

A New Haptic Interaction and Visualization Approach for Rigid Molecular Docking in Virtual Environments

Abstract

Many biological activities take place through the physicochemical interaction of two molecules. This interaction occurs when one of the molecules finds a suitable location on the surface of the other for binding. This process is known as molecular docking, and it has applications to drug design. If we can determine which drug molecule binds to a particular protein, and how the protein interacts with the bonded molecule, we can possibly enhance or inhibit its activities. This information, in turn, can be used to develop new drugs that are more effective against diseases. In this paper, we propose a new approach based on a human-computer interaction paradigm for the solution of the rigid body molecular docking problem. In our approach, a rigid ligand molecule (i.e., drug) manipulated by the user is inserted into the cavities of a rigid protein molecule to search for the binding cavity, while the molecular interaction forces are conveyed to the user via a haptic device for guidance. We developed a new visualization concept, *Active Haptic Workspace (AHW)*, for the efficient exploration of the large protein surface in high resolution using a haptic device having a small workspace. After the discovery of the true binding site and the rough alignment of the ligand molecule inside the cavity by the user, its final configuration is calculated off-line through time stepping molecular dynamics (MD) simulations. At each time step, the optimum rigid body transformations of the ligand molecule are calculated using a new approach, which minimizes the distance error between the previous rigid body coordinates of its atoms and their new coordinates calculated by the MD simulations. The simulations are continued until the ligand molecule arrives at the lowest energy configuration. Our experimental studies conducted with six human subjects testing six different molecular complexes demonstrate that given a ligand molecule and five potential binding sites on a protein surface, the subjects can successfully identify the true binding site using visual and haptic cues. Moreover, they can roughly align the ligand molecule inside the binding cavity such that the final configuration of the ligand molecule can be determined via the proposed MD simulations.

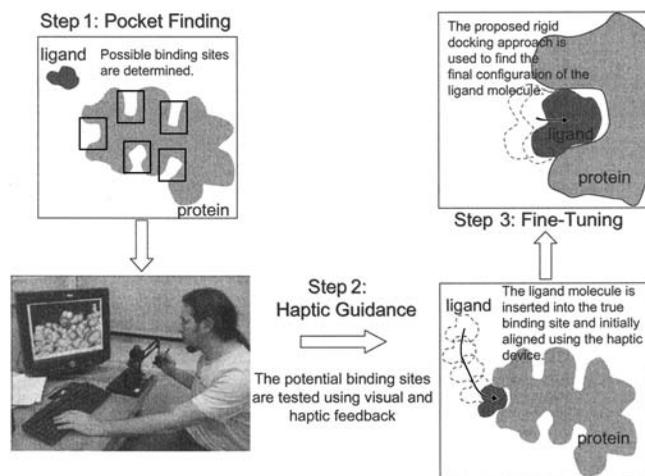


Figure 1. The steps of the proposed rigid docking approach. First, the potential binding sites are determined based on the geometry of the protein molecule using the software package Pocket (Edelsbrunner & Koehl, 2005). Then, the potential sites are tested by the user in virtual environments under the guidance of visual and haptic cues. Once the true binding site is determined and the ligand molecule is roughly aligned inside the binding cavity, the proposed rigid docking algorithm is executed to find its final configuration.

I Introduction

Molecular docking is defined as the process by which two molecules bind each other in an orientation and position determined by their geometric shape and local physicochemical properties. The geometric shapes of molecules define how well the binding surfaces complement each other, while the physicochemical properties define how well the binding strength of the interaction energies holds molecules together. Hence the goal in molecular docking studies is to determine whether two molecules interact, and if so, the binding position and orientation of molecules, such that the surface area of interaction is maximized, while the interaction energy is minimized.

The research on molecular docking has mainly focused on (a) the ligand-protein docking, and (b) the protein-protein docking problems. Ligand-protein docking involves a small molecule (the ligand) and a large molecule (the protein—also called the receptor)

while protein-protein docking involves two proteins that are approximately the same size. The computational approaches developed to solve these problems typically involve multiple stages, but different strategies in each stage (Kuntz, 1992; Halperin, Ma, Wolfson, & Nussinov, 2002; Cole, Murray, Nissink, Taylor, & Taylor, 2005). Most algorithms execute a low resolution geometric search and then a high resolution refinement stage (Kuntz, 1992). For example, if we consider the ligand-protein binding problem, one begins by searching the knobs and cavities on the surface of the protein molecule to identify the potential binding sites mostly based on shape complementarity. Then, MD simulations are typically executed to bring the ligand and receptor molecules together in the energetically most favorable final conformation. In each time step of the MD simulations, the new positions of the atoms under the influence of molecular interaction forces are calculated using Newton's Third Law until the molecules arrive at the lowest energy configuration. The success of this and similar approaches depends on several factors, including the initial geometric search, models of MD and its parameters, and scoring functions defined to find the optimum binding configuration.

More recently, virtual reality techniques have been applied to molecular simulation (Akkiraju, Edelsbrunner, Fu, & Qian, 1996; Levine et al., 1997; Anderson & Weng, 1999; Sharma et al., 2003). The earlier studies in this area involve the use of the CAVE (a room in which the user is surrounded by stereoscopic images of virtual objects rendered on the walls and the floor; Akkiraju et al., 1996; Levine et al., 1997) or a large projection screen (Anderson & Weng, 1999; Sharma et al., 2003) for an immersive and interactive visualization experience. The interaction with the molecular models has been mainly provided through a wand (a 6 DOF mouse) with no force feedback to the user. Nowadays, it is possible to explore and manipulate molecular models in virtual environments with the use of haptic devices enabling force interactions. The first work in this area dates back to the late 1980s (Ouh-Young, Pique, Hughes, Srinivasan, & Brooks, 1988; Ouh-Young, Beard, & Brooks, 1989; Brooks, Ouh-Young, Batter, & Kilpatrick, 1990). Brooks and his colleagues at UNC

simulated molecular docking in virtual environments with haptic feedback to the user. A group of experienced biochemists were asked to dock four ligand molecules to the known binding cavity of a protein molecule to discover which ligand molecule was the best choice. The experiments were repeated with and without force feedback, and the results showed that the force feedback improved the task performance of the subjects. Despite the early discovery of the potential benefits of haptic feedback in molecular simulation, more progress has not been made until recently due to the lack of rendering software and high-fidelity and low-cost commercial haptic devices. Yet there are still only a limited number of research studies in this area, which are briefly reviewed below.

In general, the use of haptics in molecular biology applications can be grouped into three areas. Haptic devices are used for (a) teaching structural molecular biology, (b) feeling interaction forces during molecular visualization, registration, docking, and interactive particle steering in virtual environments, and (c) interactive manipulation of actual molecular structures in the real world. Sankaranarayanan, Weghorst, Sanner, Gillet, and Olson (2003) have developed an augmented reality system supported by haptic feedback for teaching structural molecular biology to students. In this system, graphical virtual models are superimposed on the physical models of molecules such that the student interacts with these virtual enhancements through a haptic device while manipulating the physical model. Peterlík and Krenek (2005) investigate the conformational behavior of a molecule using a haptic device. In their approach, the haptic device is virtually coupled to a sphere having the size of a water molecule. The interactions between the sphere and the molecule change the conformation of the molecule (i.e., spatial arrangement of its atoms). The force required to make these changes is delivered back to user via the haptic device to gain an insight on the conformational behavior. Birmanns and Wriggers (2003) utilize haptic feedback for registration of low-resolution electron microscopy data with high-resolution molecular structures. They use the gradient of a cross correlation function to calculate the interaction forces and torques, which are conveyed to the user via a haptic

device for guidance in finding the optimum fit. To achieve real time and stable haptic rendering rates (on the order of 1 kHz), they utilize vector quantization techniques. Nagata, Mizushima, and Tanaka (2002) developed a VR-based simulation system that enables a user to explore the surface of a protein molecule using a globular probe that is given an electrostatic charge and manipulated by a haptic interface, to search for sites where the probe is strongly attracted to the force field. However, this system prevents the simulation of drug molecules. Lee and Lyons (2004) present a new method for smooth rendering of interaction forces between ligand and protein molecules during the simulation of molecular docking in virtual environments with haptic feedback to user. When the Lennard-Jones (hereafter LJ) potential field is used and the ligand atoms are in close proximity to the receptor atoms, the magnitude of the interaction forces increases drastically, exceeding the limits of the haptic device and leading to force instabilities. In rendering these forces, Lee and Lyons (2004) keep the gradient of the potential field unaltered when the distance between the atoms of ligand and protein molecules is greater than the sum of their van der Waals radii, but render a simple hard surface wall when they are smaller. This approach eliminates force instabilities, even in the presence of strong force gradients. Bayazit, Song, and Amato (2001) present a new framework for the solution of the ligand-protein binding problem based on the path planning techniques used in robotics. They investigate the effect of supplementary user input collected via a haptic device in identifying the low energy configurations. In their approach, the user manipulates a rigid ligand molecule around a protein molecule and samples the configuration space for low-energy configurations using the haptic device. Then the selected configurations are connected to each other via a probabilistic road map planner to find the accessibility of the binding site. In addition to helping a user to sample the configuration space, the haptic device also lets the user trace the path generated by the motion planner. Lai-Yuen and Lee (2006) also use a haptic interface to sample the search space for finding the binding site of a flexible ligand molecule. The torsional angles of the ligand molecule are allowed to change, providing flexibil-

ity to the ligand molecule. An adaptive local search method is executed to find a set of new torsional angles that result in a lower energy conformation of the ligand molecule. The molecular interaction forces are displayed to the user based on the method suggested by Lee and Lyons (2004). Stone, Gullingsrud, Schulten, and Grayson (2001) integrated a haptic device into NAMD, a public domain molecular dynamic simulation package for molecular steering. In their application, the haptic device enables the steering of molecules during MD simulations with real-time force feedback to user. The data communication between the graphics engine and NAMD is achieved through an efficient socket connection. The data transfer between the visual and haptic displays is provided through the VRPN protocol developed by Taylor et al. (2001). Stone et al. (2001) demonstrated that a molecular biologist could interactively steer a sodium ion through a gramicidin A channel while feeling the interaction forces through the haptic device. Another interesting application of the haptic devices to molecular biology is the nanoManipulator system developed at UNC. The nanoManipulator system integrates a haptic device with an atomic force microscope (AFM) for manipulation of molecular structures such as DNA (Guthold et al., 1999). Currently, direct manipulation of nanoscale objects using a standard AFM system is not possible, since the scanning and manipulation processes are done in sequence using the same AFM probe. In the nanoManipulator, the haptic interface enables the user to manipulate a DNA structure remotely and feel the interaction forces during the manipulations, while the changes, for example, in DNA shape and position, can be observed in a simulated world. The new shape and position of the DNA molecule in the virtual world is updated using the AFM scans once in every other scan only, reducing the need for frequent scanning.

In this paper, we propose new computational approaches for the solution of the ligand-protein binding problem based on the paradigms of human-computer interaction. In our approach, the user manipulates a rigid ligand molecule in virtual environments using a haptic device and explores the surface of a fixed protein molecule to find the true binding site. Then, the ligand

molecule is inserted into the binding cavity and roughly aligned with the help of force feedback. Finally, MD simulations are performed off-line to calculate the final configuration of the ligand molecule in the binding cavity. This paper has three main contributions:

1. In our approach, the final configuration of the ligand molecule inside the binding cavity is calculated off-line through time stepping MD simulations utilizing a new rigid docking algorithm. The initial alignment of the ligand molecule for the MD simulations is supplied by the user with the help of a haptic device. This rough alignment reduces the risk of being trapped in a local minimum during the MD simulations. The conventional energy minimization approaches that perform the same task without a good initial configuration of the ligand molecule typically suffer from being computationally too intensive, or from getting trapped in local minima more often (Apaydin, Guestrin, Varma, Brutlag, & Latombe, 2002). Moreover, the proposed rigid docking algorithm is computationally more efficient than the algorithms employing conventional rigid body equations for molecular docking (Timothy & Forester, 1998; Rapport, 2002).

2. For the visualization of a large protein surface to search for potential binding sites, a new haptic visualization technique, called the *Active Haptic Workspace* (AHW) was developed (Subasi & Basdogan, 2006). A protein surface contains many knobs and cavities, requiring a large scale factor to be used for effective visualization. However, the workspace of a haptic device is typically limited by the physical dimensions of its links and scaling the coordinates of a protein molecule such that it fits into the haptic workspace results in an insufficient spatial resolution for the haptic exploration. The existing methods for the visualization of a large scale object using a small haptic device mainly rely on position or rate control of the visual cursor. In our approach, when the haptic cursor (i.e., the center of the ligand molecule) is inside the AHW, the user interacts with the protein surface directly as in position control. As the haptic cursor crosses the boundaries of the AHW, the part of the protein surface being explored by the user with the haptic device is translated and/or rotated in real time at a certain rate with the help of effi-

cient coordinate transformations (instead of adjusting the velocity of the visual cursor as in rate control, which causes a mismatch between the movements of haptic and visual cursors). This gives the impression to the user that he or she is moving the haptic workspace actively over the protein surface.

3. We conducted human experiments to investigate the role of haptics in molecular docking. To our knowledge, this is the second experimental study in this area. The UNC group (Brooks et al., 1990) investigated whether the subjects could find the correct ligand molecule (among four candidates) that binds to the known binding cavity of a protein molecule using haptic feedback. We tested whether the subjects could find the correct binding cavity (among five candidates) for the given ligand and protein molecules that are known to bind to each other. We also investigated whether the initial rough alignment achieved by the subjects could be used as an input for the off-line MD simulations to calculate the final configuration of the ligand molecule in the binding cavity.

In the rest of the paper, we first formally state the ligand-protein docking problem (Section 2). Our approach to rigid docking is discussed in Section 3. The haptic rendering method and the new haptic visualization techniques are given in Sections 4 and 5, respectively. In Section 6, we present the design details and results of our molecular docking experiments. Finally, the last section concludes the study with a discussion of the performed work and the possible improvement that can be made on the current system.

2 Rigid Protein-Ligand Docking

The protein-ligand docking problem can be considered as an optimization problem where the goal is to search the configuration space for the lowest energy configuration of the ligand and protein molecules. Even if we neglect the physicochemical interactions between two molecular structures and only consider their geometric features for binding complementarity, the large search space already poses a significant problem. Both the ligand and protein molecules are characterized by a

collection of atoms and rotatable bonds (the atoms are connected to each other through covalent bonds and the rotation of two outer bonds about a central one is defined by the dihedral angle). The large degrees of freedom of the molecule arise from the rotatable bonds, since the bond lengths (distance between atom centers) and angles (angle between two consecutive bonds) do not change significantly. While a small ligand molecule typically has a few rotatable bonds, the large protein molecule may have thousands of rotatable bonds. Given that both ligand and protein molecules are not completely rigid, especially when binding, the problem of finding their lowest energy configuration becomes highly challenging. In fact, this flexibility of the molecules results in hundreds to thousands of degrees of freedom and a huge number of possible binding conformations (Teodoro, Phillips, & Kavraki, 2001). For this reason, most of the approaches to molecular docking have been limited to “rigid” docking to reduce the number of computations. In rigid docking, the binding molecules are considered as rigid objects that cannot change their spatial shape before or during the docking process. If the ligand and protein molecules are assumed to be rigid, then the geometric approaches to molecular docking can be implemented more easily. Moreover, the calculation of intrabody forces is not necessary, and longer time steps are possible in molecular dynamics simulations due to the elimination of higher frequency motions (Gillilan & Lilien, 2004). Furthermore, the majority of docking algorithms assume that the receptor molecule is fixed. Since the binding ligand molecule makes only three translational and three rotational movements in free space, these assumptions limit the search space to six-dimensional configuration space. In light of these assumptions, we formulate the rigid docking problem as

Given rigid ligand and protein molecules (L and P , where P is fixed) with their atomic coordinates in \mathbb{R}^3 , find an optimum rigid transformation $T: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ such that the potential energy of $T.L$ and P is minimized.

The total potential energy of two interacting molecules can be calculated using the contribution of bonded and non-bonded terms. If the atoms of a mole-

cule are modeled as spheres, and the bonds between them as springs, then the mechanics of spring deformation can be used to describe the effect of bonded terms. When the atoms change their position, the bonds between them stretch, bend, and twist, resulting in a flexible molecule. Since we treat the ligand and protein molecules as rigid bodies and neglect the internal interactions, the bonded terms do not contribute to the total energy. Non-bonded terms include van der Waals and electrostatic interaction energies. The interaction force between two atoms can be calculated using the gradient of these energies (Lee & Lyons, 2004). The total force acting on each atom of the ligand molecule is calculated by adding all the interaction forces between that atom and all the atoms of the protein molecule. To simulate the rigid-body behavior of the ligand molecule, the forces acting on all of its atoms must be calculated. The summation of these forces acts on the center of mass of the ligand molecule and also generates a torque about it. The formulations for rigid-body molecular dynamics are well documented in the literature (Rapport, 2002).

3 Our Approach

We propose a new approach that is simple and computationally more efficient than the standard implementation of rigid body equations for simulating the rigid-body dynamics of the ligand molecule. We initially relax the rigid-body assumption and calculate the new positions of the ligand atoms using Newton's Third Law:

$$F_i = m_i \frac{d^2 r_i}{dt^2} \quad (1)$$

where m_i is the mass of the atom, r_i is its position, and F_i is the total force acting on it. This equation can be numerically integrated to calculate the new positions of the ligand atoms in time stepping iterations. This process is known as molecular dynamics (MD) simulation (see the details in Van Gunsteren & Berendsen, 1990; Rapport, 2002). MD simulations have been commonly used in studying ligand-protein and protein-protein

docking problems. There are many commercial and public domain software packages available for running MD simulations (see the comparison in Cole et al., 2005). Finding the global minimum energy configuration of a ligand molecule using MD simulations is difficult, since the ligand molecule may easily become trapped in a local minimum while traversing the surface of a protein molecule. Hence, the quality of the results obtained from MD simulations in ligand-protein docking problems depends on the initial configuration of the ligand molecule and many other factors. If the possible initial configurations are sampled by the human operator under the guidance of visual and haptic cues, we believe that better results could be obtained.

In our approach, visual and haptic guidance is used for finding the true binding site and the rough initial alignment of the ligand molecule inside the binding cavity. The final configuration of the ligand molecule is calculated off-line via time stepping MD simulations. At each time step, the optimum rigid body transformation of the ligand molecule is calculated by minimizing the distance error between the previous rigid body coordinates of its atoms and the new coordinates calculated by Newton's Third Law, Equation 1. We formulate the computation of optimum transformation as a general least square minimization problem and state it as

Given two sets of corresponding points, find the optimum rigid transformation, $T_{4 \times 4}$, that minimizes the distance between them.

If $r_i^{t_0}$ represents the current rigid body coordinates of a ligand atom and $q_i^{t_0 + \Delta t}$ represents its new location due to the MD simulations (i.e., due to the effect of the molecular force F_i acting on it), then the optimum rigid transformation of the ligand molecule is calculated by minimizing the total distance error between the current and simulated coordinates of its atoms as

$$E = \sum_i^L \| q_i^{t_0 + \Delta t} - R r_i^{t_0} - p \| \quad (2)$$

where L is the number of atoms of the ligand molecule, E is the total distance error, and $p_{3 \times 1}$ and $R_{3 \times 3}$ are the optimum translation vector and the rotation matrix

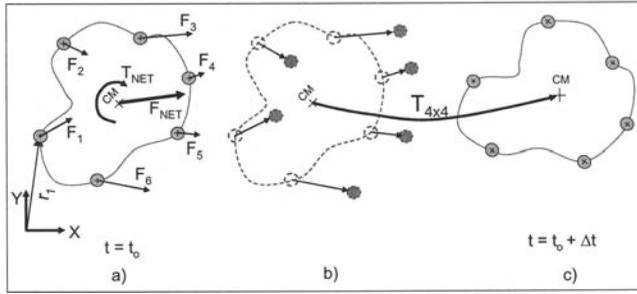


Figure 2. The proposed rigid docking approach. The ligand molecule shown in the figure is made of six atoms (a). The forces acting on the individual atoms generate a net force and torque about the center of mass. If the rigid-body equations are used directly, the new configuration of the ligand molecule can be calculated. Instead, we calculate the current coordinates of the atoms using the molecular dynamics simulations (b) and then determine the best rigid-body transformation between the current and previous coordinates that minimizes the distance error (c).

of the ligand molecule to be determined, respectively (see Figure 2).

This optimization problem can be solved using the singular value decomposition (SVD) method to obtain the optimum rotation matrix and the translation vector. Then

$$r_i^{t_0+\Delta t} = Rr_i^{t_0} + p$$

gives the new coordinates of the ligand atom, satisfying the rigid body assumption. For this purpose, we first calculate the deviations of the current and simulated coordinates from their mean values

$$\hat{r}_i^{t_0} = r_i^{t_0} - \left(\frac{\sum_i^L r_i^{t_0}}{L} \right), \quad \hat{q}_i^{t_0+\Delta t} = q_i^{t_0+\Delta t} - \left(\frac{\sum_i^L q_i^{t_0+\Delta t}}{L} \right) \quad (3)$$

and then construct a coherence matrix

$$A = \hat{R}\hat{Q}^T \quad (4)$$

where $\hat{R} = [\hat{r}_1^{t_0}, \dots, \hat{r}_L^{t_0}]$ and $\hat{Q} = [\hat{q}_1^{t_0+\Delta t}, \dots, \hat{q}_L^{t_0+\Delta t}]$. The singular value decomposition of this matrix, $[U, V, D] = SVD(A)$, enables us to calculate the optimum rotation matrix and the translation vector as

$$R_{3 \times 3} = VU^T \quad (5)$$

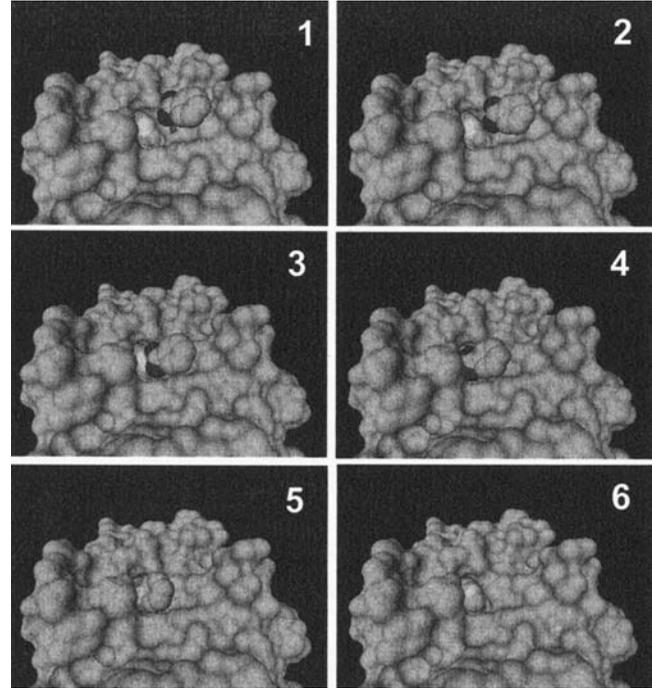


Figure 3. Snapshots showing the steps of the binding process of the Benzamidine molecule (ligand) with nine atoms to Beta-Trypsin molecule (protein) with 1,701 atoms. The ligand and protein molecules are extracted from the 3PTB complex in the protein data bank (PDB).

$$p_{3 \times 1} = \left(\frac{\sum_i^L q_i^{t_0+\Delta t}}{L} \right) - R \left(\frac{\sum_i^L r_i^{t_0}}{L} \right) \quad (6)$$

Finally, the optimum transformation matrix (Figure 2) can be constructed as

$$T_{4 \times 4} = \begin{bmatrix} R_{3 \times 3} & p_{3 \times 1} \\ 0 & 1 \end{bmatrix} \quad (7)$$

Figure 3 shows the implementation of this approach with a ligand-protein pair. As shown in the figure, the ligand molecule makes rigid body movements and enters into the binding cavity for a given initial configuration close to the binding site.

The proposed approach is computationally more efficient than the algorithms employing conventional rigid body equations for molecular docking problems. Table 1

Table 1. *The Proposed Approach Leads to Better Execution Times and RMS Error*

Ligand-protein pair	Number of ligand atoms	Number of protein atoms	Execution time (s)		RMSE (Å)	
			rigid body dynamics	Execution time (s) our approach	rigid body dynamics	RMSE (Å) our approach
Line segment cavity	6	19	31.3	25.5	13×10^{-5}	9.3×10^{-5}
Rectangular box cavity	13	19	64.7	49.4	1.5×10^{-5}	1.2×10^{-5}

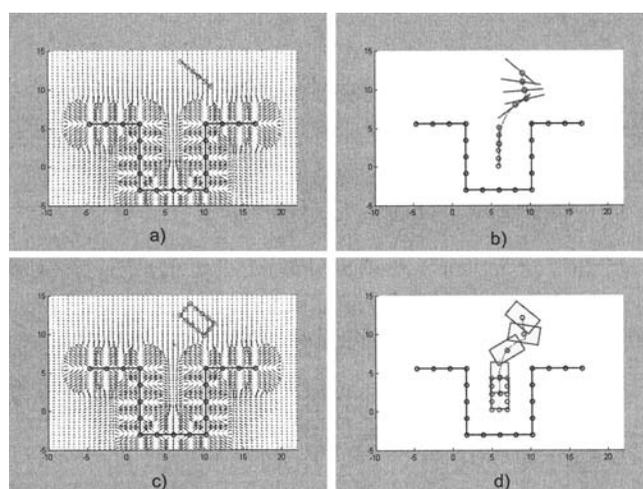


Figure 4. *Graphical comparison of the proposed rigid docking approach with the conventional rigid body solution in 2D. A line segment and a rectangular box, made of carbon atoms, are separately guided into an artificially created potential well (see windows a and c respectively). The paths calculated using our approach and the conventional rigid body simulations perfectly overlap with each other for both objects (see windows b and d).*

compares the results of the simulations performed in 2D using the proposed approach and the conventional rigid body equations. In both simulations, the simulated objects are guided into a potential well. Figure 4 shows that both objects successfully move towards the lowest energy configuration under the influence of molecular interaction forces when either the proposed approach or the conventional rigid body equations are used.

However, it is known that the methods utilizing the

gradient of potential energy for motion planning are sensitive to initial configurations and may suffer from the local minima problem. For example, the ligand molecule shown in Figure 3 traps in local minima for some initial configurations and cannot enter the binding cavity. In our approach, we start running MD simulations to search for the lowest energy configuration of the ligand molecule after it is inserted into the binding cavity and roughly aligned by the user. This initial alignment reduces the possibility of encountering a local minimum, though it does not totally eliminate it. To further reduce the risk, we perturb the configuration of the ligand molecule during the MD simulations using the well-known Metropolis method (Van Gunsteren & Berendsen, 1990). The basic idea behind this method is to add small random moves (translation and rotation) to the ligand molecule after every few MD iterations and then accept or reject the move based on a Boltzmann probability.

In Table 1, a line segment and a rectangular box made of 6 and 13 carbon atoms are separately guided into a cavity of 19 carbon atoms using the proposed approach and the conventional rigid body solution (see Figure 4). As shown in the table, the proposed approach leads to better execution times and RMS error (calculated with respect to the desired docking configuration) than the rigid body solution. Note that comparing RMS error can be a bit misleading since “numerical damping” is used in the conventional rigid body simulations to stabilize the solutions and its value affects the final configuration of the ligand molecule slightly at the level of accuracy given in the table.

For the implementation of the rigid docking approach proposed in this study, a set of parameters such as atom charges and masses and the constants of the LJ potential are necessary, which are taken from the CHARMM (Chemistry at HARvard Macromolecular Mechanics) force field model.

4 Haptic Rendering of Molecular Interactions

To provide haptic guidance to the user during molecular docking simulations, the net force arising from the molecular interactions between the atoms of the ligand and protein molecules must be calculated and reflected to the user through the haptic device. However, the computation of interaction forces involves a large number of pair-wise calculations. One must consider each atom of the ligand and loop over all the atoms of the protein to calculate the total force acting on the ligand (the same force with an opposite direction acts on the protein as well). This brute force approach requires $O(LP)$ calculations, which is not feasible for real-time implementation. To reduce the computation time, a grid-based approach is used as suggested in Lee and Lyons (2004).

In our simulations, the total force acting on the ligand molecule is displayed to the user through a haptic device. The total force acting on a ligand atom is the summation of long range attraction and repulsion forces due to LJ potential and the forces due to electrostatic charges. Due to LJ potential, the atoms initially attract each other when they are further apart, but when the distance is short, they repel each other very strongly. This sudden change, first in force direction and then in force magnitude, causes instabilities while displaying forces to a user via a haptic device. In particular, when two atoms penetrate into each other, the force magnitude reaches extremely high values in which the current haptic devices cannot render at all. If all the interaction forces are scaled down such that the penetration forces can be easily displayed, then the attraction forces become too small to be perceived by the user. For this reason, Lee and Lyons (2004) calculate van der Waals

interaction forces between two molecules until a contact occurs and then add a spring-based force component if the molecules further penetrate into each other (i.e., depth of penetration times a spring constant). Implementation of this approach requires the detection of collisions between the 3D geometric models of ligands and molecules. Haptic rendering of force interactions between two arbitrarily shaped 3D objects is not a trivial task and the research in this area is still active (see the review of haptic rendering concepts in Basdogan & Srinivasan, 2002). However, a 3D geometric model of a molecule is not arbitrary and has a certain topological structure. A common 3D surface representation of a molecule is based on spherical atoms with characteristic radii, also referred as the Connolly surface (Connolly, 1983). We take advantage of the Connolly representation to detect collisions between the atoms of the ligand and protein molecules efficiently (i.e., detecting collisions between two spheres is a straightforward task). If there is no contact between the ligand and protein molecules, the molecular interaction forces are calculated using the gradient of the potential energy function. If there is a contact, then the total force acting on the ligand molecule is the summation of the molecular forces at the point of contact and the spring forces due to the collision as suggested in Lee and Lyons (2004). To construct the Connolly surface of a molecule, a probe atom is rolled over the surface atoms of the molecule defined by their van der Waals radii while bridging the gaps via smooth surface patches.

5 Haptic Visualization of Molecular Surfaces

In order to search for the binding site, the ligand molecule is manipulated via a haptic device and the surface of the protein molecule is explored. The physical dimensions and the sensing resolution of the haptic arm used in our simulations are limited, which prevents the user from exploring the surface of a large protein molecule efficiently in a small haptic workspace. The actual dimensions of a protein molecule are on the order of a few hundred angstroms and the spatial resolution

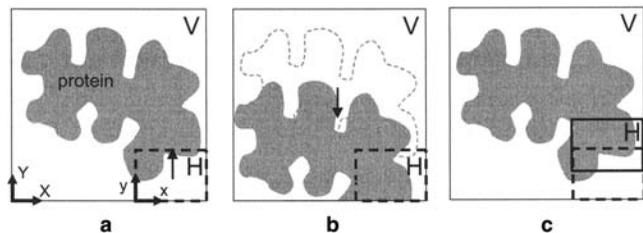


Figure 5. In this example, we show the concept of “moving” AHW for the haptic exploration of large molecular surfaces: As the user reaches the boundary of the AHW in the y direction while manipulating the ligand molecule (a), the protein molecule is moved in the $-y$ direction to allow the exploration of inaccessible parts (b). However, the user perception is that the haptic workspace moves up in the $+y$ direction (c).

achieved by scaling up the original atom coordinates such that it fits into the workspace of our haptic device is not sufficient for the detailed exploration of the potential binding sites. To resolve this problem, we propose the concept of *Active Haptic Workspace (AHW)*. The AHW is a transparent haptic subspace that actively travels in a large 3D visual workspace as the user manipulates the haptic device. This approach enables the user to explore all parts of a large protein surface interactively to search for potential binding sites in high spatial resolution while holding and manipulating the ligand molecule simultaneously (see Figure 5). As an alternative, Dominjon, Lecuyer, Burkhardt, Andrade-Barroso, and Richier (2005) propose the “bubble technique” and Conti and Khatib (2005) slowly shift the workspace of the haptic device toward the area of interaction of the haptic cursor. These approaches and ours target the same problem (i.e., visualization of a large size virtual object using a haptic device having a small workspace), but follow slightly different paths in implementation. In the bubble technique (Dominjon et al.), position control is used when the cursor is inside a spherical bubble, and rate control is used to adjust the cursor movements when it is outside the bubble. This approach was originally suggested by Hollis and Salcudean (1993) for telemanipulation of remote objects efficiently. We similarly use position control inside the AHW, but the objects being visualized are translated and rotated in real time

at a certain rate (instead of controlling the cursor velocity as in the rate control) using highly efficient coordinate transformations when the cursor is outside its boundaries. Our approach also allows the haptic visualization of the back side of the object, since the object can be rotated by pressing the switch on the haptic probe. In the approach proposed by Conti and Khatib, shifting the workspace of the device slowly when the cursor is in motion initiates the user to correct this “drift” unconsciously while executing a task with the device at the same time. However, Conti and Khatib (2005) report that the rate of drift can cause distortion between physical and visual representations at the edge of the physical workspace of the device.

In our approach, the part of the protein surface that is inside the AHW can be explored in high resolution with the haptic device when the haptic cursor (i.e., the geometric center of the ligand molecule) is inside the AHW. When the haptic cursor exceeds its predefined boundaries in one direction, the center of AHW is shifted while the protein surface is translated in the opposite direction at a certain rate. For example, if the user translates the ligand molecule in the positive x direction to explore the parts of the protein molecule that are not currently accessible by AHW, the protein surface is translated in the negative x direction in the visual workspace to make this exploration possible. In addition, if the on/off switch on the haptic stylus is activated by the user at the same time, the protein surface is rotated about the y axis in a counterclockwise direction at a certain rate to allow the exploration of the back surface. The visual transformations of the protein molecule corresponding to the translational movements of the haptic probe and the switch settings are given in Table 2.

In order to further explain the concept of AHW, we define three coordinate frames: absolute (A), visual (V), and haptic (H). Assume that there exists a mapping between any two frames defined by a transformation matrix T . For example, the visual coordinates of the ligand molecule with respect to its absolute coordinates can be defined as

$${}^V LIGAND = {}^V T^A LIGAND \quad (8)$$

Table 2. The Visual Transformations of the Protein Molecule Corresponding to the Translational Movements of the Haptic Probe (Virtually Coupled to the Ligand Molecule) and the Switch Settings

Ligand translation	Protein translation*	Protein rotation**
Left (+ x)	Right (+ x)	CCW about y axis
Right (- x)	Left (- x)	CW about y axis
Up (+ y)	Down (+ y)	CW about x axis
Down (- y)	Up (- y)	CCW about x axis
Front (+ z)	Back (+ z)	CW about z axis
Back (- z)	Front (- z)	CCW about z axis

*Rate = 5 mm/s.

**Rotation occurs if the switch is on. Rate = 36 deg/s.

where ${}^A LIGAND$ represents the initial absolute coordinates of the ligand molecule extracted from the Protein Data Bank (<http://www.pdb.org>, Bergman et al. 2000) and ${}^V LIGAND$ represents the corresponding initial visual coordinates. Since the ligand molecule is virtually coupled to the haptic stylus, its visual coordinates are updated as the stylus is manipulated

$${}^V LIGAND_{current} = ({}^V_H T^H t_{stylus}) {}^V LIGAND \quad (9)$$

where ${}^H t_{stylus}$ represents the transformation matrix of the haptic stylus, and ${}^V LIGAND_{current}$ is the current visual coordinates of the ligand molecule. This transformation is carried out by the graphics processing unit (GPU) efficiently.

We define the AHW as a subspace of the physical haptic workspace having initially coincident origin. The matrix ${}^V t_{AHW}$ represents its transformation in the visual workspace. If the user exceeds the boundaries of AHW by translating the ligand molecule, the translational component of ${}^V t_{AHW}$ is updated incrementally at a constant rate of 5 mm/s until the user stops. If the switch on the haptic stylus is pressed at the same time, then the rotation component of ${}^V t_{AHW}$ is also updated at a rate of 36 deg/s (Table 2). Hence, the ${}^V t_{AHW}$ is updated as

$${}^V t_{AHW} = {}^V t_{increment} {}^V t_{AHW} \quad (10)$$

where ${}^V t_{increment}$ represents the total incremental transformation (translation plus rotation). Now, the part of the visual model of the protein surface that is accessible by the user for active haptic exploration is calculated efficiently by GPU as

$${}^V PROTEIN_{current} = {}^V t_{AHW}^{-1} {}^V PROTEIN \quad (11)$$

where ${}^V PROTEIN_{current}$ and ${}^V PROTEIN$ are the current and initial visual coordinates of the protein molecule. In order to calculate the molecular interaction forces between the ligand and protein molecules using the constants of the force field (i.e., parameters of the CHARMM) which are defined in the absolute coordinate frame, one must calculate the current coordinates of the ligand and protein molecules in the absolute coordinate frame. In other words, the current visual coordinates of the ligand (${}^V LIGAND_{current}$) and protein (${}^V PROTEIN_{current}$) molecules must be projected back to the absolute coordinate frame to calculate the interaction forces at the current time step. However, this projection can be computationally expensive, especially if the large number of atoms of the protein molecule is considered. While the transformations in the visual domain are carried out efficiently by the GPU, changing the absolute coordinates of a molecule requires an update in the actual database processed by the CPU. An alternative solution is to keep the absolute coordinates of the protein molecule unchanged, but update the absolute coordinates of the ligand molecule with respect to the fixed protein molecule

$${}^A LIGAND_{current} = ({}^A_V T^V t_{AHW}) ({}^A_H T^H t_{stylus}) {}^A LIGAND \quad (12)$$

Since the ligand molecule typically contains a much smaller number of atoms than the protein molecule, this approach involves fewer computations and can be handled more efficiently by the CPU. Now, the molecular interaction forces between the original absolute coordinates of the protein molecule (${}^A PROTEIN$) and the current coordinates of the ligand molecule in the absolute frame (${}^A LIGAND_{current}$) can be calculated and transformed back to the haptic coordinate frame using the following transformation

$${}^H\text{Force} = \text{Scale Factor} \times {}^H T({}^V R_{\text{AHW}})^{-1}({}^V T^A \text{Force}) \quad (13)$$

where ${}^V R_{\text{AHW}}$ represents the accumulated rotation in ${}^V t_{\text{AHW}}$. The *Scale Factor* is used to scale up the molecular interaction forces to the allowable range of forces that can be displayed by the haptic device.

6 Experimental Study

In order to investigate the proposed role of haptic feedback in molecular docking and visualization, we have designed and conducted an experimental study with six human subjects. In the study conducted by Brooks et al. (1990), the subjects (experienced biochemists) were asked to dock four ligand molecules to the true binding cavity of a protein molecule to discover which ligand molecule was the best choice. Our first experiment was designed to test if the subjects could find the true binding site of a protein molecule with the help of visual and haptic cues for a given ligand molecule and five potential binding sites. Our second experiment was designed to test if the haptic feedback could be used to roughly align the ligand molecule inside the true binding cavity such that the proposed rigid docking approach can be executed off-line to find its final configuration. The subjects in our experimental study were college students with no formal training in molecular biology. The design details and the results of the both experiments are given below.

6.1 Experiment I

The aim of our first experimental study was to test if the subjects could find the true binding site of a protein molecule for a given ligand molecule and five potential binding sites. The subjects were displayed six pairs of ligand-protein complexes (Table 3).

The subjects were asked to manipulate a ligand molecule and insert it into the five different cavities of a protein molecule one by one to find the true binding cavity. During this process, they felt the molecular interaction forces through the haptic device. Among the five

Table 3. The Ligand-Protein Complexes Used in Our Experiments

LP pairs (PDB ID)	Ligand molecule	Number of ligand atoms	Number of protein atoms
1BU4	2GP	24	782
1STP	BTN	16	901
1MFA	ABE	9	1712
3VGC	SRB	19	1738
1XIG	XYL	10	3031
1CSI	OAA	9	3391

different cavities, only one of them was the true binding cavity. The candidate sites were determined using Pocket in advance. Pocket is a public domain package developed by Edelsbrunner and Koehl (2005) that uses the Alpha Shape theory to detect cavities in a protein and ranks them according to their volume and surface area. For all of the six ligand-protein (LP) pairs used in our experimental study, the true binding site was among the first five cavities returned by Pocket (see Figure 6).

Before the experiment, the subjects were instructed about the haptic device and the molecular docking problem. They were asked to read a document that describes how to find the true binding site among five given candidates using visual and haptic cues. They were also shown slides summarizing what they read in the document. The subjects were told that

1. The shape complementarity is an important factor in finding the true binding site.
2. The magnitude of the net force at the true binding site is close to zero.
3. The true binding site generates a “tunneling effect” and pulls the ligand molecule toward the binding cavity.
4. The true binding site traps the ligand molecule and does not easily let it escape.

All subjects were also trained with the 3PTB complex (see Figure 7) before the actual experiments. With the help of an expert user, they were educated to find the true binding site using the instructions given above.

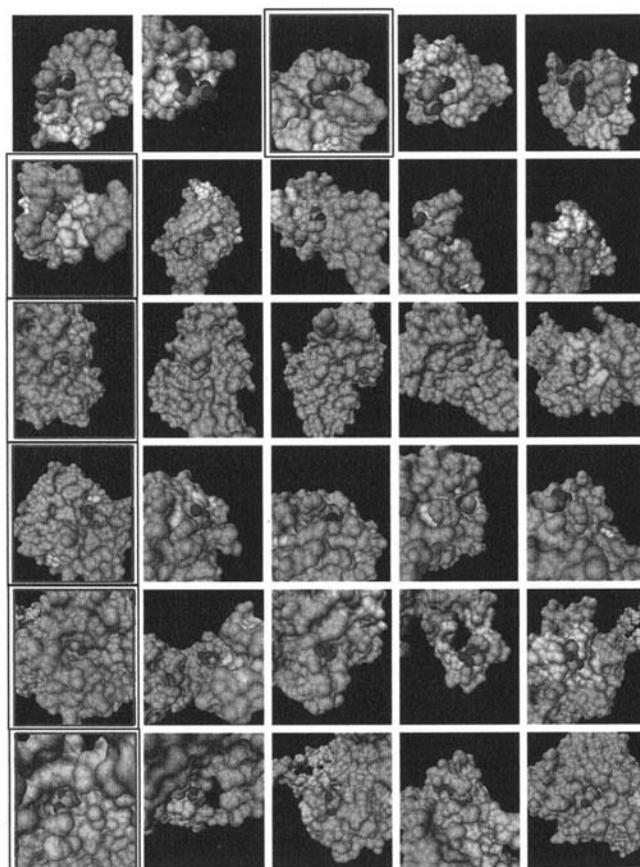


Figure 6. The molecular complexes and the cavities displayed to the subjects in the experiments. Each row (IBU4, ISTP, IMFA, 3VGC, IXIG, ICSI) corresponds to a ligand-protein couple and shows the five different cavities of the protein molecule displayed to the subjects. The true binding cavity was highlighted using a solid frame around one of the cavities in each row. The cavities are ordered in the figure from left to right according to the ranking returned by Pocket. Pocket successfully ranked the true binding site as the first cavity in all complexes except one (IBU4).

During the actual experiments, a total of 30 sites (6 pairs \times 5 cavities) were displayed to the subjects in random order. The subjects repeated the same experiment after one day rest period. The second time, the same cavities were displayed to the subjects in different order and spatial orientations to reduce the visual bias.

During the experiments, subjects were allowed to navigate freely between the five cavities labeled as A, B, C, D, and E by pressing the left and right arrows on the

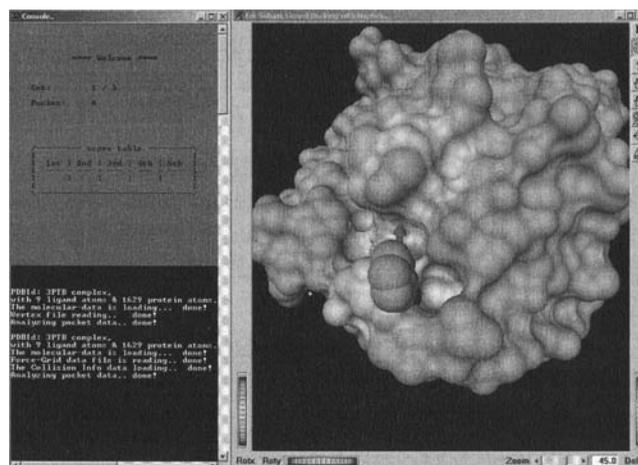


Figure 7. The screen capture of the graphical interface used in the experiments. The simulation window is displayed on the right side of the screen. On the upper left is displayed the experimental trial number, the alphabetical ID of the cavity being explored (A, B, C, D, or E), and the cavities ranked by the subject. The command window is displayed on the lower left.

keyboard. We implemented a fly-through from one cavity to the other. The polygons around the entrance and inside each cavity were highlighted with a unique color different from the color of the protein surface to help the user locate the cavities more easily. The colors assigned to the cavities were also randomized in each trial to reduce the visual bias. The subjects were asked to rank the cavities by pressing “1” (most likely binding site), “2,” “3,” “4,” and “5” (least likely binding site) keys on the keyboard. The ranking was displayed on the screen and subjects were allowed to update their ranking at any time (i.e., they were allowed to retest any of the five cavities) before start working on a new ligand-protein pair by pressing the “N” key on the keyboard.

The results of the first experiment show that the subjects successfully discovered the true binding site with the help of visual and haptic cues (Table 4).

In Table 4, the raw scores were obtained by simply counting the number of times the true binding site was discovered by the subjects. The penalty-based scores are more conservative and they were calculated by assigning a penalty score to the wrong choices made by the sub-

Table 4. *The Results of the First Experiment*

LP pairs (PDB ID)	First day (raw score)	Second day (raw score)	Combined (raw score)	First day (penalty-based score)	Second day (penalty-based score)	Combined (penalty-based score)
1BU4	5/6	5/6	10/12	7/6	8/6	15/12
1STP	6/6	6/6	12/12	6/6	6/6	12/12
1MFA	3/6	3/6	6/12	15/6	10/6	25/12
3VGC	4/6	4/6	8/12	9/6	9/6	18/12
1XIG	5/6	3/6	8/12	7/6	9/6	16/12
1CSI	3/6	3/6	6/12	9/6	10/6	19/12
Total	26/36	24/36	50/72	53/36	52/36	105/72
Success rate	72%	67%	69%	—	—	—
Average score	—	—	—	1.47	1.44	1.46

jects. For example, a penalty score of “5” was used in the calculation of total score if the subject had selected the true binding site as his or her last choice. The ranking average of the subjects is 1.47 ± 0.36 in the first set and 1.44 ± 0.27 in the second set. The average was calculated using the conservative penalty-based approach. These average values show that the subjects have ranked the true binding site as their first choice in most of the pairs. After the experiments, a simple questionnaire, made of ten questions (five questions related to the role of visual information and five questions related to the role of haptic information) was given to each subject to better understand his or her perception of the individual role of visual and haptic cues in discovering the true binding cavity. In particular, we wanted to find out if the subjects had utilized the haptic cues effectively during the experiments. All questions began with the phrase “How much did . . .” and ended with the phrase “. . . help you find the true insertion site?” (e.g., How much did the size and depth of a cavity help you find the true binding site? How much did the effect of being pulled inside a cavity help you find the true binding site?) Each question was rated on a scale that varied from 1 (very little) to 5 (very much). The results of the questionnaire, when combined with the results of the experiment, show that subjects used the haptic cues

(4.33 ± 0.8) more than the visual ones (3.67 ± 0.8) in finding the true binding cavity ($p < .05$).

6.2 Experiment II

In the second set of experiments, only the true binding cavity was displayed to the subjects for each pair, and they were asked to find the binding configuration (both position and orientation) of the ligand molecule inside the true binding cavity using the visual and haptic cues. The total energy of the interactions during haptic manipulations was displayed on the screen to help him or her find the low energy configuration. Subjects were simply told that they should adjust the position and orientation of the ligand molecule inside the binding cavity until the magnitudes of the net interaction force displayed through the haptic device and the interaction energy displayed on the screen were close to minimal values. As the subjects manipulated the ligand molecule, the total energy of the complex was calculated at the graphics update rate of 30 Hz and displayed on the screen. When the ligand molecule passed through a low energy configuration during the exploration of the binding cavity, a visual copy of the ligand molecule was left as a pointer to draw the attention of the subject (see Figure 8). When the subject discovered

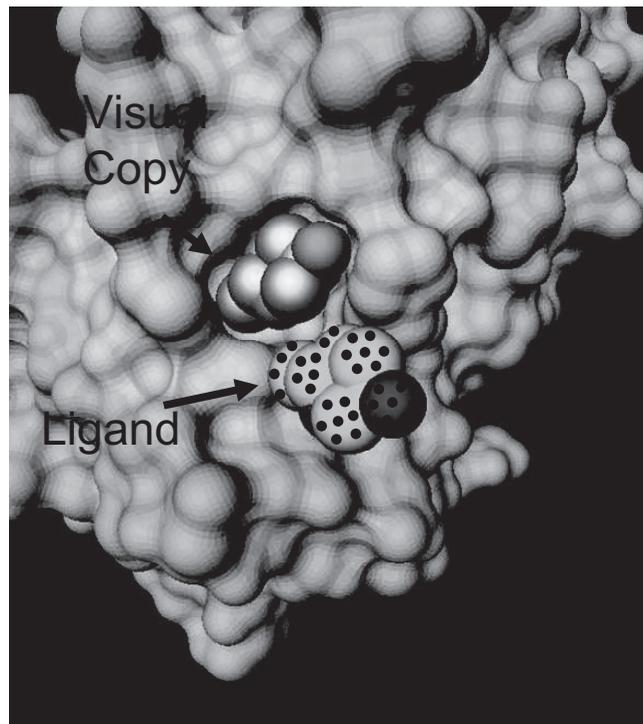


Figure 8. In the second experiment, when a ligand molecule (dotted atoms) passed through a configuration lower in energy than its current configuration during the exploration of the binding cavity, a visual copy of the ligand (shiny atoms) was displayed at that configuration to draw the attention of the subject.

a lower energy configuration than the previously stored one, the pointer was updated to inform the subject. The subjects could easily turn this feature on and off by pressing the “Backspace” key on the keyboard.

The experiments were conducted with the same molecular complexes used in the first experiment. The true binding cavity of each pair was displayed to the subjects in two different orientations to reduce the visual bias. A total of 12 cavities (6 pairs \times 2 orientations) were displayed to the subjects in random order, with the same order displayed to each subject. The subjects pressed the “Enter” key to finalize the docking process after each trial. The binding configuration of the ligand molecule obtained by the haptic exploration was saved to a text file after each trial in the form of a transformation matrix to find its final configuration later using the pro-

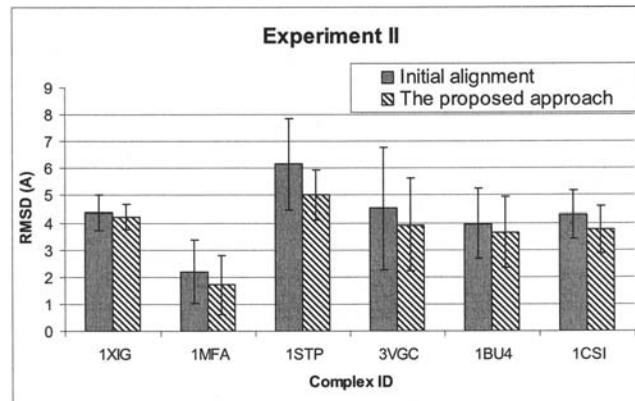


Figure 9. The results of the second experiment.

posed rigid docking approach (Section 3). The results of the second experiment show that subjects can roughly align the ligand molecule inside the binding cavity using visual and haptic cues and this initial alignment can be used by the MD simulations to further improve the results (Figure 9).

7 Conclusion and Discussion

The work presented in this paper demonstrates a novel application of human-computer interaction in molecular docking. The proposed approach aims to integrate the strengths of both worlds to achieve better results. It is argued that humans are better than computers at complex assembly and disassembly tasks involving insertion and removal of parts. In our approach, this corresponds to manipulating a ligand molecule with a haptic device in a virtual environment to search for the true insertion site on the surface of a protein molecule. The user inserts the ligand molecule into the potential cavities of the protein molecule one by one to find the true binding cavity. The true binding cavity generates a tunneling effect and pulls the ligand molecule towards the inside and does not easily let the ligand escape from there. Moreover, the net force and interaction energy are close to minimal values at the true binding configuration. To help the user explore the cavities of the protein molecule effectively and interactively using the

haptic device, we proposed the AHW concept. This approach enables the user to move the haptic workspace actively anywhere on the surface of a large protein surface for haptic visualization. This is achieved through highly efficient coordinate transformations by taking advantage of the GPU and leaving less work to the CPU. Since the workspace of haptic devices that are commercially available today is limited by their physical dimensions, the proposed approach is highly effective in haptic visualization of a large and complex 3D surface such as the surface of a protein molecule. The actual dimensions of a protein molecule are on the order of a few hundred angstroms and the spatial resolution achieved by scaling up the original atom coordinates such that the protein fits into the small haptic workspace is not sufficient for the detailed exploration of the potential binding sites.

While humans are good at executing manipulation tasks involving part insertion and removal as in finding the true insertion site of a ligand molecule, computers are superior to humans in tasks involving precision and accuracy. In our approach, the final configuration of the ligand molecule inside the true binding cavity is calculated by the computer through time stepping off-line MD simulations after it has been placed and roughly aligned by the user. In each time step, our new docking algorithm minimizes the distance error between the previous rigid body coordinates of the ligand atoms and their current coordinates obtained from the MD simulations to calculate the new rigid body coordinates. As a result, the ligand molecule makes rigid body movements and travels toward the lowest potential configuration inside the cavity. Our rigid docking algorithm is computationally more efficient than the ones utilizing conventional rigid body equations. Moreover, the ligand molecule is less likely to become trapped in a local minimum in our approach, since a good initial configuration for the MD simulations is supplied by the user via the haptic device. The conventional energy minimization approaches to perform the same task without a good initial guess typically suffers from being computationally too intensive or getting trapped in local minima more often (Apaydin et al., 2002).

The results of our human experiments with six sub-

jects testing six different molecular pairs show that haptic feedback is effective in selecting the true binding site among multiple candidates. It also contributes to the rough alignment of the ligand molecule inside the binding cavity, such that its final configuration can be determined via off-line MD simulations later. To identify the true binding site among five candidates, the subjects used both visual and haptic cues. We observed that the subjects learned to eliminate one or at most two site(s) by visual inspection alone, but it was difficult for them to find the true binding site using visual cues only. There was sometimes a geometric match between a ligand molecule and a candidate cavity, but if the ligand molecule was not pulled toward the inside of cavity or could escape from it easily after being inserted, the site was not the true binding site. For example, Pocket incorrectly ranked the true binding site for the complex 1BU4 (first row in Figure 6) as the third best candidate based on the geometric information, but the subjects were successful (10 out of 12 trials) in identifying its location correctly using visual and haptic cues together (see Table 4). We also investigated whether MD simulations could be used to find the final configuration of the ligand molecule inside the binding cavity if a good initial configuration for the simulations was supplied by the subjects effectively with the help of the haptic device. The results of the second experiment show that the proposed docking approach reduces the RMS error in all pairs, though the change is statistically significant only in some of the pairs. This area requires further investigation.

The proposed rigid docking approach has some shortcomings. First, the rigidity assumption has some disadvantages though it has been frequently used with success in docking small ligand molecules. A rigid model does not allow for energy exchange with the environment. In addition, rigid models are unable to display the conformational changes that some molecules exhibit upon binding. It is argued that conformational changes in both protein and ligand are necessary for a successful docking process. While the flexibility of a small size ligand molecule can be modeled successfully, modeling the flexibility of the protein is still far beyond the present computational capability of the existing docking programs

(Teodoro et al., 2001). Third, we neglect the contribution of hydrogen bonds in our approach. Two interacting molecules can also make a hydrogen bond (if one has a hydrogen atom and interacts with the atoms of the other) during binding, which significantly strengthens the interactions between them and can be used as a constraint to further improve the results of final docking simulations. Fourth, our haptic device does not allow the display of torques, which hinders the user from feeling coupling moments as he or she rotates the ligand molecule inside the binding cavity. The haptic device used in our experiments has 6 DOF sensing, but only 3 DOF force output capability. Hence, the user can only feel the net forces acting on the ligand molecule during the simulations. We anticipate that additional torque feedback would improve the quality of haptic feeling, especially close to the true binding configuration. Finally, the binding cavity contains local minima and thus may create a trapped situation for the ligand molecule during the MD simulations. In general, the number of local minima increases rapidly with the number of atoms. In our approach, the user finds the true binding cavity and then roughly aligns the ligand molecule inside the cavity. Hence, the MD simulations start with a good initial configuration of the ligand molecule which reduces the chance of becoming trapped in a local minimum.

Acknowledgments

We acknowledge the technical discussions made with Ozlem Keskin, Attila Gursoy, and Burak Erman of Koc University.

References

- Akkiraju, N., Edelsbrunner, H., Fu, P., & Qian, J. (1996). Viewing geometric protein structures from inside a CAVE. *IEEE Computer Graphics Applications*, 16, 58–61.
- Anderson, A., & Weng, Z. (1999). VRDD: Applying virtual reality visualization to protein docking and design. *Journal of Molecular Graphics and Modeling*, 17, 180–186.
- Apaydin, M. S., Guestrin, C., Varma, C., Brutlag, D. L., & Latombe, J. C. (2002). Stochastic roadmap simulation for the study of ligand-protein interactions. *Bioinformatics*, 18, S18–S26.
- Basdogan, C., & Srinivasan, M. A. (2002). Haptic rendering in virtual environments. In K. Stanney (Ed.), *Handbook of virtual environments* (pp. 117–134). Mahwah, NJ: Lawrence Erlbaum Associates.
- Bayazit, O. B., Song, G., & Amato, N. (2001). Ligand binding with OBPRM and haptic user input. *IEEE International Conference on Robotics and Automation*, 954–959.
- Berman, H. M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T. H., Weissig, H., et al. (2000). The protein data bank. *Nucleic Acids Research*, 28, 235–242.
- Birmanns, S., & Wriggers, W. (2003). Interactive fitting augmented by force-feedback and virtual reality. *Journal of Structural Biology*, 144, 123–131.
- Brooks, F. P., Ouh-Young, M., Batter, J. J., & Kilpatrick, P. J. (1990). Project GROPE: Haptic displays for scientific visualization. *SIGGRAPH*, 24(4), 177–185.
- Cole, J. C., Murray, C. W., Nissink, J. W., Taylor, R. D., & Taylor, R. (2005). Comparing protein-ligand docking programs is difficult. *Proteins*, 60(3), 325–332.
- Connolly, M. L. (1983). Solvent-accessible surfaces of proteins and nucleic acids. *Science*, 221, 709–713.
- Conti, F., & Khatib, O. (2005). Spanning large workspaces using small haptic devices. *Proceedings of the World Haptics Conference*, 183–188.
- Dominjon, L., Lecuyer, A., Burkhardt, J., Andrade-Barroso, G., & Richier, S. (2005). The “bubble” technique: Interacting with large virtual environments using haptic devices with limited workspace. *Proceedings of the World Haptics Conference*, 639–640.
- Edelsbrunner, H., & Koehl, P. (2005). The geometry of biomolecular solvation. *Combinatorial and computational geometry*, 52, 243–275.
- Gillilan, R. E., & Lilien, R. H. (2004). Optimization and dynamics of Protein-protein complexes using B-splines. *Journal of Computational Chemistry*, 25(13), 1630–1646.
- Guthold, M., Falvo, M., Matthews, W. G., Paulson, S., Mullin, J., Lord, S., et al. (1999). Investigation and modification of molecular structures with the nanoManipulator. *Journal of Molecular Graphics and Modeling*, 17(3–4), 187–197.
- Halperin, I., Ma, B., Wolfson, H., & Nussinov, R. (2002). Principles of docking: An overview of search algorithms and a guide to scoring functions. *Proteins*, 47(4), 409–443.
- Hollis, R., & Salcudean, S. E. (1993). Lorentz levitation tech-

- nology: A new approach to fine motion robotics, teleoperation, haptic interfaces, and vibration isolation. *Proceedings of the 5th International Symposium on Robotics Research*, 1–18.
- Kuntz, I. D. (1992). Structure-based strategies for drug design and discovery. *Science*, 257, 1078–1082.
- Lai-Yuen, S. K., Lee, Y.-S. (2006). Energy-field optimization and haptic-based molecular docking and assembly search system for computer-aided molecular design (CAMD). *The 14th IEEE International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 233–240.
- Lee, Y., & Lyons, K. (2004). Smoothing haptic interaction using molecular force calculations. *Computer-Aided Design*, 36(1), 75–90.
- Levine, D., Facello, M., Hallstrom, P., Reeder, G., Walenz, B., & Stevens, F. (1997). Stalk: An interactive system for virtual molecular docking. *Computing in Science and Engineering*, 4, 55–65.
- Nagata, H., Mizushima, H., & Tanaka, H. (2002). Concept and prototype of protein-ligand docking simulator with force feedback technology. *Bioinformatics*, 18(1), 140–146.
- Ouh-Young, M., Beard, D. V., & Brooks, F. P., Jr. (1989). Force display performs better than visual display in a simple 6-D docking task. *IEEE International Conference on Robotics and Automation*, 1462–1466.
- Ouh-Young, M., Pique, M., Hughes, J., Srinivasan, N., & Brooks, F. P., Jr. (1988). Using a manipulator for force display in molecular docking. *IEEE International Conference on Robotics and Automation*, 1824–1829.
- Peterlík, I., & Krenek, A. (2005). Haptically driven traveling through conformational space. *Proceedings of the World Haptics Conference*, 342–347.
- Pocket*: A program for detecting pockets and cavities in proteins, <http://csb.stanford.edu/~koehl/ProShape/>, September 27, 2007.
- Rapport, D. C. (2002). *The art of molecular dynamics simulation*. Cambridge, UK: Cambridge University Press.
- Sankaranarayanan, G., Weghorst, S., Sanner, M., Gillet, A., & Olson, A. (2003). Role of haptics in teaching structural molecular biology. *The 11th IEEE Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 363–366.
- Sharma, A., Nakano, A., Kalia, R. K., Vashishta, P., Kodiyalam, S., Miller, P., et al. (2003). Immersive and interactive exploration of billion-atom systems. *Presence: Teleoperators and Virtual Environments*, 12(1), 85–95.
- Stone, J. E., Gullingsrud, J., Schulten, K., & Grayson, P. (2001). A system for interactive molecular dynamics simulation. *ACM Symposium on Interactive 3D Graphics*, 191–194.
- Subasi, E., & Basdogan, C. (2006). A new approach to molecular docking in virtual environments with haptic feedback. *Proceedings of EuroHaptics Conference*, 141–145.
- Taylor, R., Hudson, T. C., Seeger, A., Weber, H., Juliano, J., & Helser, A. T. (2001). VRPN: A device-independent, network-transparent VR peripheral system. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, 55–61.
- Teodoro, M., Phillips, G. N., & Kavraki, L. (2001). Molecular docking: A problem with hundreds of thousands of degrees of freedom. *IEEE International Conference on Robotics and Automation*, 960–965.
- Timothy, R., & Forester, W. S. (1998). SHAKE, Rattle, and Roll: Efficient Constraint Algorithms for Linked Rigid Bodies. *Journal of Computational Chemistry*, 19(1), 102–111.
- Van Gunsteren, W. F., & Berendsen H. J. C. (1990). Computer simulation of molecular dynamics: Methodology, applications, and perspective in chemistry. *Angewandte Chemie International Edition in English*, 29, 992–1023.