

## *AutoDock*-based incremental docking protocol to improve docking of large ligands

Ankur Dhanik\*, John S McMurray<sup>†</sup> and Lydia Kavrakı\*<sup>‡</sup>

\*Department of Computer Science, <sup>‡</sup>Department of Bioengineering  
Rice University, Houston, Texas, 77005, Email: {ankurd, kavrakı}@rice.edu

<sup>†</sup>Department of Experimental Therapeutics  
UT MD Anderson Cancer Center, Houston, Texas, 77054, Email: jmcmurra@mdanderson.org

**Abstract**—It is well known that computer-aided docking of large ligands, with many rotatable bonds, is extremely difficult. *AutoDock* is a widely used docking program that can dock small ligands, with upto 5 or 6 rotatable bonds, accurately and quickly. Docking of larger ligands, however, is not very accurate and is computationally expensive. In this paper we present an *AutoDock*-based incremental docking protocol which docks a large ligand to its target protein in increments. A fragment of the large ligand is first chosen and then docked. Best docked conformations are incrementally grown and docked again, and this process is repeated until all the atoms of the ligand are docked. Each docking operation is performed using *AutoDock*. However, in each docking operation only a small number of rotatable bonds are allowed to rotate. We did a systematic docking study on a dataset of 73 protein-ligand complexes derived from the core set of PDBbind database. The number of rotatable bonds in the ligands vary from 7 to 30. Docking experiments were done to evaluate the docking performance of the incremental protocol in comparison to *AutoDock*'s standard protocol. Results from the study show that, on average over the dataset, docking of large ligands using our incremental protocol is 23-fold computationally faster than docking using *AutoDock*'s standard protocol and also has comparable or better accuracy. We propose that, for docking large ligands, our incremental protocol can be used as an alternative to *AutoDock*'s standard protocol.

**Keywords**-computer-aided docking; docking programs; fast docking; large flexible ligands; *AutoDock*; PDBbind; drug discovery; drug design; protein-ligand complexes; high-dimensional conformation spaces

### I. INTRODUCTION

Computer-aided docking is an important tool for gaining understanding of the binding interactions between a ligand and its target receptor (mainly, a protein) [1], [2]. Over the years, several docking programs have been developed (see representative examples [3]–[11]). Typically, docking programs explore the conformation space of the ligand to generate a conformation bound to the target protein. A conformation is usually defined by the torsional degrees of freedom around the rotatable bonds of the ligand and the translation and orientation degrees of freedom. In general, the docking programs are fairly accurate and computationally fast when docking small ligands with 5 or 6 rotatable bonds [12], [13]. However, accurate and fast docking of large ligands with many rotatable bonds is still very challenging mainly because of the high-dimensionality of the conforma-

tion spaces which need to be explored to dock such large ligands.

Our strategy for docking large ligands derives inspiration from fragment-based docking methods. Some of the fragment-based methods place several fragments of a ligand in the binding cavity and then construct the full ligand from the best placed fragments, while other methods anchor a fragment of the ligand in the binding site and then build the full ligand incrementally. Some common docking programs that utilize such fragment-based approaches include LUDI [14], FlexX [6], GROWMOL [15], HOOK [16], Q-fit [17], and SURFLEX [8]. Multistep docking strategies have also been employed for flexible ligand docking. For example, in docking approach by Wang and coworkers [18], the conformation space corresponding to the ligand's rotatable bonds and that corresponding to the rigid body (translation, orientation) degrees of freedom are explored separately. The exploration of the conformation space corresponding to the unbound ligand's rotatable bonds is done using a multistep approach where the ligand is explored a few rotatable bonds at a time. Docking programs such as DOCK [19] and MS-DOCK [20] use a similar multistep approach for exploring the ligand conformations.

In this paper, we present an incremental protocol for docking large ligands where at each increment the docking operation is done using *AutoDock* [7], [21]. *AutoDock* is an excellent non-commercial docking program that is widely used. It employs a stochastic Lamarckian genetic algorithm for computing ligand conformations and simultaneously minimizing its scoring function which approximates the thermodynamic stability of the ligand bound to the target protein. *AutoDock* performs well when docking small ligands [13]; it is accurate as well as computationally fast. For docking large ligands, *AutoDock* recommends increasing the limit on the maximum number of energy evaluations (parameter called *ga\_num\_evals*). Increasing the limit essentially amounts to a more exhaustive search for the docked conformation in the conformation space. For the large ligands considered in this paper, docking with the recommended parameter settings results in modest gains in the accuracy while significantly increasing computational time. In our docking protocol, instead of exploring a very high-dimensional conformation space of the large ligand,

we harness the strengths of *AutoDock* in a way such that it always explores low-dimensional subspaces of the conformation space.

Starting from an initial fragment of the ligand, our protocol repeatedly docks and grows fragments of the ligand. Each docking operation is done using *AutoDock*. First an initial fragment of the ligand, composed of a small number of rotatable bonds and atoms that are directly rotated by them, is chosen and all the rotatable bonds are set active. Only the active bonds are allowed to rotate in each docking operation which ensures that only the low-dimensional subspaces of the full conformation space of the ligand are explored. The initial fragment is docked and the lowest-scoring docked conformations are selected. The selected docked conformations are then grown by adding a few more rotatable bonds and corresponding atoms. The grown fragments are then docked but only a few of the rotatable bonds are set active. The best docked conformations are grown and docked again until all the atoms of the ligand are docked. We present the details of our incremental docking protocol in the Methods section.

Using our protocol, we performed a systematic docking study on a dataset of 73 protein-ligand complexes (7 to 30 rotatable bonds in the ligands) from the core set of the PDBbind database [22]. Docking experiments were done using our protocol and *AutoDock*'s standard protocol in recommended parameter settings. The docking performance was mainly evaluated on the basis of docking accuracy and computational cost. Through the study we show that our protocol results in a significantly faster docking performance as compared to *AutoDock*'s standard protocol and the average docking accuracy is comparable or better.

## II. INCREMENTAL DOCKING PROTOCOL

We present an incremental protocol for docking large ligands using *AutoDock* [7], [21]. In a typical docking operation done using *AutoDock*'s standard protocol, the conformation space of the ligand is explored to generate a docked conformation that is bound to the target protein and has minimum binding energy (computed using *AutoDock*'s scoring function). The conformation space is defined by the set of torsion angles corresponding to all rotatable bonds in the ligand and the translation and orientation degrees of freedom of the ligand. Therefore, in the case of a large ligand, with many rotatable bonds, the exploration of the conformation space for the minimum-energy docked conformation becomes very challenging. *AutoDock* recommends docking the large ligand by using the standard protocol with parameter setting that results in a more exhaustive exploration of the conformational space. Our protocol, on the other hand, docks the large ligand incrementally such that, at each increment, only a subspace of the conformation space is explored. We present a brief overview of our incremental docking protocol followed by

a detailed description.

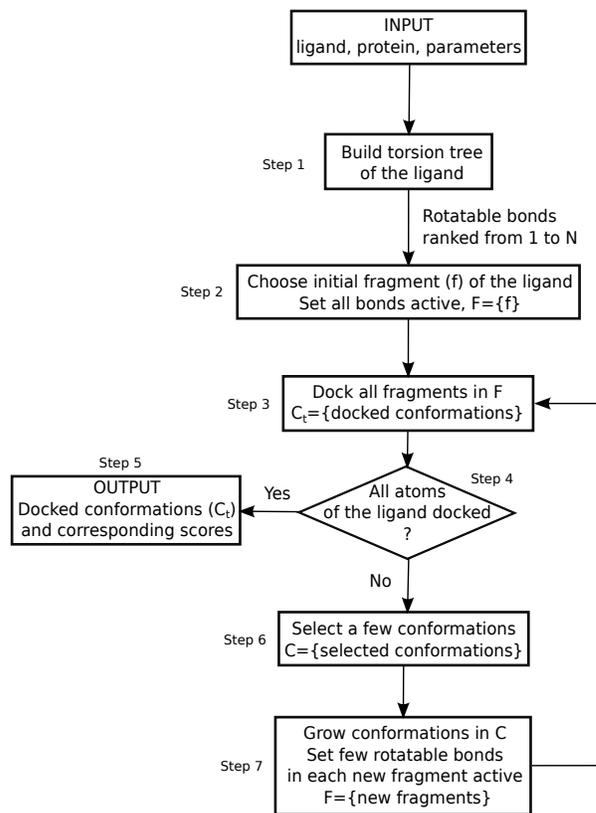


Figure 1. The flowchart shows the steps involved in our incremental docking protocol.

**Overview** Given a large ligand and a target protein, first an initial fragment of the ligand, composed of a small number of rotatable bonds and atoms directly rotated by the bonds, is chosen. All rotatable bonds in the initial fragment are set active. Only the active bonds are allowed to rotate and the fragment is docked to the target protein using *AutoDock*. Some of the docked conformations are then selected and grown by adding a few more bonds as well as atoms that are directly rotated by the new rotatable bonds. A small number of rotatable bonds in each grown conformation are set active and the conformations are docked. The docked conformations are grown and docked again, until all the atoms in the ligand are docked. Note that, in each docking operation, *AutoDock* explores a small number of rotatable bonds as well as the translation and orientation degrees of freedom, i.e., a subspace of the conformation space.

**Description** Our incremental protocol consists of several steps as shown in Figure 1. The protocol accepts as input:

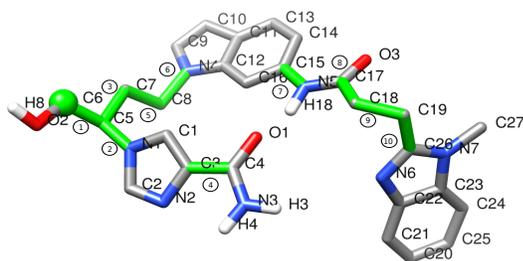


Figure 2. The top figure shows the ligand from a protein-ligand complex deposited in the PDB (1NDZ). The ligand has 10 rotatable bonds shown as green sticks. The torsion tree shown in the bottom figure is rooted at the carbon atom C6 (green sphere). An edge of the tree represents a rotatable bond and a node represents a list of the atoms that are directly rotated by the bond. The tree is traversed in a breadth-first fashion and the edges are ranked (circled numbers) by the order in which they were visited.

- a ligand structure with bonds labeled as rotatable and non-rotatable,
- a target protein structure to which the ligand has to be docked,
- the center and dimensions of the *AutoDock* grid that encompasses the binding pocket in the protein,
- $ga\_run$ , number of conformations that are output after a docking operation done using *AutoDock*,
- $ga\_num\_evals$ , maximum number of energy evaluations that are permitted by *AutoDock* when searching for the minimum of the binding energy,
- $n_s$  (maxrot size), maximum number of rotatable bonds that are allowed to rotate and are, therefore, explored in each docking operation of our protocol, and
- $k$ , number of docked conformations that are selected after a docking operation.

Note that the parameters (c)-(e) are required by *AutoDock* for a standard docking operation and the parameter names

are thus borrowed from *AutoDock* user guide [23]. The details of the steps in the incremental protocol that we propose are as follows.

*Step 1* A torsion tree is constructed where a node represents a list of atoms and an edge represents a rotatable bond of the ligand. The tree is used for ranking the rotatable bonds, and the ranked bonds and the tree nodes are utilized in the next steps of the protocol. To construct the tree, first a root atom is selected. The root atom is selected from the set of all heavy atoms in the ligand such that the initial fragment derived from the torsion tree has the largest number of hydrogen bond donors and acceptors. The idea is to identify an initial fragment of the ligand that is likely to be involved in many binding interactions, thereby increasing its docking accuracy. The root atom can also be selected such that the initial fragment contains atoms that are involved in already known binding interactions, or it can be selected randomly. In the docking study described later, we show the impact of the choice of the root atom on the docking performance of our protocol. Let  $S_1$  be a set that contains the root atom and atoms in the ligand that are connected to the root atom through a sequence of non-rotatable bonds. Let  $S_2$  be a set that contains any atom that is not contained in  $S_1$  and is bonded to at least one atom in  $S_1$ . All atoms in  $S_1$  and  $S_2$  are inserted into the root node. For each rotatable bond connected to the root atom, an edge and a node is added to the tree. The edge corresponds to the rotatable bond and the new node corresponds to the list of atoms that are directly rotated by the bond.

The leaf nodes of the tree are recursively expanded, until all atoms and rotatable bonds in the ligand have been appended to the tree. Figure 2 shows the torsion tree for the ligand from a protein-ligand complex deposited in the PDB (1NDZ). Suppose there are  $N$  edges in the torsion tree. Starting from the root node, the torsion tree is traversed in a breadth-first fashion and the edges are ranked from 1 to  $N$  in the order that they were visited.

*Step 2* The initial fragment (say,  $f$ ), first to be docked with our protocol, is composed of the atoms in the torsion tree nodes that are connected by the edges ranked from 1 to  $r_u$ , where  $r_u = \min(n_s, N)$ , and  $n_s$  is the input maxrot size. All the rotatable bonds in the initial fragment are set active and a set  $F$  is initialized such that  $F = \{f\}$ .

*Step 3-5* *AutoDock* energy grid maps are computed as is typical in docking using *AutoDock*. A map contains, for each grid point location, values of interaction energies between the protein and the various atom types in the fragment. The atom types are assigned by *AutoDock*. Since each fragment in  $F$  contains atoms of the same atom types, the maps are computed only once in this step.

Each fragment in  $F$  is then docked to the target protein. Note that, by default, at the start of a docking operation *AutoDock* assigns random values to all the active rotatable bonds and the translation and orientation degrees of freedom.

A specified number ( $ga\_run$ ) of docked conformations and their *AutoDock* scores are obtained in each docking operation. If all the atoms in the ligand are docked, then the set of all docked conformations (say,  $C_t$ ) and corresponding scores are returned as the output of our docking protocol. While all the atoms in the ligand are not docked, then *Steps 6-7* are performed.

*Step 6-7* The docked conformations in  $C_t$  are ranked by their *AutoDock* scores such that a conformation with lower score is ranked higher. The  $k$  top-ranked docked conformations are selected and each conformation is grown by adding rotatable bonds ranked from  $r_u + 1$  to  $\min(r_u + \lfloor n_s/2 \rfloor, N)$  and all the atoms (derived from the torsion tree nodes) that are directly rotated by the added bonds. In each new fragment that is created by growing the docked conformations, the newly added rotatable bonds are set active. Few ( $\lfloor n_s/2 \rfloor$ ) of the already explored bonds are also set active because re-exploration of the bonds improves the docking accuracy. Thus, the rotatable bonds ranked from  $r_u - \lfloor n_s/2 \rfloor + 1$  to  $\min(r_u + \lfloor n_s/2 \rfloor, N)$  are set active, and all the other rotatable bonds are set non-active. The set  $F$  and parameter  $r_u$  are updated such that  $F = \{\text{new fragments}\}$ , and  $r_u = \min(r_u + \lfloor n_s/2 \rfloor, N)$ . *Steps 3-7* are repeated.

The incremental growth of the ligand from a protein-ligand complex deposited in the PDB (1NDZ) is shown in Figure 3. Since our protocol involves selecting  $k$  docked conformations at each increment, and then growing and docking them again, it lends itself very easily to a parallel implementation. Each of the  $k$  conformations in *Step 6-7* are, therefore, docked in parallel.

### III. DOCKING STUDY

The purpose of our docking study is twofold: (a) to determine how the selection of the root atom and the maxrot size ( $n_s$ ) affects docking performance of our protocol, and (b) compare the docking performance of our incremental protocol with the docking performance of *AutoDock*'s [7], [21] standard protocol. Here we establish a dataset of protein-ligand complexes for our docking study and describe various docking experiments. In each docking experiment, all of the ligands in the dataset were computationally docked to their target proteins.

*Dataset* The dataset is derived from the core set of PDBbind (v2007) database [22] that contains 210 protein-ligand complexes in total. For each complex, the structures of the protein and the ligand, as well as the experimentally measured binding affinity is available. In each complex, the bonds in the ligand are labeled as rotatable or non-rotatable. Any bond that is an amide bond, or is in a cycle, or rotates only hydrogen atoms is labeled non-rotatable, and the rest of the bonds are labeled rotatable. Out of the 210 complexes, 73 complexes have greater than 6 rotatable bonds in the ligand. We consider ligands in the 73 complexes as large ligands for

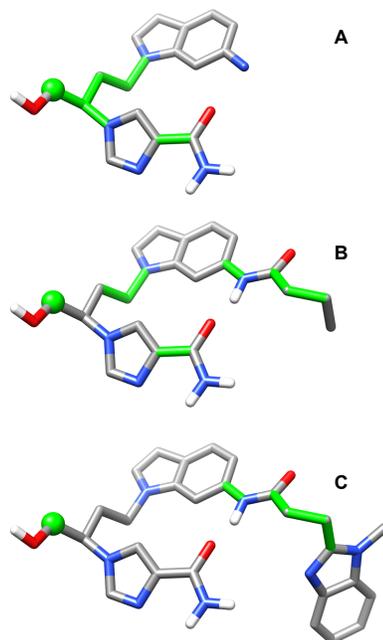


Figure 3. After the construction of the torsion tree (Figure 2), the ligand is incrementally grown and docked. For this illustration, the maxrot size ( $n_s$ ) is set to 6. (A) The first fragment has atoms that are contained in the nodes connected by the edges ranked from 1 to 6. Bonds corresponding to these edges are set active and the fragment is docked. Active bonds are shown as green sticks. (B) The docked conformations of the fragment are grown by adding atoms that are contained in the nodes connected by the edges ranked from 7 to 9. The 3 newly added bonds and the 3 bonds corresponding to the edges ranked from 4 to 6 are set active. The new fragments are docked again. (C) The fragment is grown again by adding atoms that are contained in the node connected by the edge ranked 10. The newly added bond and the bonds corresponding to the edges ranked from 7 to 9 are set active. The fragment is docked and we obtain docked conformations of the full ligand.

the purpose of our study and these 73 complexes form our dataset. The distribution of the number of rotatable bonds in the dataset is shown in Figure 4.

Each protein and ligand in the dataset is processed in the following manner. The ligand atoms are assigned atom types and Gasteiger charges [24]. All non-polar hydrogen atoms and lone-pair charges are removed. The charge of each removed non-polar hydrogen is added to the carbon atom to which it is bonded and the charge of each lone pair is added to the atom it is associated with. The protein atoms are also assigned atom types and Gasteiger charges. The non-polar hydrogen atoms and lone-pairs of the protein are processed in the same way as those of the ligand. Waters and non-standard residues are removed as well. Similar to some other docking studies [13], [25] that use the PDBbind database, structural optimization of the ligand and the protein was not performed.

*Root atom selection mode and maxrot size* For examining

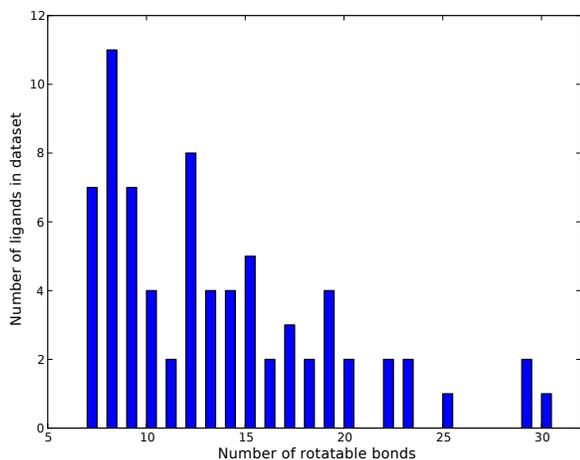


Figure 4. Distribution of number of rotatable bonds in the dataset that was used for the docking experiments.

the robustness of our proposed protocol, two different modes for the selection of the root atom, required for the construction of the torsion tree in our incremental protocol, were tested: (a) a heavy atom was picked such that the tree rooted at that atom results in the initial fragment with the largest number of hydrogen bond donors and acceptors (*std* mode), and (b) a heavy atom was picked randomly (*rand* mode). We also tested various maxrot sizes,  $n_s = 4, 6,$  and  $8$ . Using all combinations of selection modes and maxrot sizes, we performed six docking experiments using our protocol. The input parameter *ga\_num\_evals* was set to 250000, and *ga\_num* was set to 50 for the first docking operation and to 20 for subsequent docking operations. The input parameter *k* was set to 5.

**Docking performance comparison** For comparing the docking performance of our protocol and *AutoDock*’s standard protocol, we did two docking experiments. *AD<sup>inc</sup>* refers to the docking experiment done using our incremental protocol with  $n_s = 6$  and *std* mode for the selection of the root atom. Other input parameters to *AD<sup>inc</sup>* were same as in the six docking experiments discussed above. *AD* refers to the experiment with *ga\_num\_evals* set to 25 million and *ga\_run* set to 50. The settings in *AD* are the recommended settings for *AutoDock*’s standard protocol when docking large ligands and have been used previously for docking performance comparison [11].

In dockings done with *AutoDock*, the *AutoDock* grid was centered on the average coordinates of the atoms of the input ligand structure (obtained from the dataset) and the dimensions of the grid were chosen such that the grid box fully encompasses the input ligand structure. The size of each dimension of the grid box was extended by  $15\text{\AA}$ . The size of any dimension, if it was less than  $60\text{\AA}$ , was set to  $60\text{\AA}$

Table I  
ROOT ATOM SELECTION MODE AND MAXROT SIZE ( $n_s$ ).

Mode	$n_s$	$S_{CS}^{avg}$ (kcal/mol)	$R_{CS}^{avg}$ ( $\text{\AA}$ )	$S_{CR}^{avg}$ (kcal/mol)	$R_{CR}^{avg}$ ( $\text{\AA}$ )	$DT^{avg}$ (h)
<i>std</i>	4	-8.93	5.50	-6.84	3.49	41.57
<i>rand</i>	4	-8.99	5.73	-6.79	3.39	42.02
<i>std</i>	6	-9.04	5.06	-6.51	3.01	31.70
<i>rand</i>	6	-8.84	5.44	-6.56	2.99	31.58
<i>std</i>	8	-8.34	5.73	-6.17	3.15	25.26
<i>rand</i>	8	-8.27	5.76	-6.21	3.01	24.88

This table lists average *AutoDock* scores ( $S_{CS}^{avg}$ ,  $S_{CR}^{avg}$ ), average RMSD values ( $R_{CS}^{avg}$ ,  $R_{CR}^{avg}$ ), and average total docking times ( $DT^{avg}$ ).  $S_{CS}^{avg}$  and  $R_{CS}^{avg}$  correspond to the *Top-scoring* docked conformations.  $S_{CR}^{avg}$  and  $R_{CR}^{avg}$  correspond to the *Top-RMSD* docked conformations. In selection mode *std*, the root atom is selected using a heuristic and in *random* mode it is selected randomly from the set of all heavy atoms in the ligand. Different maxrot sizes are represented by  $n_s$ .

Table II  
DOCKING PERFORMANCE COMPARISON.

Experiment	$S_{CS}^{avg}$ (kcal/mol)	$R_{CS}^{avg}$ ( $\text{\AA}$ )	$S_{CR}^{avg}$ (kcal/mol)	$R_{CR}^{avg}$ ( $\text{\AA}$ )	$DT^{avg}$ (h)
<i>AD<sup>inc</sup></i>	-9.04	5.06	-6.51	3.01	31.70
<i>AD</i>	-9.66	5.17	-8.44	2.92	725.85

This table lists average *AutoDock* scores ( $S_{CS}^{avg}$ ,  $S_{CR}^{avg}$ ), average RMSD values ( $R_{CS}^{avg}$ ,  $R_{CR}^{avg}$ ), and average total docking times ( $DT^{avg}$ ).  $S_{CS}^{avg}$  and  $R_{CS}^{avg}$  correspond to the *Top-scoring* docked conformations.  $S_{CR}^{avg}$  and  $R_{CR}^{avg}$  correspond to the *Top-RMSD* docked conformations. *AD<sup>inc</sup>* refers to the experiment done using incremental protocol ( $n_s = 6$ , *std* selection mode) and *AD* (*ga\_num\_evals* = 25 million, *ga\_run* = 50) refer to the experiments done using *AutoDock*’s standard protocol.

[11]. The rest of the parameters of *AutoDock*, that are not explicitly mentioned, were set to their default values. Since *AutoDock* employs a stochastic algorithm, all experiments were repeated 5 times. Version 4.2 of *AutoDock* was used in all of the docking experiments. The docking study was done on a cluster at Rice University that has 192 computing nodes and 2304 (in total) processor cores, with each core running at 2.83 GHz.

## IV. RESULTS

Here we present the results from the various docking experiments described in Section III. Each docking of a ligand to its target protein produced multiple docked conformations of the ligand and corresponding docking scores. We computed root mean squared distance (RMSD) between the docked conformations and the input conformation of the ligand that is obtained from the structure of the protein-ligand complex in the dataset. The RMSD values were computed using only the heavy atoms of the ligand. We identified two docked conformations of each ligand: (a) a *Top-scoring* conformation (*CS*), and (b) a *Top-RMSD* conformation (*CR*). *CS* is the docked conformation that has the minimum score and the *CR* is the one with the lowest RMSD value. It is to be noted that since the dataset

contains structures of the ligands bound to their respective proteins, it is possible to compute the RMSD values. In the general case of ligands for which the structures of the bound conformations are not known, only the *Top-scoring* conformation of each ligand can be identified. The docking performance is, therefore, evaluated on the basis of the score ( $S_{CS}$ ) and RMSD value ( $R_{CS}$ ) associated with  $CS$  of each ligand, and the computational expense of docking ( $DT$ ). Better docking performance means lower values of  $R_{CS}$  and  $DT$ . We also analyze the score ( $S_{CR}$ ) and RMSD value ( $R_{CR}$ ) associated with  $CR$  of each ligand.

#### A. Root atom selection mode and maxrot size

Six docking experiments were performed to determine the effect of root atom selection mode (*rand*, *std*) and maxrot size on the docking performance ( $n_s = 4, 6, 8$ ) of our incremental protocol. The results from the six experiments are presented in Table I. The values of  $S_{CS}$ ,  $S_{CR}$ ,  $R_{CS}$ , and  $R_{CR}$  were obtained for the 73 (number of protein-ligand complexes in the dataset) dockings done in each experiment. Average of the values over the 73 dockings were computed for each experiment. Each experiment was repeated 5 times and Table I lists the average (over the 5 experiments) of the average values, denoted by  $S_{CS}^{avg}$ ,  $S_{CR}^{avg}$ ,  $R_{CS}^{avg}$ , and  $R_{CR}^{avg}$ .  $DT^{avg}$  represents the total computational expense of the 73 dockings averaged over the 5 repeated experiments.

Most accurate docking performance, i.e., lowest  $R_{CS}^{avg}$  was obtained in the experiment done with *std* selection mode and maxrot size  $n_s = 6$ . Experiments with  $n_s = 8$  were done in the lowest docking time ( $DT^{avg}$ ). The maxrot size determines number of docking operations that are performed in each docking using our protocol. A bigger maxrot size thus results in fewer docking operations per docking which means experiments with bigger maxrot size require lower docking times. Similarly, experiments with smaller maxrot size require higher docking times as in the case of experiments with  $n_s = 4$ .

The *Top-scoring* conformations were, in general, less accurate ( $R_{CS}^{avg} > R_{CR}^{avg}$  in all experiments) than *Top-RMSD* conformations. The docked conformation with the lowest score is not necessarily the one with the lowest RMSD value because the use of approximate scoring functions in docking programs results in inaccurate estimation of the thermodynamic stabilities of the complexes. Comparison of scores ( $S_{CR}^{avg} > S_{CS}^{avg}$ ) re-emphasizes that the conformations with the lowest RMSD values do not necessarily have the lowest scores.

#### B. Docking performance comparison

To compare the performance of our incremental protocol, we did docking experiments using *AutoDock*'s standard protocol as described in the Methods section. For each experiment, values of  $S_{CS}^{avg}$ ,  $S_{CR}^{avg}$ ,  $R_{CS}^{avg}$ ,  $R_{CR}^{avg}$ , and  $DT^{avg}$  were computed as above. The results from the experiments

are listed in Table II. The experiment done using our protocol took, on average, 31.70 hours to finish. It is to be noted that this is the total CPU time spent. Due to the parallel implementation of our protocol (see Methods section), the real time spent was much lower. Comparison of docking times  $DT^{avg}$  in Table II reveals that the docking time spent in the experiment done using our protocol ( $AD^{inc}$ ) is much lower than experiment *AD*. The *Top-scoring* conformations computed in  $AD^{inc}$  were overall more accurate (lower  $R_{CS}^{avg}$  value) than those computed in *AD*. As expected, in both experiments,  $R_{CR}^{avg}$  values were lower and  $S_{CR}^{avg}$  values were higher than corresponding  $R_{CS}^{avg}$  and  $S_{CS}^{avg}$  values, thus emphasizing the limitations of the scoring functions. For experiments  $AD^{inc}$  and *AD*, the medians of the 73  $R_{CS}$  values, corresponding to the 73 protein-ligand complexes in our dataset, were 4.66Å and 4.88Å respectively. The median values also reflected the better docking performance of our protocol.

Figures 5 and 6 show similar distribution of  $R_{CS}$  and  $R_{CR}$  RMSD values obtained in  $AD^{inc}$  and *AD*. In eval-

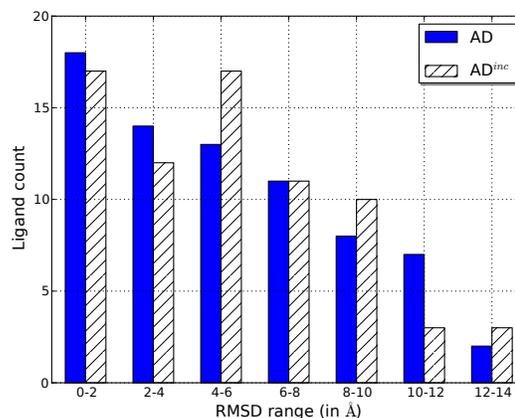


Figure 5. Distribution of  $R_{CS}$  RMSD values obtained in *AD* and  $AD^{inc}$ .

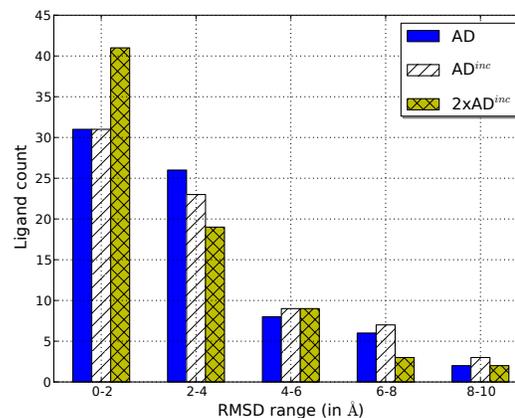


Figure 6. Distribution of  $R_{CR}$  RMSD values obtained in *AD*,  $AD^{inc}$  and a docking experiment that combined docking results from 2 runs of  $AD^{inc}$ .

uation of docking performances, a docked conformation that has a RMSD value of less than  $2.0\text{\AA}$  is considered accurate. The count of protein-ligand complexes for which  $R_{CS}$  is less than  $2.0\text{\AA}$  is low. However, in comparison, the count of protein-ligand complexes for which  $R_{CR}$  is less than  $2.0\text{\AA}$  is significantly higher which again demonstrates that *AutoDock*'s scoring function has limited success in picking the lowest RMSD conformation from a set of docked conformations.

The incremental docking protocol is clearly superior over the standard protocol as it requires 23-fold lower docking time to achieve accuracy that is comparable to the standard protocol. An advantage of faster docking time is that multiple runs of incremental docking protocol can be combined to increase the exploration of the conformation space of large ligands. For example, after combining the set of final docked conformations from 2 runs of incremental docking protocol, the count of protein-ligand complexes with  $R_{CR} < 2.0\text{\AA}$  increased by 32%. Figure 6 shows the improvement in the distribution of the  $R_{CR}$  values obtained after combining the set of final docked conformations from the 2 runs. Not only was there an improvement in the docked conformations, but the total docking time for the 2 runs was also an order of magnitude lower than the standard protocol.

## V. DISCUSSION

*AutoDock* is a popular non-commercial docking program that docks a ligand to its target protein and performs well (accurate and computationally fast) when the number of rotatable bonds in the ligand is small. Increase in the number of rotatable bonds in the ligand, however, severely affects the docking performance of *AutoDock*.

In this paper we propose an incremental docking protocol for docking large ligands that utilizes *AutoDock* for docking operations. Instead of searching for the docked conformation in the high-dimensional conformation space of a large ligand, the search is done in the subspaces of the conformation space. Starting from a carefully chosen initial fragment, the ligand is incrementally docked and grown until all of its atoms are docked. At each increment a few bonds are set active and are allowed to rotate. By limiting the number of active bonds (termed as maxrot size) to a small value, we ensure that in all the docking operations, the search for the docked conformation is done in a subspace of the conformation space.

We performed a systematic docking study with docking experiments on a dataset of 73 protein-ligand complexes (7 to 30 rotatable bonds in the ligands) from the core set of the PDBbind database. Six different experiments were done using our protocol to determine the effect of the choice of the root atom and maxrot size on the docking performance. The choice of the root atom leading to an initial fragment of the ligand with highest number of hydrogen bond donors and acceptors, and maxrot size of 6 resulted in the best docking

performance. To compare our protocol with *AutoDock*'s standard protocol, we did experiments using our protocol and *AutoDock*'s standard protocol in recommended parameter settings.

From the results of our docking study, it is clear that the accuracy of docking using our protocol is comparable to or better than the accuracy of docking using *AutoDock*'s standard protocol. The docking experiments demonstrate that our protocol is computationally fast, it is 23-fold faster than docking using *AutoDock*'s standard protocol. The parallel implementation of the protocol makes docking even more fast. Each docking operation in our protocol explores a subspace of the conformation space and is, therefore, fast. Even though our protocol consists of multiple docking operations, less computational time is spent in each individual operation and this results in lower docking time overall.

The average (over the dataset) RMSD value of the *Top-scoring* conformations computed in the experiment done using our protocol is  $5.06\text{\AA}$ . As shown in Figure 6, in the case of many protein-ligand complexes, our protocol computed conformations with  $\text{RMSD} < 2.0\text{\AA}$ . Comparison with Figure 5, however, shows that the *Top-scoring* conformations, as identified by *AutoDock*'s scoring function, are not necessarily the conformations with lowest RMSD. The need for improved scoring function is, thus, evident.

To further improve the docking protocol, we are investigating different root atom selection modes and ways of selecting conformations in *Step 6* of the protocol, e.g., selecting a combination of conformations with lowest scores as well as conformations representing the biggest clusters. The docked conformations computed by the protocol can be further refined. Molecular dynamics based refinement studies are under way and the strength of our protocol is that docked conformations required for the refinement can be quickly computed. Consensus docking programs [10], [26] that utilize multiple docking programs, such as *AutoDock* and others, can also benefit from the quick computation of docked conformations using our protocol. Such consensus docking programs can probably further improve the accuracy of docking large ligands. Thus, for docking large ligands, our incremental protocol can provide an excellent alternative to *AutoDock*'s standard protocol.

## ACKNOWLEDGMENT

This work has been supported in part by a training fellowship from the Keck Center Computational Cancer Biology Training Program of the Gulf Coast Consortia (CPRIT RP101489), The Texas Higher Education Coordinating Board (NHARP 01907), The National Cancer Institute (RO1 CA096652), The John and Ann Doerr Fund for Computational Biomedicine at Rice University, and the Data Analysis and Visualization Cyberinfrastructure funded by NSF under grant OCI-0959097.

## REFERENCES

- [1] A. C. Anderson, "The process of structure-based drug design," *Chemistry & Biology*, vol. 10, pp. 787–797, 2003.
- [2] G. Schneider, "Virtual screening: an endless staircase?" *Nature Reviews Drug Discovery*, vol. 9, no. 4, pp. 273–276, Apr. 2010.
- [3] I. Kuntz, "A geometric approach to macromolecule-ligand interactions," *Journal of Molecular Biology*, vol. 161, no. 2, pp. 269–288, Oct. 1982.
- [4] R. Abagyan, M. Totrov, and D. Kuznetsov, "ICM—a new method for protein modeling and design: applications to docking and structure prediction from the distorted native conformation," *Journal of Computational Chemistry*, vol. 15, no. 5, pp. 488–506, May 1994.
- [5] G. Jones, P. Willett, and R. C. Glen, "Molecular recognition of receptor sites using a genetic algorithm with a description of desolvation," *Journal of Molecular Biology*, vol. 245, no. 1, pp. 43–53, Jan. 1995.
- [6] M. Rarey, B. Kramer, T. Lengauer, and G. Klebe, "A fast flexible docking method using an incremental construction algorithm," *Journal of Molecular Biology*, vol. 261, no. 3, pp. 470–489, Aug. 1996.
- [7] G. M. Morris, D. S. Goodsell, R. S. Halliday, R. Huey, W. E. Hart, R. K. Belew, and A. J. Olson, "Automated docking using a Lamarckian genetic algorithm and an empirical binding free energy function," *Journal of Computational Chemistry*, vol. 19, no. 14, pp. 1639–1662, Nov. 1998.
- [8] A. N. Jain, "Morphological similarity: a 3D molecular similarity method correlated with protein-ligand recognition," *Journal of Computer-aided Molecular Design*, vol. 14, no. 2, pp. 199–213, Feb. 2000.
- [9] R. A. Friesner, J. L. Banks, R. B. Murphy, T. A. Halgren, J. J. Klicic, D. T. Mainz, M. P. Repasky, E. H. Knoll, M. Shelley, J. K. Perry, D. E. Shaw, P. Francis, and P. S. Shenkin, "Glide: a new approach for rapid, accurate docking and scoring. 1. Method and assessment of docking accuracy," *Journal of Medicinal Chemistry*, vol. 47, no. 7, pp. 1739–1749, Mar. 2004.
- [10] D. Plewczynski, "VoteDock: consensus docking method for prediction of protein-ligand interactions," *Journal of Computational Chemistry*, vol. 32, pp. 1–14, 2010.
- [11] O. Trott and A. J. Olson, "AutoDock Vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization, and multithreading," *Journal of Computational Chemistry*, vol. 31, pp. 455–461, 2010.
- [12] C. Hetényi and D. van der Spoel, "Efficient docking of peptides to proteins without prior knowledge of the binding site," *Protein Science*, vol. 11, no. 7, pp. 1729–1737, 2002.
- [13] D. Plewczynski, M. Lziniewski, R. Augustyniak, and K. Ginalska, "Can we trust docking results? Evaluation of seven commonly used programs on PDBbind database," *Journal of Computational Chemistry*, vol. 32, no. 1, pp. 742–755, 2011.
- [14] H.-J. Böhm, "The computer program LUDI: a new method for the de novo design of enzyme inhibitors," *Journal of Computer-Aided Molecular Design*, vol. 6, pp. 61–78, 1992.
- [15] R. S. Bohacek and C. McMartin, "Multiple highly diverse structures complementary to enzyme binding sites: results of extensive application of a de novo design method incorporating combinatorial growth," *Journal of the American Chemical Society*, vol. 116, no. 13, pp. 5560–5571, 1994.
- [16] M. B. Eisen, D. C. Wiley, M. Karplus, and R. E. Hubbard, "HOOK: a program for finding novel molecular architectures that satisfy the chemical and steric requirements of a macromolecule binding site," *Proteins: Structure, Function, and Bioinformatics*, vol. 19, no. 3, pp. 199–221, 1994.
- [17] R. M. Jackson, "Q-fit: a probabilistic method for docking molecular fragments by sampling low energy conformational space," *Journal of Computer-aided Molecular Design*, vol. 16, no. 1, pp. 43–57, Jan. 2002.
- [18] J. Wang, P. A. Kollman, and I. D. Kuntz, "Flexible ligand docking: a multistep strategy approach," *Proteins: Structure, Function, and Bioinformatics*, vol. 36, no. 1, pp. 1–19, Jul. 1999.
- [19] T. J. Ewing, S. Makino, A. G. Skillman, and I. D. Kuntz, "Dock 4.0: search strategies for automated molecular docking of flexible molecule databases," *Journal of Computer-Aided Molecular Design*, vol. 15, pp. 411–428, 2001.
- [20] N. Sauton, D. Lagorce, B. O. Villoutreix, and M. A. Miteva, "MS-DOCK: accurate multiple conformation generator and rigid docking protocol for multi-step virtual ligand screening," *BMC Bioinformatics*, vol. 9, p. 184, Jan. 2008.
- [21] G. M. Morris, R. Huey, W. Lindstrom, M. F. Sanner, R. K. Belew, D. S. Goodsell, and A. J. Olson, "AutoDock4 and AutoDockTools4: automated docking with selective receptor flexibility," *Journal of Computational Chemistry*, vol. 30, no. 16, pp. 2785–2791, 2009.
- [22] R. Wang, X. Fang, Y. Lu, C.-Y. Yang, and S. Wang, "The PDBbind database: methodologies and updates," *Journal of Medicinal Chemistry*, vol. 48, no. 12, pp. 4111–4119, 2005.
- [23] "Autodock 4.2 user guide." [Online]. Available: <http://autodock.scripps.edu/downloads/autodock-registration/autodock-4-2-download-page/>
- [24] J. Gasteiger and M. Marsili, "Iterative partial equalization of orbital electronegativity—a rapid access to atomic charges," *Tetrahedron*, vol. 36, no. 22, pp. 3219–3228, 1980.
- [25] X. Li, Y. Li, T. Cheng, Z. Liu, and R. Wang, "Evaluation of the performance of four molecular docking programs on a diverse set of protein-ligand complexes," *Journal of Computational Chemistry*, vol. 31, pp. 2109–2125, 2010.
- [26] R. Teramoto and H. Fukunishi, "Supervised consensus scoring for docking and virtual screening," *Journal of Chemical Information and Modeling*, vol. 47, no. 2, pp. 526–534, 2007.