus scoring: A method for improved hit rates from docking databases of three-dimensional structures into proteins. *J. Med. Chem.* **1999**, *42*, 5100–5109.
- Clark, R. D.; Strizhev, A.; Leonard, J. M.; Blake, J. F.; Matthew, J. B. Consensus scoring for ligand/protein interactions. *J. Mol. Graph. Modell.* **2002**, *20*, 281–295.
- Stahl, M.; Rarey, M. Detailed analysis of scoring functions for virtual screening. *J. Med. Chem.* **2001**, *44*, 1035–1042.
- Perola, E.; Walters, W. P.; Charifson, P. S. A detailed comparison of current docking and scoring methods on systems of pharmaceutical relevance. *Proteins* **2004**, *56*, 235–249.
- Warr, A. Virtual High-Throughput Screening: Computational Tools for Drug Discovery and Design. *Drug Discovery Des.* **2001**, *10*, 1–16.
- Hoffmann, D.; Kramer, B.; Washio, T.; Steinmetzer, T.; Rarey, M.; Lengauer, T. Two-stage method for protein–ligand docking. *J. Med. Chem.* **1999**, *42*, 4422–4433.
- Floriano, W. B.; Vaidehi, N.; Zamanakos, G.; Goddard, W. A., III. Hierarchical docking protocol for virtual ligand screening of large-molecule databases. *J. Med. Chem.* **2004**, *47*, 56–71.
- Verkhivker, G. M.; Bouzida, D.; Gehlhaar, D. K.; Rejto, P. A.; Arthurs, S.; Colson, A. B.; Freer, S. T.; Larson, V.; Luty, B. A.; Marrone, T.; Rose, P. W. Deciphering common failures in molecular docking of ligand–protein complexes. *J. Comput.-Aided Mol. Des.* **2000**, *14*, 731–751.
- Ferrara, P.; Gohlke, H.; Price, D. J.; Klebe, G.; Brooks, C. L., III. Assessing scoring functions for protein–ligand interactions. *J. Med. Chem.* **2004**, *47*, 3032–3047.
- Kuntz, I. D.; Meng, E. C.; Shoichet, B. K. Structure-based molecular design. *Acc. Chem. Res.* **1994**, *27*, 117–123.

- (31) Pearlman, D. A.; Charifson, P. S. Are free energy calculations useful in practice? A comparison with rapid scoring functions for the p38 MAP kinase protein system. *J. Med. Chem.* **2001**, *44*, 3417–3423.
- (32) Wang, J.; Morin, P.; Wang, W.; Kollman, P. A. Use of MM-PBSA in reproducing the binding free energies to HIV-1 RT of TIBO derivatives and predicting the binding mode to HIV-1 RT of efavirenz by docking and MM-PBSA. *J. Am. Chem. Soc.* **2001**, *123*, 5521–5230.
- (33) Pearlman, D. A.; Charifson, P. S. Improved scoring of ligand-protein interactions using OWFEG free energy grids. *J. Med. Chem.* **2001**, *44*, 502–511.
- (34) Chang, C. E.; Gilson, M. K. Free energy, entropy, and induced fit in host-guest recognition: calculations with the second-generation mining minima algorithm. *J. Am. Chem. Soc.* **2004**, *126*, 13156–13164.
- (35) Kalyanaraman, C.; Bernacki, K.; Jacobson, M. P. Virtual Screening against Highly Charged Active Sites: Identifying Substrates of Alpha-Beta Barrel Enzymes. *Biochemistry* **2005**, *44*, 2059–2071.
- (36) McGovern, S. L.; Shoichet, B. K. Information decay in molecular docking screens against holo, apo, and modeled conformations of enzymes. *J. Med. Chem.* **2003**, *46*, 2895–2907.
- (37) Connolly, M. L. Solvent-accessible surfaces of proteins and nucleic acids. *Science* **1983**, *221*, 709–713.
- (38) Ferrin, T. E.; Huang, C. C.; Jarvis, L. E.; Langridge, R. The MIDAS display system. *J. Mol. Graphics* **1988**, *6*, 13–27.
- (39) SYBYL; 6.7 ed.; Tripos Associates: St. Louis, MO.
- (40) Shoichet, B. K.; Kuntz, I. D. Matching chemistry and shape in molecular docking. *Protein Eng.* **1993**, *6*, 723–732.
- (41) Meng, E. C.; Shoichet, B. K.; Kuntz, I. D. Automated docking with grid-based energy evaluation. *J. Comput. Chem.* **1992**, *13*, 505–524.
- (42) Nicholls, A.; Honig, B. A Rapid Finite-Difference Algorithm, Utilizing Successive over-Relaxation to Solve the Poisson-Boltzmann Equation. *J. Comput. Chem.* **1991**, *12*, 435–445.
- (43) MDDR, 2000.2 ed.; MDL Inc.: San Leandro, CA.
- (44) Lorber, D. M.; Shoichet, B. K. Flexible ligand docking using conformational ensembles. *Protein Sci.* **1998**, *7*, 938–950.
- (45) Abagyan, R. A.; Totrov, M. M.; Kuznetsov, D. A. ICM: a new method for structure modeling and design. *J. Comput. Chem.* **1994**, *14*, 488–506.
- (46) McMartin, C.; Bohacek, R. S. QXP: Powerful, rapid computer algorithms for structure-based drug design. *J. Comput.-Aided Mol. Des.* **1997**, *11*, 333–344.
- (47) Morris, G. M.; Goodsell, D. S.; Halliday, R.; Huey, R.; Hart, W. E.; Belew, R. K.; Olson, A. J. Automated docking using a Lamarckian genetic algorithm and an empirical binding free energy function. *J. Comput. Chem.* **1998**, *19*, 1639–1662.
- (48) Wei, B. Q.; Baase, W. A.; Weaver, L. H.; Matthews, B. W.; Shoichet, B. K. A model binding site for testing scoring functions in molecular docking. *J. Mol. Biol.* **2002**, *322*, 339–355.
- (49) Jacobson, M. P.; Kaminski, G. A.; Friesner, R. A.; Rapp, C. A. Force field validation using protein side chain prediction. *J. Phys. Chem. B* **2002**, *106*, 11673–11680.
- (50) Jacobson, M. P.; Pincus, D. L.; Rapp, C. S.; Day, T. J.; Honig, B.; Shaw, D. E.; Friesner, R. A. A hierarchical approach to all-atom protein loop prediction. *Proteins* **2004**, *55*, 351–367.
- (51) Li, X.; Jacobson, M. P.; Friesner, R. A. High-resolution prediction of protein helix positions and orientations. *Proteins* **2004**, *55*, 368–382.
- (52) Jorgensen, W. L.; Maxwell, D. S.; Tirado-Rives, J. Development and Testing of the OPLS All-Atom Force Field on Conformational Energetics and Properties of Organic Liquids. *J. Am. Chem. Soc.* **1996**, *118*, 11225–11236.
- (53) Kaminski, G. A.; Friesner, R. A.; Tirado-Rives, J.; Jorgensen, W. L. Evaluation and Reparametrization of the OPLS-AA Force Field for Proteins via Comparison with Accurate Quantum Chemical Calculations on Peptides. *J. Phys. Chem. B* **2001**, *105*, 6474–6487.
- (54) Ghosh, A.; Rapp, C. S.; Friesner, R. A. Generalized Born model based on a surface integral formulation. *J. Phys. Chem. B* **1998**, *102*, 10983–10990.
- (55) Gallicchio, E.; Zhang, L. Y.; Levy, R. M. The SGB/NP hydration free energy model based on the surface generalized born solvent reaction field and novel nonpolar hydration free energy estimators. *J. Comput. Chem.* **2002**, *23*, 517–529.
- (56) Zhu, K.; Friesner, R. A.; Jacobson, M. P. Multiscale Optimization of a Truncated Newton Minimizer and Application to Proteins and Protein-Ligand Complexes. *J. Comput. Chem.* **2005**, submitted.
- (57) Xie, D. X.; Schlick, T. Efficient implementation of the truncated-Newton algorithm for large-scale chemistry applications. *SIAM J. Optimization* **1999**, *10*, 132–154.
- (58) Tuckerman, M.; Berne, B. J.; Martyna, G. J. Reversible Multiple Time Scale Molecular-Dynamics. *J. Chem. Phys.* **1992**, *97*, 1990–2001.
- (59) IMPACT, v. 2.7; Schrodinger Inc.: New York.
- (60) Jacobsson, M.; Liden, P.; Stjernschantz, E.; Boström, H.; Norinder, U. Improving structure-based virtual screening by multivariate analysis of scoring data. *J. Med. Chem.* **2003**, *46*, 5781–5789.
- (61) Toledo-Sherman, L.; Desouza, L.; Hosfield, C.; Liao, L.; Boutillier, K.; Taylor, P.; Climie, S.; McBroom-Cerajewski, L.; Moran, M. New Targets for an Old Drug—A Chemical Proteomics Approach to Unraveling the Molecular Mechanism of Action of Methotrexate. *Clin. Proteomics* **2004**, *1*, 45–68.
- (62) Meng, E. C.; Kuntz, I. D.; Abraham, D. J.; Kellogg, G. E. Evaluating docked complexes with HINT exponential functional and empirical atomic hydrophobicities. *J. Comput.-Aided Mol. Des.* **1994**, *8*, 299–306.
- (63) Zou, X.; Sun, Y.; Kuntz, I. D. Inclusion of Solvation in Ligand Binding Free Energy Calculations Using the Generalized-Born Model. *J. Am. Chem. Soc.* **1999**, *121*, 8033–8043.
- (64) Bolin, J. T.; Filman, D. J.; Matthews, D. A.; Hamlin, R. C.; Kraut, J. Crystal structures of *Escherichia coli* and *Lactobacillus casei* dihydrofolate reductase refined at 1.7 Å resolution. I. General features and binding of methotrexate. *J. Biol. Chem.* **1982**, *257*, 13650–13662.
- (65) Greasley, S. E.; Yamashita, M. M.; Cai, H.; Benkovic, S. J.; Boger, D. L.; Wilson, I. A. New insights into inhibitor design from the crystal structure and NMR studies of *Escherichia coli* GAR transformylase in complex with beta-GAR and 10-formyl-5,8,10-trideazafoolic acid. *Biochemistry* **1999**, *38*, 16783–16793.
- (66) Urzhumtsev, A.; Tete-Favier, F.; Mitschler, A.; Barbanton, J.; Barth, P.; Urzhumtseva, L.; Biellmann, J. F.; Podjarny, A.; Moras, D. A ‘specificity’ pocket inferred from the crystal structures of the complexes of aldose reductase with the pharmaceutically important inhibitors tolrestat and sorbinil. *Structure* **1997**, *5*, 601–612.
- (67) White, A. W.; Almasy, R.; Calvert, A. H.; Curtin, N. J.; Griffin, R. J.; Hostomsky, Z.; Maegley, K.; Newell, D. R.; Srinivasan, S.; Golding, B. T. Resistance-modifying agents. 9. Synthesis and biological properties of benzimidazole inhibitors of the DNA repair enzyme poly(ADP-ribose) polymerase. *J. Med. Chem.* **2000**, *43*, 4084–4097.
- (68) Fedorov, A.; Shi, W.; Kicska, G.; Fedorov, E.; Tyler, P. C.; Furneaux, R. N.; Hanson, J. C.; Gainsford, G. J.; Larese, J. Z.; Schramm, V. L.; Almo, S. C. Transition state structure of purine nucleoside phosphorylase and principles of atomic motion in enzymatic catalysis. *Biochemistry* **2001**, *40*, 853–860.
- (69) Turner, M. A.; Yuan, C. S.; Borchardt, R. T.; Hershfield, M. S.; Smith, G. D.; Howell, P. L. Structure determination of selenomethionyl S-adenosylhomocysteine hydrolase using data at a single wavelength. *Nat. Struct. Biol.* **1998**, *5*, 369–376.
- (70) Krishnan, R.; Zhang, E.; Hakansson, K.; Arni, R. K.; Tulinsky, A.; Lim-Wilby, M. S.; Levy, O. E.; Semple, J. E.; Brunck, T. K. Highly selective mechanism-based thrombin inhibitors: structures of thrombin and trypsin inhibited with rigid peptidyl aldehydes. *Biochemistry* **1998**, *37*, 12094–12103.
- (71) Dvir, H.; Wong, D. M.; Harel, M.; Barril, X.; Orozco, M.; Luque, F. J.; Munoz-Torrero, D.; Camps, P.; Rosenberry, T. L.; Silman, I.; Sussman, J. L. 3D structure of *Torpedo californica* acetylcholinesterase complexed with huprine X at 2.1 Å resolution: kinetic and molecular dynamic correlates. *Biochemistry* **2002**, *41*, 2970–2981.
- (72) Kamb, A.; Finer-Moore, J.; Calvert, A. H.; Stroud, R. M. Structural basis for recognition of polyglutamyl folates by thymidylate synthase. *Biochemistry* **1992**, *31*, 9883–9890.
- (73) Humphrey, W.; Dalke, A.; Schulten, K. VMD—visual molecular dynamics. *J. Mol. Graphics* **1996**, *14*, 33–38.