Ligand Binding with OBPRM and User Input*

O. Burchan Bayazit Guang Song Nancy M. Amato
{burchanb,gsong,amato}@cs.tamu.edu
Department of Computer Science, Texas A&M University
College Station, TX 77843-3112

Abstract

*In this paper, we present a framework for studying ligand binding which is based on techniques recently developed in the robotics motion planning community. We are interested in locating binding sites on the protein for a ligand molecule. Our work investigates the performance of a fully automated motion planner, as well as the effects of supplementary user input collected using a haptic device. Our results applying an obstacle-based probabilistic roadmap motion planning algorithm (*OBPRM*) to some protein-ligand complexes are encouraging. The framework successfully identified potential binding sites for all complexes studied. We find that user input helps the planner, and a haptic device helps the user to understand the protein structure by enabling them to feel the difficult-to-visualize forces.*

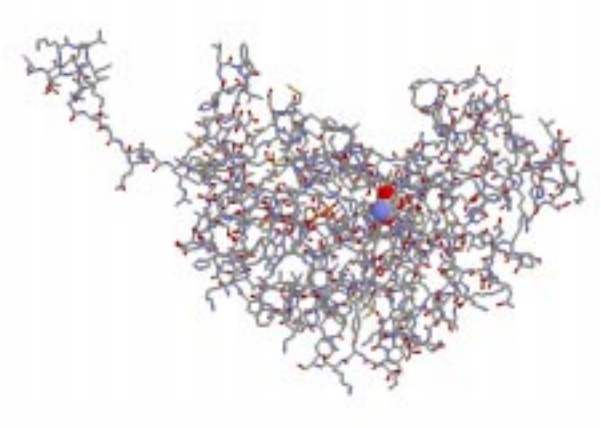

Figure 1: A protein in wireframe with a space-fill ligand. PDB file: 1LDM.pdb.

1 Introduction

Efficiency of a drug molecule is measured by its ability to find a specific position and orientation inside a protein. This process is called binding (or docking) and the drug molecule is often referred as a ligand. The binding configuration should satisfy some constraints based on geometric, electrostatic, and chemical reactions between the ligand and protein atoms. A good binding site should also be accessible to the ligand which must reach it from an outside location. This makes the path to the binding site important, and motivates the use of motion planning to study this problem.

Most researchers investigating automated docking methods simplify the problem by treating the ligand and protein molecules as rigid bodies. Our experiments show that this simplification negatively impacts our ability to identify the binding site. However, if the molecules are flexible (articulated bodies), the dimensionality of the problem becomes very high, making a deterministic search of the configuration space infeasible.

In this work, we apply *probabilistic roadmap* (PRM) motion planning methods to protein-ligand binding. PRMs have been extremely successful for problems with high-dimensional configuration spaces, and moreover, are very simple to apply, requiring only the ability to randomly generate points in C-space, and then test them for feasibility. Besides ligand binding [16], it has been applied also to study protein folding problems [17, 18] as well. The configuration of the ligand in the binding site has low potential energy, and so the usual PRM feasibility test (collision) is replaced by a test for low potential energy. In this study, we use an obstacle-based PRM called OBPRM [1] that generates configurations close to the surface of the protein. Since our goal, the binding site, is a local minima of the potential, we used a gradient-descent like algorithm to drive the roadmap nodes towards local minimum. To recognize potential binding sites, we construct 'local roadmaps' around (most of) the roadmap nodes and give each candidate node a score based on an analysis of its local roadmap. Our approach was able to generate configurations close to the true binding site in the three tested protein-ligand complexes. Two of these were studied in [16], where in one case their approach failed to generate a configuration in the binding site.

While our results are very encouraging, a fully automated approach suffers from the known problems of PRMs. The most significant of which is the *narrow passage* problem [3], so called because it is difficult for the planner to

*This research supported in part by NSF CAREER Award CCR-9624315 (with REU Supplement), NSF Grants IIS-9619850 (with REU Supplement), ACI-9872126, EIA-9975018, EIA-9805823, and EIA-9810937, and by the Texas Higher Education Coordinating Board under grant ARP-036327-017. Bayazit is supported in part by the Turkish Ministry of Education.

sample important configurations in a narrow C-space region. In molecular docking, the narrow passage becomes a passage through high potential areas. Another challenge with PRMs is that it is hard for the user to visualize and understand the progress made by the planner. In molecular docking, this problem becomes even more significant since the familiar three-dimensional workspace is replaced with potential energy landscape. To address these problems, we add a human operator who explores the energy landscape with a haptic device. The user selects some important configurations of the ligand with the haptic device which are passed to the planner for further processing. Since feeling the potential energy forces is not sufficient to understand the planner's progress, roadmap visualization methods similar to [2] are also employed.

In his paper we will use the terms configuration and conformation interchangeably.

2 Previous Work

Many automated docking algorithms have been proposed, such as, AutoDock [10], Dock [5], FlexX [15], and FTDock [4]. However useful they are, their success ratio is often relatively low (2-20%) [20]. Also, in most cases these algorithms make the simplifying assumption that the ligand is rigid.

Recently, a new approach which makes use of the PRM paradigm for protein-ligand binding was proposed in [16]. Their results were very promising, and are in large part the motivation for this work. They uniformly sample configurations, and resample more densely around a small subset of low potential configurations from the original sample. In essence, they randomly sampled configurations, identified those with low energy, and further explored those regions. Since the goal of the obstacle-based PRM used in our work is to generate, directly, configurations near the protein's surface, we believe that it will provide a better sample of this area than the uniform sampling approach of [16]. As will be seen in Section 7, some evidence in favor of our approach is our success in generating configurations in the binding site for a test case in which [16] failed.

Some researchers, such as SMD [6], STALK [7], and SCULPT [19], have added a human operator into the decision process. A revolutionary approach taken in [12] was to include a haptic interface and let the user feel the force applied on the ligand during the docking process. They also show in [12] that force feedback performed better than visual display in the docking process. In [13], an interactive molecular docking simulator was used for interactive molecular design. It is shown in [21] that SCULPT can be improved by adding a haptic device.

Unfortunately, there has not been much work on improving automated planners with human input. In [2], we used a haptic device together with a PRM planner on motion planning problems typically of maintainability studies in mechanical CAD designs. Our results showed that human insight improved the planner's ability to solve problems involving narrow passages in C-space.

3 Potential Energy

A key requirement in simulating the ligand binding process is to have a good model of the interaction between a ligand and its receptor, a protein. Proteins usually consist of thousands of atoms, whereas the ligand typically consists of tens. While both protein and ligand are in fact flexible, we will model the protein as a rigid body and the ligand as a flexible body as was done in [16]. In particular, the ligand is modeled as an articulated body with a free base, where the torsional movement of atomic bonds is modeled by revolute (1 dof) joints (bond angles and lengths are fixed). Thus, the base atom requires 6 dof and each additional torsional movement requires 1 dof.

During binding, the ligand can be viewed as moving in the potential field created by the protein's atoms, searching for a stable low potential configuration. Since a complete, accurate model of the potential is computationally infeasible, we approximate the potential by the van der Waals terms only, excluding other components such as Coulomb potential and solvent effects. While this is an extreme simplification, the results obtained are quite impressive. In particular, the potential we use is

$$U = \sum_{atom\ pair\ i,j} (A/r_{ij}^{12} - B/r_{ij}^{6}),$$

where r_{ij} is the distance between protein atom i and ligand atom j. Parameters are taken from [8].

Even with our simplified potential, computing the potential energy of a ligand configuration is expensive since it includes a contribution for each atom pair. To optimize this computation we use a grid-based energy calculation. In this approach, the protein is covered by a three-dimensional grid, and each grid point stores the potential applied by all protein atoms at that point [14]. During potential energy calculations, the value stored at the closest grid point to each ligand atom is used as an approximation of potential for that atom. The accuracy of this method directly depends on the size of the grid. Our experiments indicate that a 0.5 Å grid is sufficient for realistic force feedback for our haptic device.

4 Haptic Interaction for Molecular Docking

Our prototype system consists of a PHANToM haptic device (6 dof) output) [9], and a Xeon 550Mhz NT Workstation. The operator uses the PHANToM to manipulate a *rigid* ligand around a virtual protein.

A challenge with the molecular docking is to compute the potential energies of the configurations fast enough for haptic interaction. To achieve this, we used a grid-based

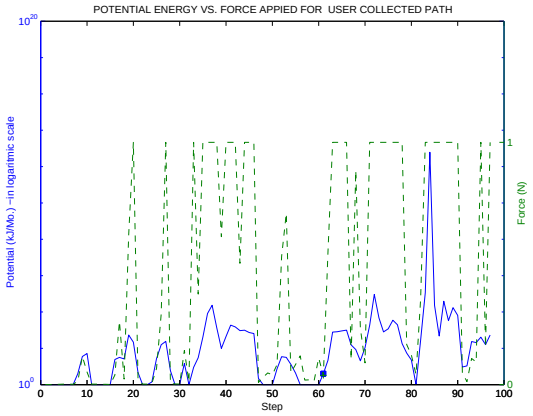

Figure 2: Correlation between potential energy (logaritmic left y-axis, in solid line) and user-felt force (right y-axis, in dashed line).

force calculation as described above, which was also used in [11]. The efficiency of this method can be observed in Figure 2, where both the real potential energy and the force the user felt (based on a grid) while collecting a path of ligand configurations are shown. In the figure, the potential is shown with the solid line and its (logarithmic) y-axis is represented on the left . Similarly the user felt force is shown with the dashed line and the its y-axis is represented on the right. To stay in the PHANToM's force feedback limits, any computed force greater than 1N was reduced to 1N without changing its direction, which resulted in a repulsive force in the correct direction. The binding configuration is represented by a small rectangle in both lines. The correlation between the figures shows the user feels low repulsive forces when the real potential energy is low. Note that at the binding configuration, the potential energy is very small and the user felt almost no force, hence with visual and haptic feedback the user can estimate the position of a binding site.

Our haptic system also lets the user investigate a path or roadmap generated by the automated planner. The user can visualize a path, select a configuration in the path, and investigate the potential field around that configuration. Similarly, the user can visualize a roadmap. Roadmap configurations are represented by points in the three-dimensional workspace and roadmap edges by lines connecting them [2]. The user can investigate a specific roadmap configuration by moving the haptic pointer close to a roadmap configuration causing the configuration to pop up visually.

5 Generating Binding Site Candidates

Good binding configurations are thought to be configurations at local minima in the potential landscape, and moreover, should be surrounded by high potentials so that the ligand will tend to stay in the site once it gets there [16].

Our goal is to automatically generate such configurations of the ligand. Our strategy is to first generate nodes near the potential energy surface, and then push the generated nodes to their local minima. Nodes can be generated automatically with our OBPRM planner, or manually by the user with the haptic device. Next, these nodes are connected into a roadmap that will be used to determine the accessibility of a site (this is very useful for eliminating 'decoy' nodes that are inaccessible, but otherwise appear to be good candidates). Finally, we will compute a weight for all low potential configurations contained in large connected components of the roadmap (see Section 6). Details of each step are discussed below.

Automatic Node Generation. OBPRM[1] is designed to generate contact and surface nodes. We model each protein atom as a unit cube (1 Å), and each ligand atom as a sphere with 1 Å radius. Considering only the base atoms of the articulated ligand model, we generate a collision-free (geometrically) node for the base. Then, we assign random values for the joint angles. Next, we evaluate the potential of the ligand's configuration and keep the node only if its potential is less than some fixed value $E_{\max}$. A fairly high value of 5,000 KJ/Mol was chosen for $E_{\max}$. However, after applying gradient descent to the nodes (described below), most remaining nodes have potentials below zero.

User Generated Nodes. We also experiment with operator collected nodes using a visual and haptic interface. Our previous work with CAD models has shown significant improvement when user input is given to the automatic planner [2]. A user would collect a node when he/she felt like it was local minimum in the potential and also saw it was in some sort of pocket on the protein. The ligand was restricted to a rigid body for the haptic interaction.

Approximate Gradient Descent. Since a binding configuration is commonly thought to be a local minima surrounded by higher potentials, it is natural to push our generated nodes to local minima. In particular, we perform an approximate gradient descent as follows. For each generated node, we sample uniformly twice the ligands dof nearby nodes, and select the lowest potential new node for iteration. This process is repeated until an iteration limit or a local minima is reached (when all sampled nodes have higher potential, usually in 2-4 iterations).

Node Connection. After gradient descent is performed, all local minima nodes are added to the roadmap (with the previously generated nodes). Then, we attempt to connect each node to its k nearest neighbors along the straight-line in C-space between them (we use $k = 10$). An edge between the node and its neighbor is added to the roadmap if it passes some validity check. The only difference here is that our validity check uses potential evaluation instead of collision detection.

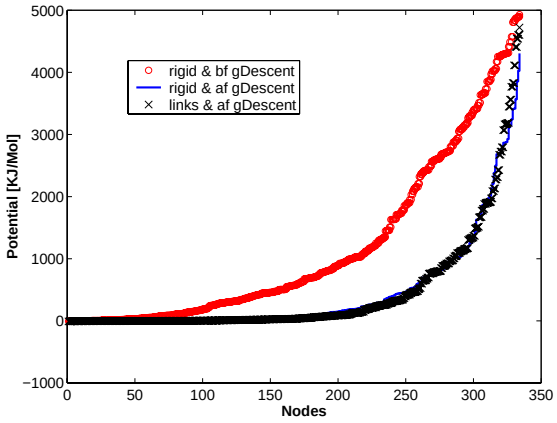

Figure 3: The potential distribution of all roadmap nodes before and after gradient descent for 1LDM protein-ligand complex.

6 Recognizing Binding Sites

In order to identify configurations in the binding site, as well as to determine promising configurations for further processing, we need to rank (score) configurations.

In the PRM based method presented in [16], they first identify the 10 roadmap nodes with the lowest potential, and then further evaluate them with a score function called average path weight. They successfully found the binding sites for two of the three examples they studied.

Our approach in this paper is similar to theirs, but has two significant differences.

First, after the gradient descent most of the roadmap nodes end up having very low energies, as shown in Figure 3, and the approximation used for the potential can easily change the relative order of the nodes. So potential energy alone cannot be used as a selection criteria. However, one factor that can be used to filter nodes is accessibility. We use the roadmap's connectivity to filter out inaccessible configurations, and in particular, only evaluate nodes contained in the largest connected component(s) of the roadmap.

Second, we use a different weight function. Since the binding site is a *local* feature, we observe there is no need to evaluate path weights globally as in [16]. Instead, we propose building a 'local roadmap' for each candidate node and evaluating path weights in it. In particular, we first uniformly sample n nodes within a fixed distance r from the candidate node, and then define the weight, or score, of the candidate as the average potential of all the sampled nodes. The potential is truncated to 5,000 KJ/Mol if it has a larger value.

7 Experiments

We tested our approach on three different protein-ligand complexes obtained from the Protein Data Bank at `http://www.rcsb.org/pdb/`, Two of these complexes, 1LDM (M_4-Lactate Dehydrogenase-Oxamate) and 1STP (Streptavidin-Biotin) , were studied in [16]. The third, 1A5Z (L-Lactate Dehydrogenase-Oxamate), was chosen based on the complexity of protein's potential field which created many isolated areas. Some details of these complexes are shown in Table 2.

In our experiments we investigate the following questions:

- Does the approximation of the ligand molecule as a rigid object affect our results?
- Can an OBPRM-based automated planner generate nodes in binding sites?
- Can a user provide helpful information to an automated planner?

To answer these questions we perform three different experiments. (i) In the first experiment, the ligand is treated as a rigid body. The planner generates configurations, and finds and ranks potential binding sites using these configurations. (ii) In the second experiment, the user collects configurations using the PHANToM by manipulating a rigid ligand. The automated planner then fine-tunes these configurations, now treating them as articulated bodies, and evaluates the possible binding sites. (iii) In the third experiment, we investigate the behavior of the automated planner when the ligand is flexible, i.e., articulated.

The running time and roadmap statistics are shown in Table 1. The number of configurations the user collected varies and can be found in Table 2.

7.1 Results

A summary of our results is shown in Table 3. For each experiment for a protein-ligand complex, we list the top five scoring configurations identified. The first column represents the rank of the configuration. The 0 ranked configuration is the binding configuration for the respective complex. The RMSD column shows the Root Mean Square Deviation with respect to the known binding site. The score of that configuration is computed as described in Section 6. The last column represents the potential energy of the complex at that configuration. Based on the rigid body experiments of 1A5Z and 1STP (Table 3(a) and 3(c)) we did not consider the rigid body representation for 1LDM.

As can be seen, in almost all cases, at least one the top five scoring configurations is very close (2-3 Å) to the true binding configuration. This suggests that our automated planner is effective at finding configurations close to the binding site. In addition, our success in identifying the binding site for 1STP shows an advantage of our OBPRM based approach over the results presented in [16], where the closest configuration identified was 13 Å from the true binding site.

Based on the automated results for the rigid body and articulated representations, for 1A5Z and 1STP in Table 3,

Roadmap Statistics						
Model	Gen. time	Conn. time	Total time	# node(init #, # local min added)	# edge	bigCC
1A5Z Rigid	55.	70	229.8	1429 (1000, 429)	3784	1182
1A5Z User	75	101	199	1684 (1666, 18)	3045	1304
1A5Z Artic.	75	203	657.9	2417 (1666, 751)	6253	2071
1LDM User	853	141	279	2011 (2000, 11)	4141	1643
1LDM Artic.	85	299	1089	2912 (2000, 912)	8472	2579
1STP Rigid	104	369	822	1771(1000, 771)	5589	1738
1STP User	162	163	343	1674 (1666, 8)	3115	1504
1STP Artic.	162	447	2095	3020 (1666, 1354)	7971	2909

Table 1: Roadmap statistics. Total time includes node generation time, connection time, and node score calculation time. All times are in seconds.

	Protein	***Ligand***		**# Cfgs user**
Pair ID	**# Atom**	**#Atom**	**dof**	**collected**
1LDM	2544	6	7	37
1A5Z	2416	6	7	129
1STP	1001	16	11	92

Table 2: Statistics for different protein-ligand complexes. The table also shows number of configurations that are collected by the user during the node generation phase.

it is clear that better results were obtained using the articulated representation, which did incur some additional computational costs as seen in Table 1. The articulated representation was slightly better for 1A5Z, where the difference in the degrees of freedom between the rigid and articulated versions is small. When these differences increase the performance gap increases, as with 1STP. In general, an articulated ligand can move easier then a rigid ligand in narrower areas by taking on the shape of the surrounding potential surface. Our comparison of the rigid and articulated representation is potentially relevant for other docking methods, since representing the ligand as a rigid body is a commonly used simplification.

The results for the fully automated articulated model and those starting from user generated nodes are comparable. In 1A5Z and 1LDM the user was able to find the closest configurations to the binding site. However, in 1STP the automated planner reached a closer configuration. This may be related to the fact that the user has to work with a rigid version of the ligand, while the automated planer takes advantage of the high-dimensional space. The large difference in the performance of the rigid body and the articulated body representations for the 1STP complex supports this idea.

An advantage of user input is that, as shown in Table 1, it was noticeably faster than the fully automated method since the user provided a much smaller sample of nodes. Fewer nodes were provided since the user only collected nodes in/near pockets. A disadvantage is that the user input suffers the users' bias. It is very tedious for the user to explore every corner of the protein to find a binding site, therefore he/she may miss the true binding site. In addition, a user may be misled by some 'decoy' sites based on only visual and/or force feedback as in 1A5Z where several nodes lie in a site far away.

Finally we remark that our energy filtering keeps many nodes that are in the flat region of the potential distribution plot. Therefore, we avoid throwing away nodes near binding sites even though their potential may be slightly higher, which could very easily be caused by inaccuracy of our potential function.

8 Conclusions and Future Work

Our results show that our approach to molecular docking is promising. In the examples we studied, we were able to find and recognize configurations in the true binding site. However, further processing, perhaps with a more accurate potential function, will be needed to find an exact binding configuration. For example, candidates identified by our method might be used as input for other docking programs that perform detailed simulations, such as molecular dynamics methods. We also saw that better results were obtained with an articulated representation for the ligand, as opposed to the commonly used rigid body simplification. User input was seen to improve efficiency, and moreover, haptic feedback was observed to help the user better understand the problem.

Our future work includes finding a better representation for potential energy formulations, concentrating on the binding sites and reaching the binding configurations, and improving the haptic interface by letting the user collect configurations for a flexible ligand.

References

[1] N. M. Amato, O. B. Bayazit, L. K. Dale, C. V. Jones, and D. Vallejo. OBPRM: An obstacle-based PRM for 3D workspaces. In *Proc. Int. Workshop on Algorithmic Foundations of Robotics (WAFR)*, pages 155–168, 1998.

[2] O. B. Bayazit, G. Song, and N. M. Amato. Enhancing randomized motion planners: Exploring with haptic hints. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 529–536, 2000.

[3] D. Hsu, L. Kavraki, J-C. Latombe, R. Motwani, and S. Sorkin. On finding narrow passages with probabilistic roadmap planners. In *Proc. Int. Workshop on Algorithmic Foundations of Robotics (WAFR)*, 1998.

[4] E. Katchalski-Katzir, I. Shariv, M. Eisenstein, A.A. Friesem, and C. Aflalo I.A. Vakser. Molecular surface recognition: Determination of geometric fit between proteins and their ligands by correlation

1A5Z									
	RIGID			USER INPUT (RIGID)			ARTICULATED (7dof)		
Rank	RMSD(Å)	Score	Pot.	RMSD(Å)	Score	Potential	RMSD(Å)	Score	Pot.
0	**0.00**	**3743**	**35.51**	**0.00**	**3743**	**35.51**	**0.00**	**3743**	**35.51**
1	5.74	4694	0.09	2.08	4138	-8.783	5.82	4658	19.92
2	5.20	4462	24.81	1.94	3670	-5.19	7.31	4488	41.43
3	4.10	4422	0.81	15.90	3316	-5.209	2.31	4459	13.67
4	9.27	4376	30.36	16.89	2682	22.207	5.63	4393	-9.59
5	4.35	4290	-7.79	16.11	2413	-5.95	2.42	4342	19.88

(a)

1LDM						
	USER INPUT (RIGID)			ARTICULATED (7dof)		
Rank	RMSD(Å)	Score	Pot.	RMSD(Å)	Score	Potential
0	**0.00**	**4570**	**5.00**	**0.00**	**4570**	**5.00**
1	1.52	4652	-6.06	5.03	4634	30.96
2	2.93	4564	-6.73	2.03	4593	23.35
3	3.79	4541	-7.04	3.07	4577	48.12
4	2.35	4499	-6.88	4.14	4475	20.95
5	2.11	4463	-8.70	1.66	4418	43.34

(b)

1STP									
	RIGID			USER INPUT (RIGID)			ARTICULATED (11dof)		
Rank	RMSD(Å)	Score	Pot.	RMSD(Å)	Score	Pot.	RMSD(Å	Score	Pot.
0	**0.00**	**4552**	**-5.21**	**0.00**	**4552**	**-5.21**	**0.00**	**4552**	**-5.21**
1	4.21	4127	48.79	5.50	3431	49.22	2.56	4609	0.37
2	11.40	3934	-12.73	3.77	3426	11.10	5.05	4284	36.44
3	10.63	3921	7.94	7.32	3425	-6.47	3.44	4151	-5.79
4	12.42	3894	-10.51	4.99	3287	48.80	4.11	4014	18.31
5	11.29	3876	38.19	6.90	3211	44.05	7.41	3952	22.68

(c)

Table 3: Performance of the scoring function on each protein. The tables show the results of three different kind of experiments; OBPRM with the ligand as a rigid body, OBPRM with the ligand as an articulated body and user collected configurations. For 1LDM, we only experimented on OBPRM with the articulated ligand and user collected configurations.

techniques. In *Natl. Acad. Sci. USA*, volume 89, pages 2195–2199, 1992.

[5] I.D. Kuntz. Structure-based strategies for drug design and discovery. *Science*, 257:1078–1082, 1992.

[6] Jonathan Leech, Jan F. Prins, and Jan Hermans. Smd: Virtual steering of molecular dynamics for protein design. *IEEE Computational Science and Engineering*, pages 38–45, 1996.

[7] D. Levine, M. Facello, P. Hallstorm, G. Reeder, B. Walenz, and F. Stevens. Stalk: An interactive system for virtual molecular docking. *IEEE Computational Science & Engineering*, April-June:55–66, 1997.

[8] M. Levitt. Protein folding by restrained energy minimization and molecular dynamics. *J. Mol. Bio.*, 170:723–764, 1983.

[9] T. H. Massie and J. K. Salisbury. The PHANToM haptic interface: A device for probing virtual objects. In *Int. Mechanical Engineering Exposition and Congress, DSC 55-1*, pages 295–302, Chicago, 1994. C. J. Radcliffe, ed., ASME.

[10] 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 empirical binding free energy function. *J. Computational Chemistry*, 19:1639–1662, 1998.

[11] M. Ouh-young. *Force Display in Molecular Docking, TR90-004*. PhD thesis, Univ. of North Carolina at Chapel Hill, Chapel Hill, NC, USA, 1990.

[12] Ming Ouh-young, David V. Beard, and Frederick P. Brooks, Jr. Force display performs better than visual display in a simple 6-d docking task. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 1462–1466, 1989.

[13] Ruth Pachter and James A. Lupo. Virtual reality for materials design. In *The 14th Annual Dayton Section Symposium AESS/IEEE*, pages 51–56, 1997.

[14] N. Pattabiraman, M. Levitt, T. E. Ferrin, and R. Langridge. Computer graphics in real time docking with energy calculation and minimization. *J. of Computational Chemistry*, 6:432–436, 1985.

[15] M. Rarey, B. Kramer, T. Lengauer, and G. Klebe. A fast flexible docking method using an incremental construction algorithm. *J. Molecular Biology*, 261:470–489, 1996.

[16] A.P. Singh, J.C. Latombe, and D.L. Brutlag. A motion planning approach to flexible ligand binding. In *7th Int. Conf. on Intelligent Systems for Molecular Biology (ISMB)*, pages 252–261, 1999.

[17] G. Song and N. M. Amato. A motion planning approach to folding: From paper craft to protein folding. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2001. To appear.

[18] G. Song and N. M. Amato. Using motion planning to study protein folding pathways. In *Proc. Int. Conf. Comput. Molecular Biology (RECOMB)*, pages 278–287, 2001.

[19] M.C. Surles, J.S. Richardson, D.C. Richardson, and F.P. Brooks, Jr. Sculpting proteins interactively: Continual energy minimization embedded in a graphical modeling system. *Protein Science*, 3:198–210, 1994.

[20] Maxim Totrov and Ruben Abagyan. Derivation of sensitive discrimination potential for virtual ligand screening. In *Third Annual International Conference on Computational Molecular Biology*, pages 312–320, 1999.

[21] Len Wanger. Haptically enhanced molecular modeling: A case study. In *Proceedings of the Third PHANTOM Users Group Workshop*, 1998.