Crowdsourced Molecular Docking Using Path-Planning and Haptic Devices

Torin Adamson, John Baxter, Kasra Manavi, Bruna Jacobson, and Lydia Tapia
Department of Computer Science
University of New Mexico
Albuquerque, New Mexico, 87131
Email: toriadam@cs.unm.edu, automata@unm.edu, kazaz@cs.unm.edu, bjacobson@unm.edu, tapia@cs.unm.edu

***Abstract*—Many biological processes, including immune recognition, enzyme catalysis, and molecular signaling, are driven by molecular binding. The binding of a protein receptor to a ligand is still an open problem in biological sciences. We present a method, the Haptic-Guided Ligand Docking Tool, that explores the binding between a protein receptor and ligand by utilizing crowdsourced haptic-guided motion planning. Haptic touch devices enable users to feel the interactions of the two molecules as they move the ligand into an appropriate binding site on the receptor. We demonstrate the method on a system critical for human immune response, the binding of ligands to Major Histocompatibility Complex (MHC) molecules. From the multiple runs by the users, our approach is able to construct a global roadmap that finds low energy paths to molecular docking sites. We find that the crowdsourcing results are comparable to the highly-biased case of roadmaps generated by Gaussian conformational sampling around a known docked ligand conformation.**

I. Introduction

Many biological processes rely on the interactions between small-molecule ligands and large protein receptors. Ligand-receptor binding prediction is therefore critical to biochemical engineering. Ligands bound in a low energy conformation are said to be *docked*. Binding affinities are controlled by the energetic feasibility of transition from undocked to docked conformations. Motion planning techniques can predict interactions between molecules efficiently without falling into local minima traps [12] while screening techniques only consider the validity of final docked conformations.

Molecular docking methods are more computationally efficient than Molecular Dynamics (MD) methods for predicting ligand and receptor binding. Efficiency from reduced-complexity models make interactive components (e.g., haptics) feasible. Many molecule docking tools keep receptors and/or ligands rigid during simulation. While such simplifications prevent identification of some properly docked conformations, they enable interactivity by improving runtime efficiency.

Haptic-based molecular docking simulators can be efficient tools for studying molecular interactions and docking. In this work we develop a method to generate ligand docking pathways using motion planning techniques integrated with realtime haptic simulations, similar to previous realtime haptic adaptations running on commodity hardware [6]. Probabilistic Roadmap Methods (PRMs) have been applied to motion planning problems with high-dimensional space of conformation, such as protein folding [29, 1] and ligand binding [3] applications. PRMs are easy to apply, only needing the ability to sample conformations and evaluate conformational feasibility. Conformations with the ligand bound to the receptor are low in potential energy, allowing for the use of potential energy to test for feasibility.

Finding solutions to biological problems using crowdsourcing has been successfully implemented in popular programs such as Foldit [22, 11] and Folding@home [26, 4] for protein folding. Crowdsourcing utilizes the public to generate a large set of data, and in some applications as users become more familiar with the software, they can improve their performance which may result in achieving low energy conformations.

In this study, we construct pathways of ligands moving into binding sites generated using two different methods of construction. In the first method, users operate a haptic device with force-feedback to explore the feasible space of conformations. Multiple user inputs are combined to generate a roadmap that can be traversed to identify feasible paths. The second method uses a Gaussian distribution sampling to generate conformations, centered around a known docked conformation of the ligand. As this method is biased towards the native binding site, it provides us with a "best-case scenario" result that we can use to compare to the outcomes from the Haptic-Guided Ligand Docking Tool.

This novel approach also incorporates a reduced polygon model for visualization efficiency, intuitive docking, and small data storage while computing energy and forces from the all-atom structures. These simplifications and multi-threading allows for fast real-time use on commodity hardware. The haptic device used in this work is the Novint Falcon®, a commodity haptic device with three degrees of freedom. However, the Haptic-Guided Ligand Docking Tool is not limited to one kind of haptic device as it applies general methods and is therefore expandable to multiple platforms.

II. Related Work

A. Molecular Docking

Molecular simulations and physical experiments are costly and time-consuming, so fast accurate methods sampling and scoring ligand-receptor binding candidates help prioritize limited scientific resources. A wide array of molecular docking algorithms and tools have been developed (see [20]).

Early docking tools used ligand-receptor cavity geometry complementarity to both guide sampling and score feasibility [23]. Later tools incorporate atomic force fields into scoring [15]. Methods to discover and measure ligand/cavity complementarity have also become more sophisticated [9].

Relative to rigid body approximations, accounting for ligand and receptor flexibility greatly increases docking problem dimensionality. Techniques applied to the problem of docking with flexible molecules include incremental construction by tree search [28], genetic algorithms [21], and Monte Carlo energy minimization [25]. Receptor flexibility may also be handled as ligand binding to rigid receptor ensembles [10].

Most docking tools automate conformation sampling; in contrast, our method utilizes user guidance to discover feasible ligand trajectories. The force field used to score samples additionally provides haptic feedback informing the user's search. We present a system that merges conformations discovered by separate users to produce better quality trajectories than any individual user produced.

B. Molecular Docking With Haptics

Haptic devices can enhance the operator's intuition and understanding of molecular binding processes [5]. They can assist automated docking algorithms with "hints" from user-sampled ligand conformations combined with existing sampled conformations into a single roadmap [2]. Haptic feedback has also been incorporated into docking simulations with adaptive user control of the flexibility/performance tradeoff [6]. Haptic controls directly manipulating ligand position have been compared against force-based control schemes [7]. These devices have also been applied to the control of "probe" objects such as water molecules to discover solvent accessible locations [30].

Molecular docking tools developed for different device capabilities include 3-DOF force feedback and 6-DOF force/torque feedback [19]. Collaborating users simultaneously manipulating separate molecules has been considered [17]; the method presented here, by contrast, synthesizes conformation sampling performed independently by multiple users into ligand binding trajectories.

We implement haptic force-feedback as the gradient of potential energy, with force scaled and time-smoothed to prevent unstable haptic feedback. Unlike [17], there are no dead zones; haptic feedback always reflects the potential energy gradient. Responsiveness is maintained on commodity hardware by computing potentials and haptic-feedback on separate threads. A complex energy model is handled similarly in realtime in [13]; but this was not performed on commodity hardware.

C. Motion Planning With Haptics

Haptics are well suited for integration with motion planning problems and have been used to give guiding hints to existing motion planners [3, 16]. These devices can also be used by guide or train the user [31, 24].

In [31], Kautham path planning (configuration space cellular decomposition/classification) is used to generate a "local channel" (path). This local channel is used to generate haptic force feedback to aid the user during execution of a task.

In our method, force on the ligand is felt since potential energy gradient generates the haptic force-feedback. The user makes the decisions for ligand movement based on this force-feedback, taking advantage of their intuition during docking. Sampled ligand conformations in this process are used in probabilistic roadmap construction.

III. Implementation and Methods

A. Molecular Simulation

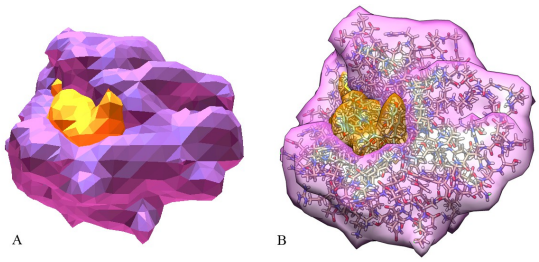

Fig. 1. Ligand (orange) with receptor (purple). The potential docking cavity can be seen as an actual cavity in the receptor's isosurface. (A) The actual appearence of ligand Tel1P and receptor MHC in Haptic-Guided Ligand Docking Tool; (B) underlying atoms with transparent isosurface of ligand Tel1P bound in MHC.

In our simulation, the receptor and the ligand are represented as rigid body isosurfaces with static internal atoms. This rigid body representation reduces complexity for runtime performance but prevents Haptic-Guided Ligand Docking Tool from finding docked conformations that require flexibility. The Tel1P and MHC chains from 3H9S combined have 3027 atoms. The receptor is fixed in place while the ligand is free for the user to move. The underlying set of atoms shown in Figure 1(B) are used for the potential energy approximation, but are only shown to the user as isosurfaces in Figure 1(A). Drawing only the isosurface representation decreases the time spent drawing the scene and simplifies the problem visually for the user. The colors chosen for the molecules in Haptic-Guided Ligand Docking Tool are arbitrary and only used to visually differentiate between the ligand and receptor. The isosurface models were generated from PDB files using Chimera with a resolution setting of 2 for the ligand and 3 for the receptor [27]. The ligand atoms translate and rotate as a rigid body when the user moves the ligand. The structures for the human class I MHC molecule (receptor) bound to Tel1p (ligand) were taken from the RCSB Protein Data Bank [8] (PDB 3H9S). MHC is used due to its diverse binding and importance in immune system activation. The missing hydrogen atoms (due to the nature of X-Ray Crystallography) were inserted using Chimera [27] and its built-in "Add H" tool.

B. Energy Approximation Function

We calculate intermolecular potential energy U_{inter} (1) between receptor R and ligand L as the sum of all pairwise

electrostatic U_{es} (2) and Lennard-Jones U_{vdw} (3) atomic interactions of receptor atoms i and ligand atoms j:

$$U_{\text{inter}}(R,L) = \sum_{i}^{R}\sum_{j}^{L} U_{\text{es}}(i,j) + U_{\text{vdw}}(i,j), \tag{1}$$

$$U_{\text{es}}(i,j) = C\frac{q_i q_j}{r_{ij}}, \tag{2}$$

$$U_{\text{vdw}}(i,j) = \sqrt{\varepsilon_i \varepsilon_j}\left[\left(\frac{\rho_i+\rho_j}{r_{ij}}\right)^{12} - 2\left(\frac{\rho_i+\rho_j}{r_{ij}}\right)^{6}\right]. \tag{3}$$

In the above equations, r_{ij} is interatomic distance and C the electrostatic constant. Our current implementation uses values for partial charges q_i, Lennard-Jones well depths ε_i, and Lennard-Jones minimal distances $\rho_i = 2^{\frac{1}{6}}\sigma_i$ from the AMBER force field [14].

The intermolecular potential energy is used directly to rank ligand-receptor conformations. (Because we make the rigid body assumption for both molecules, intramolecular interactions are not calculated.) This value is displayed to users which they are challenged to reduce as much as possible. By this mechanism, users are encouraged to manipulate the ligand to discover local and global potential minima.

The force approximation used for feedback is calculated from the gradient of the potential approximation. For torque, the cross product between each ligand atom's displacement vector (from the center of mass) and the force from the interaction between the ligand atom and each receptor atom can be used, similar to [18]. However, torque and force are not handled at the same time due to the limitations of the particular three axis haptic device. The operator can hold one button down for translation movement and force feedback, or a different button for angular movement and torque feedback. These calculations are done using an all-atom cloud model between ligand and receptor.

High DOF haptic devices may not be available for some users. The method to convert user input and force feedback can be adapted for other devices that support various methods of input and force feedback. A device might allow the user to touch the ligand on screen, or might have a vibration motor for force feedback. A deployment of this method for a large amount of users would have to adapt to many input/output devices available to the users.

C. Force Feedback

The energy potential approximation is highly sensitive to the position of the atoms due to the nature of the Lennard-Jones potential. Because of these large differences between low and high energy approximation values, a logarithmic scaling function (4) is used to bring the values into a smaller range for force-feedback. Energy values can be scaled logarithmitically to reduce the sudden differences in energy as seen in Figure 2:

$$E_s(E) = \begin{cases} \ln(E), & \text{if } E \geq 1 \\ -\ln(2-E), & \text{if } E < 1 \end{cases}, \tag{4}$$

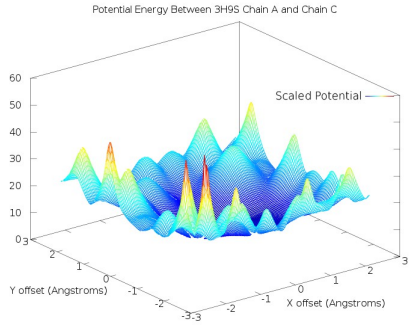

Fig. 2. Scaled potential energy approximation near the native bound conformation of human class I MHC receptor bound to ligand Tel1P (PDB 3H9S)

where E is the original energy value and E_s is energy rescaled. To prevent sudden device tremors and confusion with force-feedback, the feedback vector is time-smoothed according to:

$$\vec{v}_t = S(\vec{v}_i - \vec{v}_{t0}), \tag{5}$$

$\vec{v}_t$ is the new force-feedback direction, $\vec{v}_i$ is the eventual target feedback direction, and $\vec{v}_{t0}$ is the current force-feedback direction. S is an adjustable factor that increases or decreases the speed at which $\vec{v}_t$ reaches $\vec{v}_i$. This proportional time delay applied to each frame of the program maintains the need for quick changes in feedback for extreme energy differences, as well as smooths out the general noise in the scaled energy approximation. Higher values of S "tighten" the force-feedback with $S = 1$ resulting in no time delay. A value of $S = \frac{1}{2}$ was used in the study to balance stability with responsiveness.

After scaling and proportional time delay, the resulting vector is passed to the haptic device for output. With all of these combined, a maximum threshold for force output through the haptic device is unnecessary. The force-feedback effect strength scale can be adjusted for the user's comfort. In addition to force-feedback scaling, the user can zoom in and out to increase or decrease the device sensitivity.

The 3 degree of freedom haptic force feedback device described may not be available to all users, so adaptations may be needed for other input devices. Vibration force feedback motors are present in many mobile phones and video game controllers, thus requiring the three dimensional force feedback vector to change to a scalar response.

Fig. 3. Users operating the haptic device and laptop running Haptic-Guided Ligand Docking Tool.

D. Roadmaps

Roadmaps are constructed similar to the PRM method. However, specific user ligand conformations are used as input conformations. First, conformations are sampled from Haptic-Guided Ligand Docking Tool by recording user ligand conformations where each conformation is no more than 0.1Å apart. Ligand conformations of an energy greater than a high potential energy threshold E_{MAX} are not recorded.

An edge between two ligand conformations (c_1, c_2) is weighted by a function of the difference between the maximum potential energy among interpolated ligand conformations between the start and end conformation, $c_1 = s_0, s_1, ..., s_n = c_2$, and the initial potential energy $E(c_1)$. The edge weight, $W_{i,j}$, is $\ln(\Delta E + 1)$ where the difference in energy, ΔE, is $max(E(s_0), ..., E(s_n)) - E(c_1)$. Therefore, edges of decreasing potential energy are given a weight of 0, otherwise the weight reflects an energetic traversal cost. This is needed to identify shortest paths using Djikstra's algorithm. Edges are calculated for every pair of conformations in both directions. New roadmaps are built from existing roadmaps by appending them with new user sets using the incremental roadmap generation method [32].

IV. Results

A. Performance

In order to quantify the computation time for the interactive system, we captured runtimes that reflected the potential calculation and the impact of model rendering and resolution. Our method uses multiple threads to improve overall performance and provide smooth force-feedback. One thread repeatedly updates potential energy, averaging about 12 calculations per second. Force-feedback and scene drawing are handled on another thread, using the most recent results of the potential energy thread to calculate force. The scene is drawn using the isosurface representation for higher performance.

TABLE I
Runtime Performance of Main Thread With Isosurface Model Polygon Count

Resolution (Chimera setting)	Polygons in Isosurfaces	Time per Frame (ms)
-	0 (No Drawing)	18
3	3160	21
2	7840	23
1	60184	51

The polygon count in the isosurface vs. main thread performance can be seen in Table I. Recall that model resolution can be adjusted without affecting the energy and force calculations. In Table I, the first entry is the baseline performance, force feedback and program overhead, when drawing no models. We also studied the impact of model resolution. In Table I resolution corresponds to the isosurface model resolution setting from Chimera. Polygon triangles are more efficient for GPUs to render than realistic atomic spheres, of which Tel1P and MHC combined have 3027 atoms. Together, this multi-threaded environment produces a real-time sensation of the atomic forces modeled in the energy approximation on commodity hardware with a visual and haptic touch feedback frame rate of about 42 frames per second (about 23 milliseconds per visual and haptic touch feedback frame) using 61.4MB memory on a commodity laptop with an AMD A6-5200 APU chipset with a 4-core CPU 2GHz clock rate and Radeon HD 4800 GPU.

B. Haptic-Guided User Sampled Conformations

In order to capture conformations for roadmap generation, data from three users with two runs per user were recorded. Each user run contributed 1000 conformations, shown in Figure 4. It is interesting to notice that different strategies are implemented by users, represented by distinct colors in Figure 4. The inset shows details around negative energy conformations. There are particular locations where users would focus exploration before investigating other locations, such as "User 3" (blue) between 2 and 4 Å RMSD and the cluster generated by several users near the native conformation.Also, the area around the native conformation had been densely explored despite the fact that users were not explicitly aware of the native conformation location.

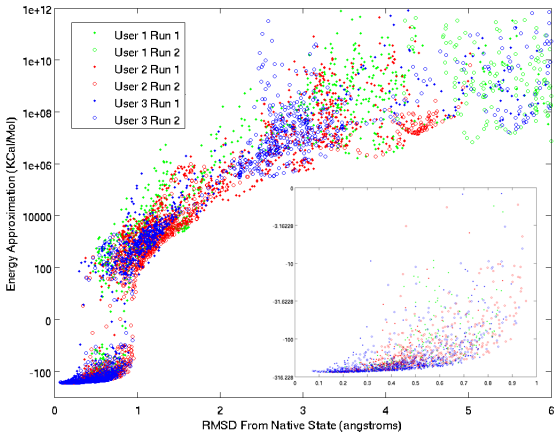

Fig. 4. RMSD and potential energy for automatic and manual haptic-guided conformations. The energy is shown on a logarithmic scale.

C. Multi-user Roadmaps

Once we obtain sets of different conformations from all user runs, they can be combined to generate roadmaps, as explained in Section III-D. The samples from Section IV-B were used to build the roadmap. Each subsequent user extended the roadmap iteratively. Table II shows the number of conformations and edges created for each run.

For comparison against the roadmap of haptic-guided conformations, a roadmap was built with 6000 Gaussian distributed rigid-body ligand conformations centered at a mean of 0Å when the ligand is in a docked conformation, with 5Å in translational and 5° in rotational standard deviations. 88,758 weighted edges were created between these conformations.

All roadmap conformations were connected by nearest neighbors ($K = 10$) using the scaled euclidean distance metric,

TABLE II
ROADMAPS FROM HAPTIC-GUIDED CONFORMATIONS AND GAUSSIAN DISTRIBUTED CONFORMATIONS.

Data type	Cumulative sets (User,Run)	Conformation Count	Edge Count
Haptics	(1,1)	1000	12516
	(1,1),(2,1)	2000	25246
	(1,1),(2,1),(1,2)	3000	37424
	(1,1),(2,1),(1,2),(3,1)	4000	50240
	(1,1),(2,1),(1,2),(3,1),(2,2)	5000	63126
	(1,1),(2,1),(1,2),(3,1),(2,2),(3,2)	6000	75750
Gaussian	-	6000	88758

and then a K-pair = 10 component connection method to connect "islands" of connected conformations.

Roadmaps of user sampled conformations are created by incremental construction, using the method presented in [32]. Building roadmaps incrementally also requires less computation and is online, enabling the roadmap constructor to recieve new user sets during runtime. After the first roadmap of 1000 samples is built, we create successive roadmaps by importing another set of user conformations and connecting.

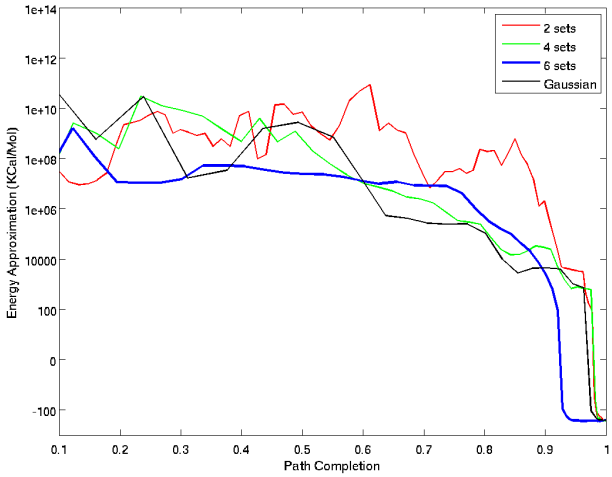

Fig. 5. Potential energy through the same query for constructed roadmaps. Path completion is the normalized amount of total RMSD traveled in query path. The energy is shown on a logarithmic scale.

Finally, queries were performed from a start ligand conformation about 5.02Å RMSD distance from the known native conformation of 3H9S using Djikstra's shortest weighted path method. As user sets were combined into larger roadmaps, the resulting query path became smoother with less peaks or energy barriers to overcome as seen in Figure 5. The path resulting from 6 user sets, shown in blue, also has the least potential energy as it approaches the native conformation.

The resulting query from the Gaussian roadmap, displayed in black, also contains less pronounced energy peaks. However, recall that the Gaussian ligand conformation samples were generated with a centered mean around the known native conformation while the users did not know the precise native conformation and had only the force feedback to guide them. Therefore, the Gaussian sampler could not be applied to new ligand receptor pairs of unknown native conformation where the program doesn't require a known native conformation.

V. CONCLUSION AND FUTURE WORK

To our knowledge, this is the first work that investigates molecular docking by combining haptics and crowdsourced multi-user path planning. Roadmaps built from a small number of haptic-guided user trajectories produced paths of low, smooth potential energy to the native conformation where an automated sampling method required full knowledge of the native conformation. We are currently undertaking a larger crowdsourcing study and expanding the Haptic-Guided Ligand Docking Tool to other input devices. By supporting various commodity input/output hardware devices, a larger user base could be leveraged. The iterative nature of roadmap construction can be integrated easily with streams of new user conformation sets and resulting roadmap quality can be monitored over time. Also, we are exploring kinematic linkage extensions that will allow us to represent molecular flexibility. Implementing molecular flexibility would enable Haptic-Guided Ligand Docking Tool to find conformations that a rigid body model couldn't.

REFERENCES

[1] N. M. Amato and G. Song. Using motion planning to study protein folding pathways. 9(2):149–168, 2002. Special issue of Int. Conf. Comput. Molecular Biology (RECOMB) 2001.

[2] O. B. Bayazit, G. Song, and N. M. Amato. Ligand binding with OBPRM and user input. In *IEEE Int. Conf. on Rob. and Auto.*, pages 954–959, 2001.

[3] O. Burchan Bayazit, Guang Song, and Nancy M. Amato. Providing haptic 'hints' to automatic motion planners. Technical report, In Proc. of the Phantom User's Group Workshop, Oct 1999.

[4] A.L. Beberg, D.L. Ensign, G. Jayachandran, S. Khaliq, and V.S. Pande. Folding@home: Lessons from eight years of volunteer distributed computing. In *Parallel Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on*, pages 1–8, May 2009.

[5] Petter Bivall, Shaaron Ainsworth, and Lena A.E. Tibell. Do haptic representations help complex molecular learning? *Wiley Periodicals, Inc. Sci. Ed.*, 95:700–719, March 2011.

[6] Aude Bolopion, Barthélemy Cagneau, Stephane Redon, and Stéphane Régnier. Haptic feedback for molecular simulation. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 237–242, Oct 2009.

[7] Aude Bolopion, Barthélemy Cagneau, Stephane Redon, and Stéphane Régnier. Comparing position and force control for interactive molecular simulators with haptic feedback. *Journal of Molecular Graphics and Modelling*, 29(2):280–289, 2010.

[8] OY Borbulevych, KH Piepenbrink, BE Gloor, DR Scott, RF Sommese, DK Cole, AK Sewell, and BM Baker. T cell receptor cross-reactivity directed by antigen-dependant tuning of peptide-MHC molecular flexibility. *Immunity*, 31:885–896, 2009.

[9] Brian Y. Chen and Barry Honig. VASP: A volumetric analysis of surface properties yields insights into protein-ligand binding specificity. *PLoS Comput Biol*, 6(8):e1000881, 8 2010.

[10] Holger ClauBen, Christian Buning, Matthias Rarey, and Thomas Lengauer. FlexE: Efficient molecular docking considering protein structure variations. *Journal of Molecular Biology*, 308(2):377–395, 2001.

[11] Seth Cooper, Firas Khatib, Adrien Treuille, Janos Barbero, Jeehyung Lee, Michael Beenen, Andrew Leaver-Fay, David Baker, and Zoran Popovic. Predicting protein structures with a multiplayer online game. 466:756–760, 2010.

[12] J Cortés, T Siméon, V Ruiz de Angulo, D Guieysse, M Remaud-Siméon, and V Tran. A path planning approach for computing large-amplitude motions of flexible molecules. *Bioinformatics (Oxford, England)*, 21 Suppl 1:i116–i125, 2005.

[13] Bruno Daunay, Alain Micaelli, and Stephane Regnier. Energy-field reconstruction for haptic-based molecular docking using energy minimization processes. pages 2704–2709, 2007.

[14] Yong Duan, Chun Wu, Shibasish Chowdhury, Mathew C. Lee, Guoming Xiong, Wei Zhang, Rong Yang, Piotr Cielplak, Ray Luo, Taisung Lee, James Caldwell, Junmei Wang, and Peter Kollman. A point-charge force field for molecular mechanics simulations of proteins based on condensed-phase quantum mechanical calculations. *Wiley Periodicals, Inc. J. Comput. Chem.*, 24:1999–2013, July 2003.

[15] Todd J. A. Ewing and Irwin D. Kuntz. Critical evaluation of search algorithms for automated molecular docking and database screening. *Journal of Computational Chemistry*, 18(9):1175–1189, 1997.

[16] Xuejian He and Yonghua Chen. Haptic-aided robot path planning based on virtual tele-operation. *Robotics and Computer-Integrated Manufacturing*, 25(4–5):792–803, 2009.

[17] X. Hou, Olga Sourina, and S. Klimenko. Visual haptic-based collaborative molecular docking. In *The 15th International Conference on Biomedical Engineering*, volume 43, pages 360–363. Springer International Publishing, 2014.

[18] Xiyuan Hou and O. Sourina. Haptic rendering algorithm for biomolecular docking with torque force. In *Cyberworlds (CW), 2010 International Conference on*, pages 25–31, Oct 2010.

[19] Xiyuan Hou and Olga Sourina. Six degree-of-freedom haptic rendering for biomolecular docking. In *Transactions on Computational Science XII*, volume 6670, pages 98–117. Springer Berlin Heidelberg, 2011.

[20] Sheng-You Huang and Xiaoqin Zou. Advances and challenges in protein-ligand docking. *International Journal of Molecular Sciences*, 11(8):3016–3034, 2010.

[21] Gareth Jones, Peter Willett, Robert C Glen, Andrew R Leach, and Robin Taylor. Development and validation of a genetic algorithm for flexible docking. *Journal of Molecular Biology*, 267(3):727–748, 1997.

[22] Firas Khatib, Seth Cooper, Michael D. Tyka, Kefan Xu, Ilya Makedon, Zoran Popovic, and David Baker. Algorithm discovery by protein folding game players. *Proceedings of the National Academy of Sciences*, 2011.

[23] Irwin D. Kuntz, Jeffrey M. Blaney, Stuart J. Oatley, Robert Langridge, and Thomas E. Ferrin. A geometric approach to macromolecule-ligand interactions. *Journal of Molecular Biology*, 161(2):269–288, 1982.

[24] M. Lawitzky, J.R. Medina, Dongheui Lee, and S. Hirche. Feedback motion planning and learning from demonstration in physical robotic assistance: differences and synergies. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 3646–3652, Oct 2012.

[25] Jens Meiler and David Baker. Rosettaligand: Protein-small molecule docking with full side-chain flexibility. *Proteins: Structure, Function, and Bioinformatics*, 65(3):538–548, 2006.

[26] Vijay S. Pande, Ian Baker, Jarrod Chapman, Sidney P. Elmer, Siraj Khaliq, Stefan M. Larson, Young Min Rhee, Michael R. Shirts, Christopher D. Snow, Eric J. Sorin, and Bojan Zagrovic. Atomistic protein folding simulations on the submillisecond time scale using worldwide distributed computing. *Biopolymers*, 68(1):91–109, 2003.

[27] EF Pettersen, TD Goddard, CC Huang, GS Couch, DM Greenblatt, EC Meng, and TE Ferrin. UCSF Chimera–a visualization system for exploratory research and analysis. *J. Comput. Chem.*, 25:1605–1612, October 2004.

[28] Matthias Rarey, Bernd Kramer, Thomas Lengauer, and Gerhard Klebe. A fast flexible docking method using an incremental construction algorithm. *Journal of Molecular Biology*, 261(3):470–489, 1996.

[29] G. Song and N. M. Amato. A motion planning approach to folding: From paper craft to protein folding. pages 948–953, 2001.

[30] Matthew B. Stocks, Steven Hayward, and Stephen D. Laycock. Interacting with the biomolecular solvent accessible surface via a haptic feedback device. *BMC Structural Biology*, 9, October 2009.

[31] C. Vázquez, J. Rosell, L. Chirinos, and O.A. Domínguez. Haptic primitives guidance based on the Kautham path planner. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 4686–4691, Oct 2010.

[32] Dawen Xie, Marco Morales, Roger Pearce, Shawna Thomas, Jyh-Ming Lien, and Nancy M. Amato. Incremental map generation (IMG). In *Algorithmic Foundation of Robotics VII*, pages 53–68. Springer, Berlin/Heidelberg, 2008. book contains the proceedings of the International Workshop on the Algorithmic Foundations of Robotics (WAFR), New York City, 2006.