

Haptic Virtual Fixtures for Robot-Assisted Manipulation

Jake J. Abbott, Panadda Marayong, and Allison M. Okamura*

Department of Mechanical Engineering, The Johns Hopkins University
{jake.abbott, pmarayong, aokamura}@jhu.edu

Abstract. Haptic virtual fixtures are software-generated force and position signals applied to human operators in order to improve the safety, accuracy, and speed of robot-assisted manipulation tasks. Virtual fixtures are effective and intuitive because they capitalize on both the accuracy of robotic systems and the intelligence of human operators. In this paper, we present the design, analysis, and implementation of two categories of virtual fixtures: *guidance* virtual fixtures, which assist the user in moving the manipulator along desired paths or surfaces in the workspace, and *forbidden-region* virtual fixtures, which prevent the manipulator from entering into forbidden regions of the workspace. Virtual fixtures are analyzed in the context of both cooperative manipulation and telemanipulation systems, considering issues related to stability, passivity, human modeling, and applications.

1 Introduction

Haptic virtual fixtures are software-generated force and position signals applied to human operators via robotic devices. Virtual fixtures help humans perform robot-assisted manipulation tasks by limiting movement into restricted regions and/or influencing movement along desired paths. By capitalizing on the accuracy of robotic systems, while maintaining a degree of operator control, human-machine systems with virtual fixtures can achieve safer and faster operation. To visualize the benefits of virtual fixtures, consider a common physical fixture: a ruler. A straight line drawn by a human with the help of a ruler is drawn faster and straighter than a line drawn freehand. Similarly, a robot can apply forces or positions to a human operator to help him or her draw a straight line. However, a robot (or haptic device) has the additional flexibility to provide assistance of varying type, level, and geometry.

Virtual fixtures show great promise for tasks that require better-than-human levels of accuracy and precision, but also require the intelligence provided by a human directly in the control loop. Traditional cooperative manipulation or telemanipulation systems make up for many of the limitations of autonomous robots (e.g., limitations in artificial intelligence, sensor-data interpretation, and environment modeling), but the performance of such systems is still fundamentally constrained by human capabilities. Virtual

* This work is supported by NSF grants #ITR-0205318 and #IIS-0347464.

fixtures, on the other hand, provide an excellent balance between autonomy and direct human control. Virtual fixtures can act as safety constraints by keeping the manipulator from entering into potentially dangerous regions of the workspace, or as macros that assist a human user in carrying out a structured task. Applications for virtual fixtures include robot-assisted surgery, difficult assembly tasks, and inspection and manipulation tasks in dangerous environments.

Virtual fixtures can be applied to two types of robotic manipulation systems: *cooperative manipulators* and *telemanipulators*. In cooperative manipulation, the human uses a robotic device to directly manipulate an environment. In telemanipulation, a human operator manipulates a master robotic device, and a slave robot manipulates an environment while following the commands of the master.

In general, the robots used in these systems can be of the impedance or the admittance type [3]. Robots of the impedance type, such as typical haptic devices, are backdrivable with low friction and inertia, and have force-source actuators. Robots of the admittance type, such as typical industrial robots, are nonbackdrivable and have velocity-source actuators. The velocity is controlled with a high-bandwidth low-level controller, and is assumed to be independent of applied external forces. Figure 1(a) shows the Johns Hopkins University Steady-Hand Robot [15], an admittance-type cooperative manipulator designed for microsurgical procedures. Figure 1(b) shows the da Vinci[®] Surgical System (Intuitive Surgical, Inc.), an impedance-type telemanipulator designed for minimally invasive surgical procedures. The virtual fixtures created and studied in our lab are designed explicitly for systems such as these.

A number of studies have been performed for virtual fixtures on cooperative manipulation and telemanipulation systems of both the impedance and admittance types [4, 10, 12, 13, 16]. It is also possible to implement virtual fixtures with passive robotic devices, such as “Cobots” [9, 11], that require

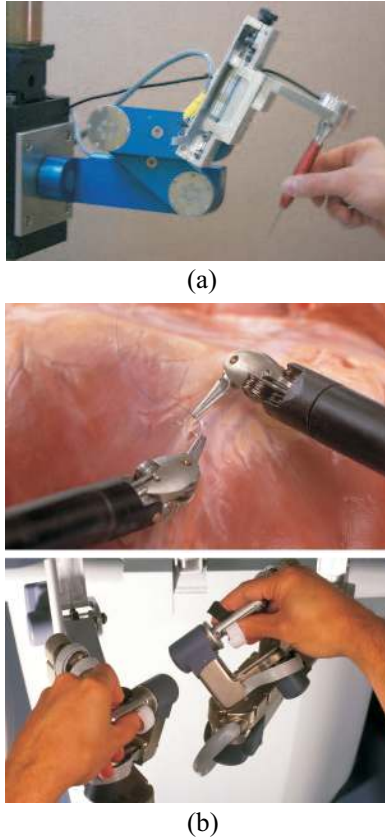


Fig. 1. (a) The Johns Hopkins University Steady-Hand Robot. (b) The da Vinci[®] Surgical System (image used with the permission of Intuitive Surgical, Inc.).




Fig. 2. (a) Guidance virtual fixtures assist in guiding the robot along desired paths. (b) Forbidden-region virtual fixtures help keep the robot out of forbidden regions.

specific hardware and therefore cannot be overlaid on existing robotic manipulation systems. While virtual fixtures have taken a variety of forms, all can be described by two categories: *guidance* virtual fixtures, which assist the user in moving the robot manipulator along desired paths or surfaces in the workspace, and *forbidden-region* virtual fixtures [12], which prevent the robot manipulator from entering into forbidden regions of the workspace. These virtual fixture types are illustrated in Fig. 2.

Prior work on virtual fixtures has been largely *ad hoc*, with significant reliance on particular applications. Thus, in this paper, we attempt to unify the past and present research in the field by considering the design, analysis, and application of virtual fixtures to various system types. In Sections 2 and 3, we discuss how guidance virtual fixtures and forbidden-region virtual fixtures, respectively, can be used for task assistance in both cooperative manipulation and telemanipulation. Next, in Section 4, we discuss in detail the issues involved with safe and functional implementation of virtual fixtures. Finally, in Section 5, we present a set of interesting topics for future work in this field of research.

2 Guidance Virtual Fixtures

Guidance virtual fixtures (GVFs) assist the user in moving the robot manipulator along desired paths or surfaces in the workspace. GVFs can be of either the impedance or admittance type [3]. Impedance-type GVFs act as potential fields, actively influencing the movement of the robotic manipulator. These impedance methods can lead to unexpected and undesirable movements of the manipulator, so we have chosen to focus on GVFs of the admittance type.

Admittance control typically takes the form $\mathbf{v} = K_a \mathbf{f}$, where \mathbf{f} is the user's applied force vector, K_a is an admittance gain matrix, and \mathbf{v} is the output velocity vector. This control scheme is sometimes referred to as proportional-velocity control. Admittance control has the desirable property that the velocity of the manipulator is proportional to the applied force, so the manipulator does not move if the user does not apply a force. In addition, slow robot movement is achieved with a soft touch. Admittance-type GVFs are very natural with admittance-type cooperative systems, but can also be implemented on impedance-type telemanipulation systems with a novel Pseudo-admittance control law [1].

2.1 GVFs for Cooperative Manipulation

In an admittance-type cooperative manipulation system, the robot motion is proportional to the user’s applied force, which is measured by a force sensor. To create GVFs, an instantaneous preferred direction is defined based on the position of the robot relative to the desired path or surface. The applied force is then decomposed into components in the preferred direction and in other, non-preferred directions. By eliminating the commanded motion due to the applied force in the non-preferred directions, we create a passive guidance along the preferred direction. Implementing GVFs in this fashion essentially makes the admittance gain matrix K_a both state and input dependent. Details of this GVF method can be found in [4].

Varying the response to the non-preferred force component creates different levels of guidance. *Hard* guidance refers to GVFs where none or almost none of the non-preferred force component is permitted, leaving the user with no or little freedom to deviate from the preferred path. Alternatively, *soft* GVFs give the user the freedom to move away from the path by allowing some motion in the non-preferred directions. We conducted an experiment with the JHU Steady-Hand Robot to evaluate the effect of GVF admittance on user performance, including accuracy and execution time [8]. Three tasks (Path Following, Off-path Targeting, and Avoidance) were selected to represent a broader class of motions that can occur in a real task execution. GVFs were used with varying admittance to keep the user on the preferred path, in this case a sine curve on a horizontal plane.

Figure 3 shows the robot trajectories during the Off-path Targeting and Avoidance tasks, with three levels of guidance. In the Targeting task, the users were instructed to reach the target located on the perimeter of the circle outlined in gray. In the Avoidance task, the users avoided the area by trying to follow along the circle perimeter. Robot trajectories in the Path




Fig. 3. Robot trajectories in the Targeting task (top) and the Avoidance task (bottom) with JHU Steady-Hand Robot.

Following task were similar to the portions seen outside the circular area in the two off-path tasks shown in Fig. 3. In the Path Following task, the users performed the task more accurately (with statistical significance) with GVFs, though not significantly faster. In the off-path tasks, the users had to fight against the GVF guidance to complete the desired motion. This represents situations where the virtual fixture is incorrectly placed and the user wishes to override the guidance. As expected, users take significantly longer to perform off-path tasks with increased guidance. Error also increases slightly. Despite the classic time-accuracy tradeoff, the experiment shows that GVFs can improve both time and accuracy simultaneously, while still allowing some independent user motion. More detailed descriptions of the experiment and the results can be found in [8]. GVF implementation for tasks in 3-D were also explored in Dewan *et al.* [5], where the tool was guided along a user-defined desired surface. In this experiment, stereo cameras were used to reconstruct the workspace and track the tool position and orientation.

2.2 GVFs for Telemanipulation

In telemanipulation, good position correspondence between the master and slave robots is desirable to create a sense of telepresence for the user. However, it is actually the slave manipulator that we wish to guide using GVFs, and master movements in its corresponding workspace are somewhat less important.

Unlike cooperative manipulation systems, telemanipulation systems are typically designed as impedance-type systems (that is, the master is an impedance-type haptic device, while the slave manipulator can be of either the impedance or admittance type). For these systems, we do not control the velocity of the system directly (due to force-source actuation), so we cannot implement admittance control directly. We have developed a novel telemanipulation control algorithm called Pseudo-admittance control [1] that mimics admittance control on impedance-type telemanipulators, and extends the GVFs described in Section 2.1 and [4] to telemanipulation. Pseudo-admittance makes use of a proxy, which exists only in software, that can be commanded to move under admittance control.

Under Pseudo-admittance control, the master servos to the slave position, while the slave servos to the proxy position, as illustrated in Fig. 4. The proxy moves under admittance control, using the force of the master's servo controller as its input force \mathbf{f} . GVFs are then implemented by attenuating the commanded velocity in non-preferred directions, as described in Section 2.1. Figure 4 shows the experimental results from two PHANTOM[®] robots (SensAble Technologies, Inc.) configured for Pseudo-admittance control. Using different levels of guidance (i.e., modifying the calculation of the preferred direction and the attenuation of velocities in the non-preferred directions), the slave is guided to a preferred plane in the workspace, but the user retains ultimate control to move the slave anywhere in the workspace.

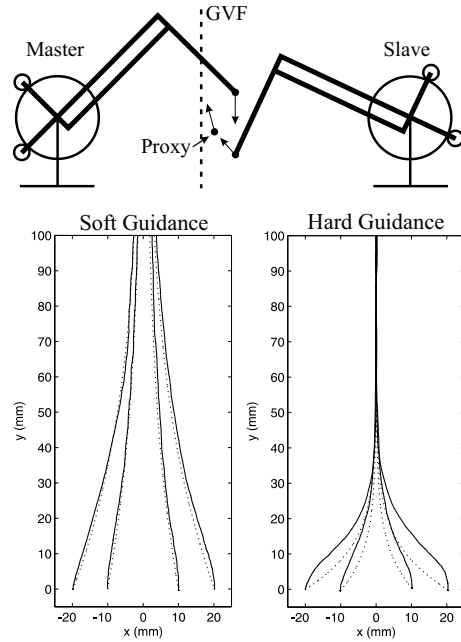


Fig. 4. Guidance virtual fixtures implemented on two PHANTOM[®] robots (top) configured for Pseudo-admittance Bilateral Telemanipulation [1]. Experimental data (bottom), with master (—) and proxy (···) trajectories, are shown for two levels of guidance. The slave serves to the proxy. The GVF is on the plane $x = 0$. The user applies a force approximately in the positive y direction, and the manipulator is guided by the GVF.

3 Forbidden-Region Virtual Fixtures

Forbidden-region virtual fixtures (FRVFs) prevent the robot manipulator from entering into forbidden regions of the workspace. They have an on/off nature, such that they have no effect on the robot when it is not interacting with the FRVF. As with GVFs, FRVFs can be of either the impedance or admittance type. Impedance-type FRVFs take the form of “virtual walls,” which are commonly employed and studied for haptic virtual environments, and are typically implemented as simple spring-damper surfaces. These are penalty-based methods, so the force generated by the FRVF is proportional to the manipulator’s penetration of the FRVF (i.e., some penetration is necessary to engage the FRVF). Admittance-type FRVFs are simply implemented by not commanding any manipulator motion into the forbidden region.

3.1 FRVFs for Cooperative Manipulation

FRVFs can be viewed as a subclass of GVF for an admittance-controlled cooperative manipulator. The FRVFs are trivial to implement, by simply eliminating any commanded motion into the forbidden region. Inherently, the forbidden region is the non-preferred direction defined in the GVFs.

Examples of FRVFs in cooperative systems are highlighted in [5] and [7]. In Dewan *et al.* [5], the virtual fixtures constrained the user to move along the shortest path between the current tool position and a predefined target

on the surface. The robot admittance gain was turned to zero once the target was reached. Li and Taylor [7] combined both GVFs and FRVFs in creating anatomy-based motion constraints for a path-following task in a constrained workspace. The algorithm uses the robot kinematics, the user's force input, and a 3-D geometric model of the workspace to generate virtual fixtures and an optimal set of joint displacements to guide the tool tip along a path while preventing the tool shaft from entering into forbidden regions.

The user may want the option to intentionally move past the FRVF if it is deemed necessary. The GVFs implemented in Section 2.1 left the user with ultimate control to move the manipulator away from the desired path, but it is not clear if it makes sense to create admittance-type FRVFs that allow some motion into the forbidden region. In one sense, an admittance-type FRVF that acts in this way is not a FRVF at all. It may be possible though, through state-and-input-dependent adaptation of the admittance-gain matrix, to implement FRVFs that allow some penetration into the forbidden region while retaining their functional purpose.

3.2 FRVFs for Telemanipulation

As with the GVFs of Section 2.2, in telemanipulation we are only really concerned with penetration of the slave manipulator into the forbidden region. Penetration of the master device into the corresponding region of its workspace is somewhat inconsequential.

Impedance-type FRVFs can be implemented on telemanipulators by overlaying a penalty-based virtual wall on the existing telemanipulation controller. It is possible to implement the virtual wall on either the master or the slave side (or both simultaneously). Both have the effect of reducing movement of the slave into the forbidden region. However, each presents a different haptic experience for the user, depending on the underlying telemanipulation controller, and each provides different levels of disturbance rejection (depending on the location of the disturbance) [1].

It is also possible to implement admittance-type FRVFs through the use of a proxy. If the slave manipulator serves to a proxy, rather than directly servoing to the master, then we can influence slave movement in forbidden regions by adapting the dynamic properties of the proxy. When the master is not interacting with the FRVF, the proxy is made to follow the master exactly. When the master moves beyond the FRVF, we attenuate the movement of the proxy past the FRVF (including removing the penetration completely).

Both types of FRVF act by attenuating slave movement into the forbidden region, while allowing the user to move the slave into the forbidden region if desired. The amount of attenuation, and consequently user control, is governed by system gains. Admittance-type FRVFs can be made to be infinitely stiff, but the stiffness of an impedance-type FRVF is ultimately limited by stability constraints. The stability of impedance-type FRVFs, under stability and passivity considerations, is explored in detail in [1].

4 Virtual Fixture Design Considerations

Prior work in virtual fixtures has focused primarily on application-specific virtual-fixture geometries and user performance of specific tasks. This section highlights a number of additional design considerations that are important for progress in this field; researchers have only recently begun to examine these issues.

One fundamental design problem is to determine the best type of underlying system for a virtual-fixture application. Cooperative manipulation systems are intuitive to use, due to the natural hand-eye coordination that comes from directly manipulating the tool. The sense of telepresence felt with a telemanipulator is limited by the position error in the system, as well as the quality of the visual and haptic feedback provided to the user. Admittance-type cooperative systems also have desirable “steady-hand” properties; the user’s hand is literally steadied by holding onto the rigid, slow-moving robot. This behavior must be mimicked on an impedance-type telemanipulator; the slave manipulator can be controlled to move slowly, but a backdrivable master device is not as capable of steadying the hand of the user. However, telemanipulators provide not only the ability to manipulate distant environments, but also the ability to provide scaling in both position and force. Force scaling is also possible with cooperative manipulation [14], although an additional force sensor or accurate environmental model is needed to obtain the contact force. It is important, in general, to consider whether force sensing is necessary and practical in terms of size, cost, and environment compatibility.

System performance also depends on the accuracy of the task geometry definition. For example, a computer vision system can be used to reconstruct the workspace and define the geometry of the virtual fixtures. The accuracy of the virtual fixtures defined depends on the resolution of the vision system, calibrations, and the accuracy of the tracking algorithm, which can be sensitive to changing light conditions and occlusions. The designer of a virtual fixture must be able to predict the sensitivity of system performance to inaccuracies in virtual-fixture geometry definition and develop mechanisms to correct for errors. It may be necessary to build in enough user control to compensate for inaccuracies in the virtual-fixture geometry, as was discussed in Section 2.1.

In cooperative systems, unmodeled robot dynamics, such as joint and link flexibility, can introduce significant tool positioning error, especially for micro-scale tasks. Joint and link flexibility add unactuated degrees of freedom to the robot. A human actively and directly manipulating the tool exacerbates the difficulty of error correction. A hand dynamic model could be added to better predict the system response near a virtual fixture, and adjust the controller appropriately to compensate for the error. This issue is being tested with 1-DOF system with link flexibility (Fig. 5(a)), which is equipped with position sensing from both an optical encoder on the linear stage and computer vision sensing of the position of the tool.

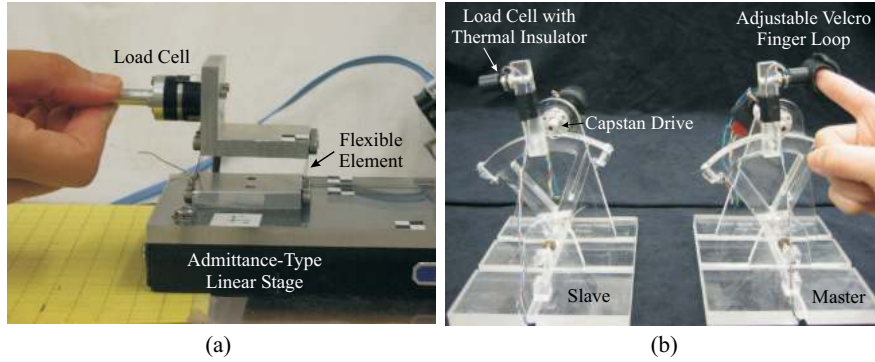


Fig. 5. 1-DOF experimental systems. (a) Admittance-type cooperative manipulator for the study of the effect of link flexibility on virtual-fixture performance. (b) Impedance-type telemanipulator for the study of FRVF stability.

Another major concern in the design of virtual fixtures for impedance-type telemanipulators is stability. Because of their backdrivable force-source actuators, these systems are prone to instability if the control-system gains are too high. This makes stable and effective virtual fixtures conflicting goals. We have investigated the stability of FRVFs, considering effects of friction, sampling, and quantization, using both equilibrium stability analysis [1] and passivity analysis [2]. We used a 1-DOF system, shown in Fig. 5(b), for this purpose. It is possible to design a FRVF to be passive, with the additional assumption of human passivity being sufficient for system stability. However, as shown in Fig. 6, we found that including an explicit model of potential human users can lead to stability predictions that are significantly less conservative than simply requiring passivity of the FRVF.

It is tempting to model the human user as an exogenous input to the system, for the purpose of stability analysis, but in general, the dynamics of the human user are part of the closed-loop feedback system. However, it is also reasonable to assume that for certain slow-moving systems, the human user is essentially unaffected by the movement of the system. An initial study in our lab shows that, for an admittance-type cooperative manipulator, it is the velocity of the robot, and not the admittance gain, that directly affects human force control precision [17]. Thus, by restricting the velocity of the manipulator, it may be possible to consider the human user as an exogenous input, greatly




Fig. 6. Experimental stability limits on master FRVF stiffness, for large- and small-handed users, compared to predicted stability limits and passivity of the FRVF.

simplifying system stability analysis. More research is needed to better understand the role of the human user in the total system response.

As illustrated above, it is not always obvious when dynamic modeling of the human user is necessary or desirable in virtual-fixture design and analysis. Most of the prior work on virtual fixtures has excluded modeling of the human user. In addition to mechanical modeling, experimental results of GVFs in cooperative systems suggest that human intent and psychophysics may also affect GVF performance. Selecting an appropriate level of guidance is required for optimal performance, and the selection is task dependent. Having a high level of guidance increases error and time for tasks that require off-path motions, though it significantly improves both time and error during path-following. An optimal GVF selection was explored in [8]. Artificial intelligence can also be added to adjust the GVF based on the user’s intent. For example, Li and Okamura [6] used Hidden Markov Models to recognize user motions and provide appropriate GVF assistance in a combined curve-following and object-avoidance task.

5 Summary and Future Work

This paper described methods for design and implementation of haptic virtual fixtures on a number of different underlying platforms. Through analysis and experiments, we show that virtual fixtures can improve human-machine performance, while allowing the user to maintain ultimate control over the task execution.

There are a number of critical questions that provide important topics for future research in this field. For example, what is the best virtual-fixture geometry for a given task? How does the human user interpret the combination of haptic cues coming from the manipulated environment and the virtual fixture? Does this lead to haptic confusion, affecting the user’s sense of immersion in the task? If the virtual fixture geometry and/or gains vary in time, not only could it lead to confusion on the part of the user, but it also complicates stability analysis. Can virtual fixtures be used as training devices for complicated tasks, and then eventually be removed, much like training wheels on a bicycle? To what extent does the human need to be included in the analysis of these systems? It is desirable to say as much as possible about the robotic system itself, without needing to consider human dynamics. Is it possible to apply what we have learned thus far to the design of force virtual fixtures, which assist the user in applying the proper force to the manipulated environment?

It is important that we generalize the research in this field across systems and tasks, so that knowledge gained in individual research efforts can advance the field as a whole. Virtual fixtures will no doubt facilitate robot-assisted tasks that were previously impossible, but this nascent field is rich with interesting research topics that must be explored before human-machine systems can capitalize on the full benefit of virtual fixtures.

References

1. J. J. Abbott. *Virtual Fixtures for Bilateral Telem Manipulation*. PhD thesis, Department of Mechanical Engineering, The Johns Hopkins University, 2005.
2. J. J. Abbott and A. M. Okamura. Effects of position quantization and sampling rate on virtual-wall passivity. *IEEE Trans. Robotics*, 2005.
3. R. J. Adams and B. Hannaford. Stable haptic interaction with virtual environments. *IEEE Trans. Robotics and Automation*, 15(3):465–474, 1999.
4. A. Bettini *et al.* Vision-assisted control for manipulation using virtual fixtures. *IEEE Trans. Robotics*, 20(6):953–966, 2004.
5. M. Dewan *et al.* Vision-based assistance for ophthalmic micro-surgery. In *Proc. 7th Int'l. Conf. on Medical Image Computing and Computer-Assisted Intervention*, pages 49–57, 2004.
6. M. Li and A. M. Okamura. Recognition of operator motions for real-time assistance using virtual fixtures. In *Proc. 11th Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*, pages 125–131, 2003.
7. M. Li and R. H. Taylor. Spatial motion constraints in medical robot using virtual fixtures generated by anatomy. In *Proc. IEEE Int'l. Conf. on Robotics and Automation*, pages 1270–1275, 2004.
8. P. Marayong and A. M. Okamura. Speed-accuracy characteristics of human-machine cooperative manipulation using virtual fixtures with variable admittance. *Human Factors*, 46(3):518–532, 2004.
9. C. A. Moore *et al.* Cobot implementation of virtual paths and 3-D virtual surfaces. *IEEE Trans. Robotics and Automation*, 19(2):347–351, 2003.
10. S. Park *et al.* Virtual fixtures for robotic cardiac surgery. In *Proc. 4th Int'l. Conf. on Medical Image Computing and Computer-Assisted Intervention*, pages 1419–1420, 2001.
11. Y. S. Park *et al.* Enhanced teleoperation for D & D. In *Proc. IEEE Int'l. Conf. on Robotics and Automation*, pages 3702–3707, 2004.
12. S. Payandeh and Z. Stanisic. On application of virtual fixtures as an aid for telemanipulation and training. In *Proc. 10th Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*, pages 18–23, 2002.
13. L. Rosenberg. Virtual fixtures: Perceptual tools for telerobotic manipulation. In *Proc. IEEE Virtual Reality Int'l. Symposium*, pages 76–82, 1993.
14. J. Roy *et al.* Haptic feedback augmentation through position based adaptive force scaling: Theory and experiment. In *Proc. IEEE/RSJ Int'l. Conf. on Intelligent Robots and Systems*, pages 2911–2919, 2002.
15. R. Taylor *et al.* Steady-hand robotic system for microsurgical augmentation. *Int'l. J. Robotics Research*, 18(12):1201–1210, 1999.
16. N. Turro and O. Khatib. Haptically augmented teleoperation. In *Proc. 7th Int'l. Symposium on Experimental Robotics*, pages 1–10, 2000.
17. M. Wu *et al.* Effect of velocity on human force control. In *Proc. World Haptics Conf.*, pages 73–79, 2005.