Position-Based Virtual Fixtures for Visual Control of a Handheld Micromanipulator

Brian C. Becker, *Student Member, IEEE*, Robert A. MacLachlan, *Member, IEEE*, Louis A. Lobes, Jr.,
Gregory D. Hager, *Fellow, IEEE*, Cameron N. Riviere, *Senior Member, IEEE*

Abstract—Performing micromanipulation and delicate operations in sub-mm workspaces is difficult because of destabilizing tremor and imprecise targeting. Robotic aid combined with filtering techniques that suppress tremor frequency bands increases performance. However, if knowledge of the operator's goals is available, virtual fixtures have been shown to improve performance even more. In this paper, we derive a position-based virtual fixtures framework for active handheld micromanipulators. For applicability in surgical environments, the fixtures are generated in real-time from microscope video during the procedure. Additionally, we develop motion scaling behavior around virtual fixtures as a simple and direct extension to our formulation. We demonstrate that virtual fixtures outperform tremor cancellation algorithms on a set of artificial but medically relevant tasks: hold still, circle tracing, move-and-hold, and vein tracing ($p < 0.05$).

Index Terms—Micro/nano robots, dexterous manipulation, motion control, medical robots and systems, vision-based control

I. INTRODUCTION

IN biology and microsurgery, proper manipulation of extremely small anatomical features often requires precision and dexterity that exceeds the capability of an unaided operator. Physiological tremor, or high-frequency involuntary hand movements with amplitudes of over 100 μm [1], is a large contributing factor to the difficulty of micromanipulation. Furthermore, lower frequency drift in gross hand positioning over time, caused by poor depth perception through the microscope and the relatively low bandwidth of the eye-hand feedback loop, reduces pointing accuracy [2]. Various robotic platforms such as the Robot Assisted MicroSurgery (RAMS) system [3], the Johns Hopkins SteadyHand [4], and the Eye-Robot [5] have been proposed to deal with these problems in micromanipulation. In general, robotic control aids for micromanipulation can be grouped into three categories: tremor compensation, motion scaling, and virtual fixtures.

Tremor compensation is a key component of many surgical robots [6, 7], the goal usually being to suppress frequency bands dominated by tremor in order to eliminate unwanted motion while preserving the operator's intended movements. Cooperative robots such as SteadyHand accomplish this mechanically by utilizing the inherent stiffness of the robot arm to damp high-frequency movement. In robots where the actuation is independent of the hand movement, tremor compensation can be achieved by inserting filters between the hand motion and the drive mechanism. In either case, reduction in the higher-frequency tremor while retaining the operator's gross lower-frequency movements is the end goal.

Motion scaling, often used in conjunction with tremor suppression, scales movements so the tip only moves a fraction of the hand movement. For example, a scale factor of ½ would transform hand movement of 2 mm to a 1-mm movement at the robotic tip. This scaling behavior is a key feature in master/slave robots [8, 9], but is more difficult in cooperative robots, which are controlled by force at the hand/robot interface, rather than hand motion [4]. Recent developments in tremor compensation for handheld devices describe a low-pass shelving filter [2], which is a hybrid tremor suppression filter that affords what can be viewed as relative motion scaling. An additional 30-50% reduction in tremor compared to a simple low-pass filter is reported.

Virtual fixtures, in contrast to tremor compensation or motion scaling behaviors that operate in the general case, instead aim to improve specific motions or tasks [10]. Similar to a ruler that aids in the specific instance of drawing straight lines, virtual fixtures can be thought of as software guides that constrain the robot motion in specific ways. Proposed originally by Rosenberg [11] as a method to overlay abstract
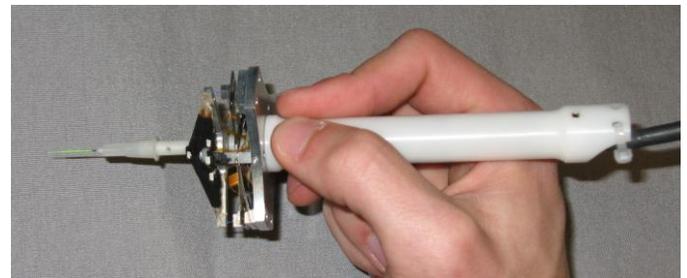
Fig. 1. Micron micromanipulator, shown without housing to illustrate the piezoelectric motors between the handle and tip, which enable the tip to actuate independently of the handle (or hand) motion.

sensory information onto a force-reflecting master workspace, virtual fixtures were developed to address latency in teleoperation tasks. They were also found to reduce cognitive load and increase performance. Successful efforts have been reported in using virtual fixtures with the SteadyHand robot for surgically-relevant tasks [12, 13].

A third class of micromanipulators is active handheld micromanipulators, best exemplified by the Micron device [2] depicted in Fig. 1. Micron incorporates motors between the handle and the tip, effectively creating an active handheld tool in which the end-effector can actuate independently of the hand motion within a limited range of motion (1x1x0.5 mm). With this capability, active tremor suppression [2], relative motion scaling, and simple virtual fixtures [14] have been individually investigated. Recently we proposed a position-based virtual fixture framework for handheld micromanipulators that integrates both tremor compensation and motion scaling [15]. This paper expands upon that work by formulating a more general spline representation for virtual fixtures, integrating vision-based control with dense stereo vision, using a new feedforward controller [16], introducing extensions from planar surfaces to naturally curved surfaces, adding a new vein-tracing experiment, and reporting results from multiple surgeons. Section II covers background material and introduces Micron and the system setup. Section III presents the derivation of position-based virtual fixtures for Micron. In Section IV, we test our system with medically relevant experiments. Finally, we conclude in Section V with a discussion of results and directions for future work.

II. BACKGROUND

Virtual fixtures and the concept of robotics constraints have gained popularity with the practical application of guiding the tip of robotic end-effectors during surgical manipulation tasks. Funda et al. [17] presented an optimization technique for commanding multi-linkage robotic arms via joystick, while keeping joints within their limits, maintaining orientation of the end-effector, and imposing task-space constraints on position. Kapoor et al. [18] extended Funda's work in constrained optimization to jointly handle simultaneous primitives such as maintaining a direction, rotating around a line, and staying above a plane, with a weighted optimization. Davies et al. [19] demonstrated the ACROBOT (Active Constraint Robot) with virtual fixtures to prevent the surgical cutter from entering pre-defined avoidance zones. Everett et al. [9] showed decreased operation time and operator fatigue in telemanipulating objects viewed from a camera by dynamically scaling the motion based on the relative position and orientation of the robotic end-effector.

In search of more general frameworks, Kumar et al. [20] decomposes surgical tasks into basic system states where the surgeon initiates primitive actions with a foot pedal to execute subroutines to position, orient, approach, insert, and the like. Bettini et al. [10] introduce a comprehensive virtual fixtures framework that translates forces applied to a shared manipulator arm into end-effector velocities after applying motion constraints, demonstrating a number of useful
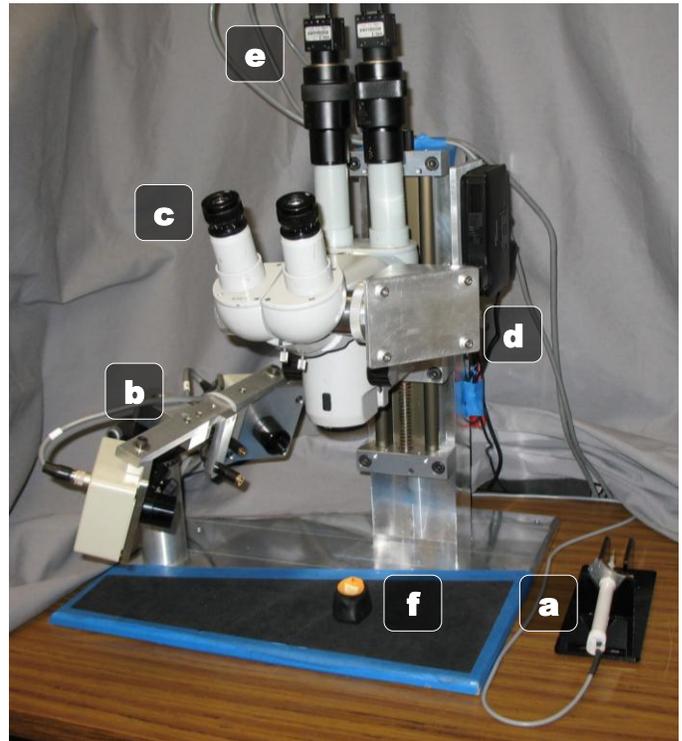
Fig. 2. System setup with (a) Micron, (b) ASAP position sensors, (c) surgical microscope, (d) image injection system, (e) stereo cameras, and (f) phantom half eyeball.

behaviors such as trajectory following, positioning, and volumetric restrictions. Hard virtual fixtures disregard any forces in non-preferred directions, in order to prevent deviations from the fixtures, while soft virtual fixtures allow a mixture between human commands and the active virtual fixture. Such virtual fixtures and robotic pose constraints can be generated from cameras observing the surgical workspace, leading to research in using virtual fixtures to enable the end-effector to track surfaces [13, 21], follow trajectories in tight spaces [22], and aid in simulated surgical procedures [12, 22].

A significant amount of research has been accomplished in applying virtual fixtures to various robotic platforms and micromanipulation problems [10, 18, 19]. In most formulations of virtual fixtures, the user manipulates a robot arm that is attached to the instrument directly or remotely via teleoperation. Forces on the robot arm are transformed to velocity commands at the instrument tip. These velocities are shaped first in software via constraints as specified by the active virtual fixture. If the robot is non-backdriveable (such as the Johns Hopkins SteadyHand), strict adherence to the virtual fixture can be enforced by ensuring that velocity components are zero in directions that move the tip away from the fixture.

Unlike most existing virtual-fixture-enabled robots, Micron is not manipulated by the operator through the application of forces to a joystick control or robot arm. It is a fully handheld device that purely senses position; thus, the input to the virtual fixtures must be handle (or hand) motion. This fundamental difference, the use of position instead of force as the control input, necessitates the development of a unique formulation of

virtual fixtures specifically designed for this class of handheld micromanipulators.

A. System Setup

The robotic system used in this research is Micron, a previously reported 3 DOF micromanipulator [2]. Unlike a shared robotic arm or a master/slave system, Micron is fully handheld and has motors positioned between the handle and the instrument tip as depicted in Fig. 1. The motors are three Thunder® piezoelectric actuators arranged in a radially symmetric pattern, allowing for tip movement independent of hand motion within a region of approximately 1x1x0.5-mm, centered on the handle.

Low-latency, high-bandwidth positioning information is obtained from custom optical tracking hardware named ASAP [23]. Three LEDs on the actuated shaft of the instrument allow for 6 DOF tracking of Micron's tip. An additional LED attached to the handle provides information about raw or uncompensated hand movement. Position sensitive detectors (PSDs) detect the LEDs within a 4 cm^3 workspace at a rate of 2 kHz, with measurement accuracies of less than 10 μm RMS.

The operator uses Micron under a surgical microscope (Zeiss OPMI 1). A 27 gauge hypodermic needle (418 μm outer diameter) is attached to Micron. A 127 μm nitinol wire is inserted into the needle to serve as a very fine instrument tip. The tip is tapered to a blunt point and painted for enhanced contrast. Stereo cameras (PointGrey Flea2) attached to the microscope view the same workspace as the operator and capture 1024x768 resolution video at 30 Hz. The cameras track both the tip of the instrument and anatomical targets. The entire setup can be seen in Fig. 2.

B. Camera Calibration and Adaptive Registration

Vision-based control of the manipulator is necessary because the anatomy of interest is localized via intraoperative imagery. However, traditional image-based visual servoing techniques [24, 25] operating at the camera rates of 30 Hz cannot provide the high bandwidth and low latencies needed for active control of handheld micromanipulators. Instead, the control loop must use the much faster 2 kHz ASAP positioning system. Thus, virtual fixtures generated from vision must be transferred from the stereo images to the ASAP coordinate frame, requiring calibration and registration of the cameras.

Calibration similar to [26] is performed by collecting roughly 30 s of the 2D tip positions $p \in [x, y, 1]^T$ in the image with the corresponding 3D tip positions $P \in [X, Y, Z, 1]^T$ as measured by ASAP. The projective camera matrix $M \in \mathbb{R}^{3 \times 4}$ that satisfies the homogenous relationship $p \equiv MP$ is estimated using the Direct Linear Transform [27]. After the initial 30-s calibration procedure, an adaptive registration dynamically adjusts the calibration with each new camera frame to account for any nonlinearities and microscope focusing.

$$M \leftarrow M + \eta(p - MP)P^T \tag{1}$$

where η is a small gain (e.g. $\eta = 10^{-6}$). Given the calibration

and a corresponding 2D point in the left and right images, the 3D point in the ASAP frame can be estimated using least squares solution of the projection error [27]. Thus, an up-to-date, invertible mapping always exists between the cameras and ASAP coordinate frames. This allows anatomy localized and tracked in the cameras to be reconstructed in the 3D ASAP coordinate system, so that virtual fixtures can be generated and enforced at 2 kHz.

C. Feedforward Control

With high-rate, low-bandwidth 3D ASAP sensors, Micron performs position-based closed-loop control using an Internal Model Control (IMC) approach [2]. When a 3D goal (setpoint) is specified in handheld operation, tremor is 90% rejected. However, the remaining 10% of tremor represents 10-20 μm deviations of the tip position from the goal, making it the largest contributing factor in the error reported for hard virtual fixtures [15]. Since there is need to cannulate extremely small retinal vessels [28, 29], this error is significant.

To further reduce this error, a feedforward controller is designed to anticipate and correct deviations from the setpoint before they are seen by the feedback loop. A constant-acceleration Kalman filter is run at 2 kHz to optimally estimate the position, velocity, and acceleration states of the tip and handle. The velocity of the handle, which represents tremor velocity, can be used to estimate feedback error caused by the 3-ms delay in the position controller. To counteract this error, a feedforward controller that integrates velocity predictions from the Kalman filter was incorporated into the existing IMC feedback position controller. This dual feedforward+feedback strategy yields an additional 50% reduction in error, for a total tremor rejection rate of 95%. Details on the design and implementation are available in [16].

III. POSITION-BASED VIRTUAL FIXTURES

The tip of Micron in 3D space is defined as $P_T \in \mathbb{R}^3$. As Micron actuates, the tip will move independently of the hand (or handle) motion, so it is important to define a measure of how much the tip has moved. We define the null position as the 3D tip position $P_N \in \mathbb{R}^3$ under the assumption that Micron is off; i.e., P_N exactly reflects the hand motion. One can think of P_N as being mechanically tied to the handle; thus $P_T = P_N$ unless Micron actuates the tip. The remainder of this section is dedicated to devising virtual fixtures that modify the behavior of the tip P_T while using the handle motion P_N as the indicator
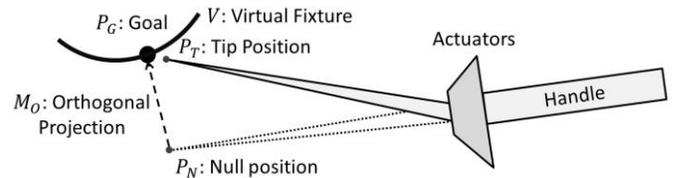
Fig. 3. Example of handheld micromanipulation with position-based virtual fixtures, which drives the tip position P_T to a goal position P_G on the virtual fixture V. The goal position is calculated by the orthogonal projection M_O of the null position P_N. The null position is the location of the tip position if the actuators were turned off. It is a natural selection for choosing the goal position because it is nominally the center of the actuator range of motion and exactly reflects the hand motion. Note: not to scale.

of operator intention. Fig. 3 graphically depicts the difference between the null position P_N and the actual tip position P_T.

A. Problem Definition

Represented as a subspace defined in 3D Euclidean space by the stereo vision cameras, the virtual fixture must constrain the tip position of Micron, using only the hand motion of the operator to guide and position the instrument tip on the virtual fixture. In the case of hard virtual fixtures, the tip of Micron should always lie on the subspace representing the fixture; in the case of soft virtual fixtures, the error between the hand motion and the virtual fixture should be scaled. In both cases, tremor compensation should smooth the tip movement.

B. Point Virtual Fixture

We begin the formulation of virtual fixtures for Micron by considering the simplest fixture possible: fixing Micron's tip to a single goal point $P_G \in \mathbb{R}^3$ in space. While the point virtual fixture is active, the control system should enforce the relationship $P_T = P_G$, regardless of where or how the operator moves the handle (within Micron's range of motion). We define the very simple control law:

$$P_T = P_G \qquad (2)$$

We use approximated linear inverse kinematics with an IMC feedback+velocity feedforward controller and a notch filter at the manipulator resonance to drive the tip to the goal position. Controller latency is 3 ms and settling time is at most 50 ms if the goal point is within Micron's current range of motion. Once the point virtual fixture is activated, Micron uses a 1 kHz controller to drive and maintain the tip on the goal. In response to shifting anatomy, moving the virtual fixture is possible by setting a new goal point P_G. To avoid high-frequency oscillations, large changes in P_G should be smoothed with either filtration (lowpass, Kalman, etc.) or trajectory planning.

C. Higher-Order Subspace Virtual Fixtures

Higher-order subspaces can be built on top of the point virtual fixture to obtain more interesting behaviors. Each additional level adds a degree of freedom to the tip motion. For instance, the point virtual fixture restricts all motion. However, a line virtual fixture allows the tip to freely travel
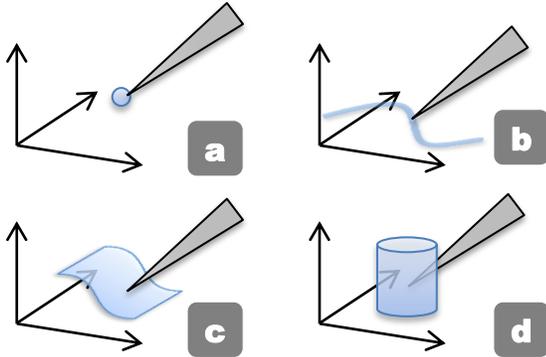
Fig. 4. Virtual fixtures constraining the tip to a subspace with increasing degrees of freedom: (a) point, (b) curve, (c) surface, and (d) volume.

along a line while restricting motion orthogonal to the line. Likewise, we can define the hierarchy of virtual fixtures as seen Fig. 4: point, curve, surface, and volume. To reconnect the need for these virtual fixtures back to medical applicability, it is worth considering possible medical relevance for each virtual fixture:

- **Point** (0 DOF): Steadying cannula during injection
- **Curve** (1 DOF): Following path for laser ablation, guiding suture/needle along a blood vessel
- **Surface** (2 DOF): Maintaining a constant standoff distance, navigating in narrow crevices
- **Volume** (3 DOF): Restricting the tip volumetrically to prevent tissue contact outside of "safe" areas

Although not a topic covered in this paper, such virtual fixture formulations extend to higher degrees of freedom to include orientation virtual fixtures. Each virtual fixture in the hierarchy can degenerate to a special case of a lower DOF fixture. For instance, a line virtual fixture can be implemented with a point virtual fixture that moves along the line.

In fact, we assert that all virtual fixtures can be implemented easily with just the point virtual fixture. The key to implementing higher-order fixtures with point fixtures is selecting the correct point on the higher-order virtual fixture. Algebraically, any point on a fixture will satisfy the constraint imposed by the fixture. Geometrically, it is most intuitive to select a point on the virtual fixture as close as possible to the operator's intended position. Thus for an arbitrary virtual fixture V, the nearest goal position P_G is the orthogonal projection of P_N onto the virtual fixture V. We define this orthogonal projection by the mapping $\boldsymbol{M_O}: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that selects the goal point P_G as the closest point from the null position P_N to the virtual fixture V:

$$P_G = \boldsymbol{M_O}(V, P_N). \qquad (3)$$

For simple geometric structures (lines, planes, circles, cylinders, etc.), analytic solutions for $\boldsymbol{M_O}$ exist. For more complex shapes, numerical solutions or iterative approximations to $\boldsymbol{M_O}$ are available [30].

D. Generalizable Spline Virtual Fixtures

When applying the position-based virtual fixtures framework to surgical environments, simple geometric shapes will not suffice for representation. A more general fixture representation is necessary, due to the complex nature of anatomical structures. In many fields, such as graphics, meshes are used to approximate arbitrarily complex objects. While such an approach has merits and in fact explicitly encodes surface normal information into the mesh, very fine meshes are needed to represent anatomy which is non-planar. Furthermore, the surface normals necessary for the orthogonal projection $\boldsymbol{M_O}$ are discontinuous at the vertices of the mesh, adding an extra level of complexity to the mapping. Because of the unwieldy nature of meshes and their large storage requirement for fine approximations to curved anatomy, meshes were not considered for the virtual fixture framework.

A more suitable representation is one of piecewise curves, or splines, which have the beneficial property of representing continuous curves and surfaces with low dimensionality. Additionally, the orthogonal mapping M_O can be calculated numerically [31]. Fitting splines can be done directly in image space [32], or after a point-cloud representation of the anatomy of interest has been extracted. Because virtual fixtures must be available in the 3D ASAP coordinate system, the latter approach is taken. Although a variety of spline representations may be used, we use basis-splines (B-splines) to describe general anatomy because of their versatile yet simple representation, with widespread support for fitting and evaluation. Surfaces such as the retina can be represented by a B-spline surface. Encoding a notion of "side" to the surface allows for volumetric fixtures, such as forbidding the instrument to pass through the B-spline surface virtual fixture representing the retina.

E. Tremor Suppression

In point virtual fixtures, all degrees of freedom are proscribed by the fixture, eliminating unwanted tremor. However, for higher-order virtual fixtures, tremor parallel to the subspace is not affected by the orthogonal projection M_O. In fact, tremor is evident in all 3 DOF for volume virtual fixtures when the tip is fully inside the volume. To reduce tremor in these higher DOF fixtures, one may use a tremor suppression filter $F_T^n : \mathbb{R}^n \to \mathbb{R}^n$ $n \in \{1,2,3\}$. For instance, we can insert the tremor suppression filter between the 3D null position P_N and the orthogonal projection M_O:

$$P_G = M_O(V, F_T^3(P_N)) \tag{4}$$

Alternatively, the tip position can be re-parameterized on the subspace representing the virtual fixture with a lower degree of freedom. In this case, we redefine the mapping $M_O^n : \mathbb{R}^3 \to \mathbb{R}^n, n \in \{1,2\}$ as transforming P_N in 3D space to a nD re-parameterized goal position on the virtual fixture via orthogonal projection. For instance, P_G on a curve can be re-parameterized by arc length and filtered with a one-dimensional tremor suppression filter F_T^1. For an n DOF virtual fixture, we can calculate the goal point as:

$$P_G = F_T^n(M_O^n(V, P_N)) \tag{5}$$

For experiments in this paper, we use (5) for circle tracing, to prevent any nonlinearity in the orthogonal projection from affecting the tremor suppression, and (4) for all other fixtures. It is also important to note that noise in virtual fixture placement can inject high-frequency movement similar to tremor. Therefore, if virtual fixtures are being generated by video, high-frequency components should be eliminated. We accomplish this with simple low-pass filtration, since the anatomy does not move rapidly.

F. Motion Scaling

So far we have described hard fixtures where the tip position cannot deviate from the constraint imposed by the virtual fixture. However, it is also possible to derive soft virtual fixtures that share control between the virtual fixture
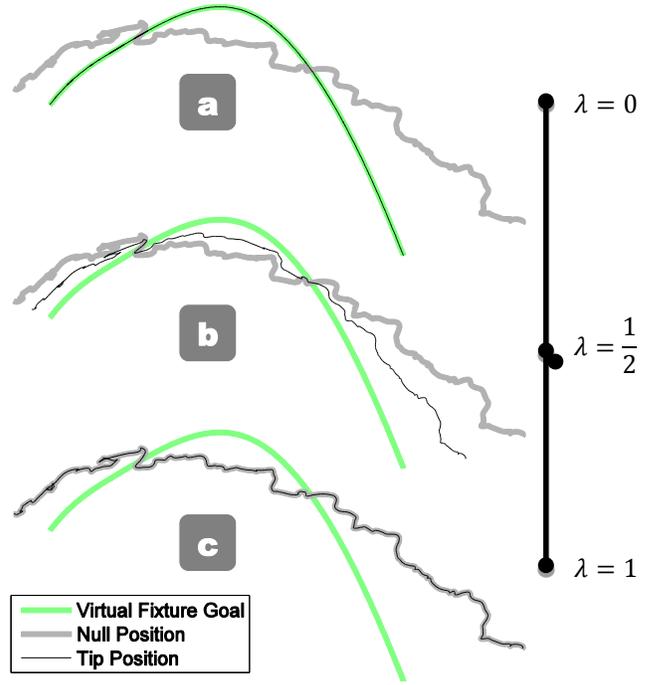
Fig. 5. Effect of motion scaling parameter λ on the tip position. (a) Hard virtual fixture ($\lambda = 0$). (b) Soft virtual fixture with $\lambda = 1/2$; control is shared 50% between virtual fixture and hand motion. (c) Fixtures off ($\lambda = 1$).

and the operator. If hard virtual fixtures are analogous to a stiff metal ruler, then soft virtual fixtures can be likened to a compliant foam ruler that aids in drawing lines but can be partially overridden by the operator. An additional parameter $\lambda \in [0,1]$ defines how much the operator can override the virtual fixture. In our formulation, λ represents the proportion of the hand motion P_N vs. the goal point P_G that Micron uses to actuate the tip:

$$P_T = (1 - \lambda)P_G + \lambda P_N \tag{6}$$

In essence, λ functions as a weighted average of the goal and null position. $\lambda = 0$ corresponds to a hard virtual fixture where the operator can only control where on the virtual fixture the tip goes, and $\lambda = 1$ disables virtual fixtures entirely and leaves the operator in complete control (Fig. 5).

However, for values of λ between 0 and 1, (6) can be directly manipulated into a motion scaling paradigm:

$$P_T = P_G + \lambda(P_N - P_G) \tag{7}$$

$$e = P_N - P_G \tag{8}$$

$$P_T = P_G + \lambda e \tag{9}$$

By introducing the error e as the difference between the null position P_N controlled by the operator and the goal point P_G calculated by the virtual fixture, it is clear the parameter λ directly represents the motion scaling factor. For example, if $\lambda = 1/2$, then all hand motions that deviate from the virtual fixture will be scaled by one half. Thus, we can see that sharing control between the virtual fixture and the operator yields motion-scaling behavior (see Fig. 5b for an example).

G. Generalized Control Law

To summarize the control law that incorporates virtual fixtures, tremor suppression, and motion scaling into a handheld micromanipulator framework based on position control, we assume for the sake of simplicity that tremor suppression on the null position is sufficient and re-parameterization is not necessary. Assuming all operations are in 3D, the generalized control law then becomes:

$$P_G = M_O(V, F_T(P_N)) \tag{10}$$

$$e = F_T(P_N) - P_G \tag{11}$$

$$P_T = P_G + \lambda e \tag{12}$$

First, (10) selects the goal point on the virtual fixture closest to the tremor suppressed null position. Eq. (11) then calculates the error between what the virtual fixture and where the operator is currently pointing. Finally, (12) drives the tip to either the virtual fixture or, if λ is non-zero, scales the error to achieve motion scaling about the fixture.

H. Combining Virtual Fixtures

Virtual fixtures can be easily combined into compound virtual fixtures by dynamically selecting which fixture is active. For example, in Section IV, we demonstrate a box virtual fixture built from four plane segments. One difficulty encountered when combining virtual fixtures occurs when the mapping M_O becomes degenerate. For instance, when two lines or planes form a right angle, the null position P_N might be positioned equidistant between both virtual fixtures. Noise in sensor measurement might subsequently introduce tip position oscillations as multiple virtual fixtures compete for control of the tip. One solution to prevent the tip from oscillating between virtual fixtures is to disable all fixtures except the current one until the tip nears the edge, so that only when the tip reaches the boundary between fixtures can Micron transition from one virtual fixture to another. This approach was used for our multi-fixture experiments.
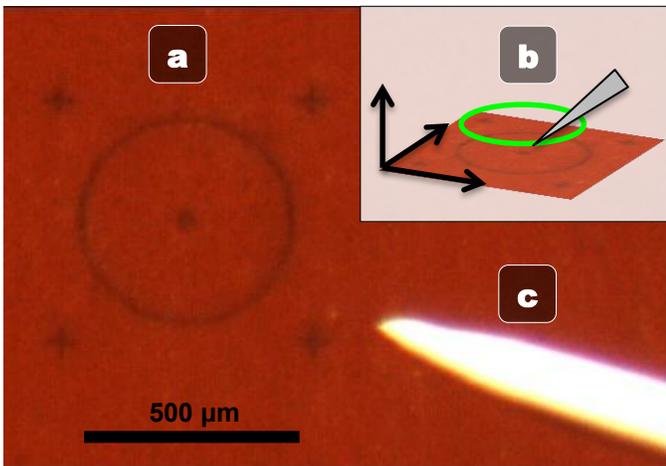
Fig. 7. (a) Laser etched target onto rubber surface (b) Generating a 3D circle virtual fixture from the tracked target (c) White-painted, tapered tip of Micron.
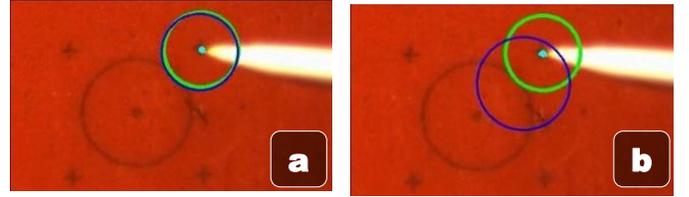
Fig. 6. Micron enforcing a point hard virtual fixture with visual cues: green represents the goal, blue represents the null position. (a) Nearly coincident circles is desirable, indicating Micron is near the center of the range of motion. (b) The position and size of the blue circle indicates drift. Here, the operator has drifted left, down, and upwards in the X, Y, and Z directions, respectively. If Micron was turned off, the tip would snap to the blue circle.

I. Visual Cues to Maintain Eye-Hand Coordination

A final consideration in our formulation of position-based virtual fixtures is one of practical implementation. The input to virtual fixtures is the null position, which is measured by the ASAP sensors but is otherwise unseen to the operator. Thus, only the tip position can be used by the operator to guide the tool. Hard virtual fixtures fix the tip to the virtual fixture regardless of any hand movement, removing all error and effectively disrupting the eye-hand coordination feedback loop. With the eye-hand feedback loop broken, the operator will unknowingly drift away from the virtual fixture until the virtual fixture is no longer within Micron's range of motion, at which point active control becomes impossible. Thus it is imperative to provide the operator with a surrogate sense of error to restore eye-hand coordination.

In order to prevent this unbounded drifting behavior, we display visual cues that indicate the 3D location of the unseen null position. As shown in Fig. 6, we choose visual cues in the form of two circles: a green one to show the goal location and a blue one to show the null position, which reflects the actual hand movements. The distance between the circle centers represents the error Micron is currently eliminating. Z error is displayed by varying the radius of the null position circle, e.g. a growing radius represents upwards drift. The operator is instructed to keep the two circles roughly coincident to prevent drifting and saturation of the motors. In all tests, the operator used these visual cues to maintain positioning (the learning curve was found to be minimal).

IV. EVALUATION

Several experiments were performed with Micron to validate the proposed virtual fixture framework. Two scenarios were investigated: simple, synthetic tasks on a rubber slide (similar to [2]), and a more complex vein-tracing task in an eyeball phantom. As shown in Fig. 7, the rubber slide had a laser-etched target consisting of four crosses arranged as the corners of a 600 μm square, with a circle of 500 μm diameter in the middle. To simulate a more realistic task in a more relevant surgical environment, an eye phantom with a curved vein was constructed on the bottom half of a 42-mm diameter sphere (see Fig. 8). The curved "vein" was a red painted hair taped on a yellow paper background that was firmly glued to the inside of a ping-pong ball. Because of the differences in scale, 29X magnification was used for the rubber slide experiments, and 11X magnification was used for
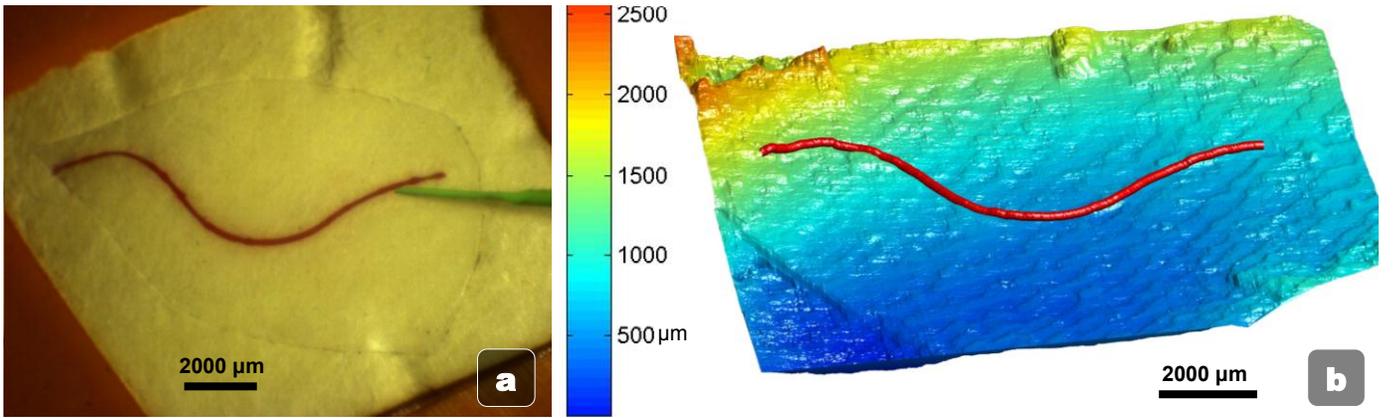
Fig. 8. (a) Half eyeball phantom with a dyed hair tapped to yellow paper affixed to the inside of a ping pong ball with green-painted tip of Micron visible. (b) 3D reconstruction and 3D spline representation of the vein, with dark blue representing low Z positions and dark orange representing high Z positions. Although the vein lies on the surface, it is visualized with red.

the eyeball phantom.

A. Experimental Procedure

Four sets of experiments were run for each task, the first three being synthetic tasks on the rubber slide similar to [2] and the last mimicking a more realistic surgical task of cannulating a blood vessel.

- **Hold Still**: Hold the tip of the instrument 500 μm above a fiducial on a plane for 30 s. A point fixture located above the top right cross was used.
- **Circle Tracing**: Trace a circle with a 500 μm offset from rubber surface two times. A 3D circle fixture derived from the tracked rubber target was used.
- **Move and Hold**: Move between the four corners of a square sequentially, pausing at each corner. Four plane segments oriented vertically and connecting each of the four corners were used, forming a compound box fixture. Because of the vertical orientation of the planes, the Z-movement is unconstrained, allowing the operator to choose the vertical distance from the surface.
- **Vein Tracing**: Follow a 10 mm segment of a curved vein on the inside of a sphere, maintaining a 500 μm vertical distance from the spherical surface. The vein was represented with a B-spline in 3D space.

Experiments were performed in random order to alleviate ordering effects. Visual cues were displayed on a 3D monitor for the synthetic tasks and in the microscope via an image injection system for the vein tracing task. Four different scenarios were tested:

- **Unaided**: Micron was turned off.
- **Aided with Shelving Filter**: Micron was turned on with the tremor suppressing shelving filter from [2].
- **Aided with Soft Virtual Fixtures**: Micron was turned on and using virtual fixtures with the motion scaling factor $\lambda = 1/5$, so errors were reduced by 5X.
- **Aided with Hard Virtual Fixtures**: Micron was turned on with virtual fixtures but no motion scaling ($\lambda = 0$).

B. Generating Fixtures from Vision

The discussion of virtual fixtures has assumed the availability of each fixture. In reality, virtual fixtures are generated from anatomical structures perceived by the vision system. In addition to generating fixtures, the vision system is also responsible for maintaining the adaptive registration by tracking the tip position. For ease of tracking, the Micron nitinol tip is painted with acrylic paint. Simple image processing techniques including thresholding and blob tracking allow the vision system to accurately locate the instrument tip. The OpenCV library [33] is used to generate virtual fixtures from visual information. While the tracking and virtual fixture formulation presented here could handle moving fixtures, this paper does not explicitly explore this aspect.

1) Rubber Slide Virtual Fixtures

Stereo cameras attached to the microscope operate at 54 Hz with a custom resolution of 504x324. Hierarchical template matching is used to locate the target. The template matching method is enhanced to handle occlusions of the target caused by the tip of the instrument and its shadow, yielding very fast and robust tracking. Once the target is located in both images, the center of the target can be backprojected into 3D using the camera registration. Individual virtual fixtures (point, circle, box) are derived from the 3D location of the target and the relative distances of the components. All tracking and virtual fixture generation happens at 50 Hz.

2) Eye-Phantom Vein Virtual Fixtures

Although the same tip tracker is used to localize the instrument and maintain adaptive registration, the vision system for the eyeball phantom was expanded to handle the

TABLE I
MAX ERROR IN POSITIONING TASKS

| Task | Unaided | Aided | | |
	Micron Off (μm)	Shelving Filter (μm)	Soft Fixtures (μm)	Hard Fixtures (μm)
Hold Still	**480±214**	**261±99**	99± 40	87±41
Circle Tracing	**710±500**	**711±262**	302±260	105±72
Move/Hold	**326±501**	**305±167**	132±96	126±43
Vein Tracing	**701±177**	492±121	415±226	383±177

Mean of the 3D max error across all seven trials of each combination of task and scenario. Significantly reduced results compared to the hard fixtures case are bolded (p < 0.05). Because vein tracing experiments used lower magnification and larger motions were necessary (10 mm), saturation was encountered more frequently than in other tasks.

Fig. 9. Mean 3D RMS error across seven trials of each combination of task and scenario, with standard deviation error bars. Hard fixtures significantly reduce error compared to unaided and shelving filter scenarios ($p < 0.05$).

Fig. 10. Mean 2D RMS error across seven trials of each combination of task and scenario, with standard deviation. Hard fixtures significantly reduce error compared to unaided and shelving filter scenarios ($p < 0.05$).

more complex geometry of the experiment. Higher-resolution 1024x768 video is analyzed at 30 Hz. Lifting the planar assumption and moving into an environment with curved surfaces necessitates a good 3D understanding of the environment. We use stereo vision to densely reconstruct the surface of the eyeball phantom retina using the fast Semi-Global Block Matching (SGBM) algorithm [34]. Although a quadratic or spherical fit could be attempted, we leave the disparity map in a raw pixel-wise representation. The vein is localized with color trackers, filtered with blob analysis, and skeletonized with morphological and distance transforms. The vein is then fit to a B-spline in the XY image space and XZ direction (based on the disparity map). Corresponding points on the spline in the left and right images are backprojected to form a full 3D representation of the vein on the surface (see Fig. 8).

The anatomy of greatest interest is often directly underneath

of the tip of the instrument and occluded from view. Because the tip occludes both the vein when fitting the spline and the surface when calculating the disparity map, the initial unoccluded spline representation is stored. When occlusion of the vein is detected, Iterative Closest Point (ICP) [35] is used to register the occluded vein with the initial vein and fill in occluded parts. The GNU Scientific Library (GSL) library [36] is used to then fit the splines. Occlusion of the disparity map is handled by linearly interpolating over occluded patches. Full 1024x768 stereo depth maps and 3D spline representations are updated at 2 Hz and provided to the controller which runs at 2 kHz.

C. Results

Each set of synthetic experiments on the rubber slide was performed by three surgeons at least once, for a total of seven sets. The vein tracing experiment in the eyeball phantom was performed by a single vitreoretinal surgeon for a total of seven

Fig. 11. Hold Still results (a) unaided (b) aided with shelving filter (c) aided with soft fixtures (d) aided with hard fixtures.
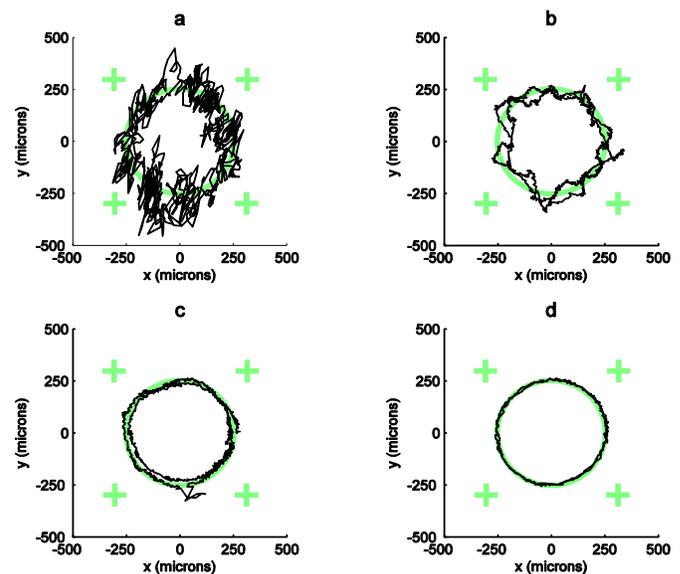
Fig. 12. Circle Tracing results (a) unaided (b) aided with shelving filter (c) aided with soft fixtures (d) aided with hard fixtures.
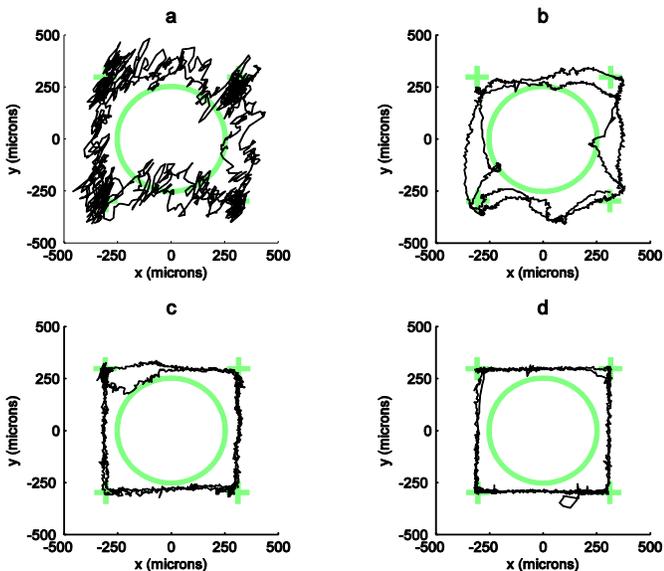
Fig. 13. Move and Hold results (a) unaided (b) aided with shelving filter (c) aided with soft fixtures (d) aided with hard fixtures. Deviations from the virtual fixtures in (c) and (d) result from saturation of the actuators, caused by tremor or drift in excess of the range of motion of Micron.

Fig. 14. Vein Tracing results (a) unaided (b) aided with shelving filter (c) aided with soft fixtures (d) aided with hard fixtures.

sets. All error was measured as the Euclidean distance between the tip position sensed by the ASAP optical trackers and the closest point on the virtual fixture as generated by the registered stereo cameras. Error was recorded at 2 kHz.

Figs. 9-10 display 3D and 2D mean RMS error with standard deviation across all three trials for each combination of task and scenario. It is important to note that the Move and Hold scenario has the same 3D and 2D error because the vertical virtual fixtures do not restrict Z-movement. Additionally, the Vein Tracing task has a higher incident of saturation and thus more error, which is partially because of the lower magnification and the greater tip traversal required (almost 10 mm, compared to the 0.5-mm circle). In fact, 98% of the error during the Vein Tracing task with hard virtual fixtures occurs while the manipulator is in saturation.

Maximum 3D error for each set of trials is presented in Table I. In all cases, mean and maximum error for hard virtual fixtures is significantly less than in the unaided case for both 2D and 3D measurements ($p < 0.05$) as calculated by a two-tailed t-test. Mean error of hard virtual fixtures for both 2D and 3D is significantly less than the state-of-the-art shelving filter from [2] ($p < 0.05$). Figs. 11-14 show one trial of each task and scenario with a trace of the 2D tip position overlaid in black on the target represented by thick light green lines.

V. DISCUSSION

In this paper, we have presented a derivation of virtual fixtures that depends not on forces applied to a robot arm, but on motion of the handle of the instrument. This new position-based virtual fixture formulation is necessary for the class of handheld micromanipulators such as Micron that use handle positions as input to the control system. Virtual fixtures are generated in real time from stereo cameras attached to the microscope, providing task-dependent behaviors to the operator. Visual cues are displayed to the operator to maintain eye-hand coordination and help guide the operator during the procedure.
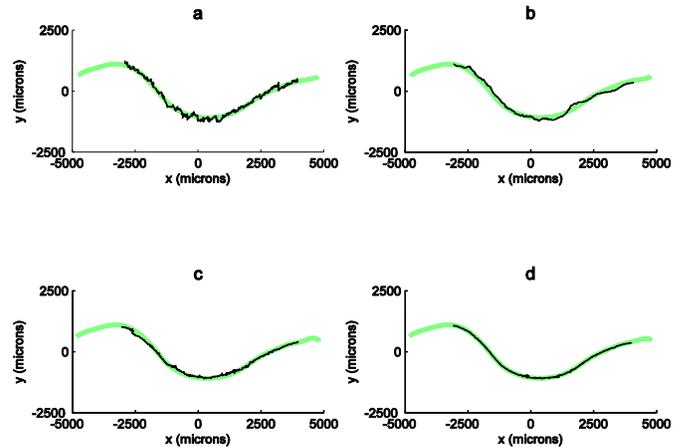
Hard virtual fixtures totally constrain the tip, enforcing constraints that should never be violated, e.g. forbidden areas or snap-to behaviors. Implementation of soft virtual fixtures under our formulation directly results in intuitive motion scaling. In all cases, tremor suppression filters compensate for unwanted high frequency motion. Furthermore, the virtual fixture framework easily adapts to general parameterizations such as splines to model complex anatomy. Using Micron as a test platform, virtual fixtures have been validated with medically relevant artificial tests such as vein tracing. Compared to unaided performance and state-of-the-art tremor compensation [2], virtual fixtures significantly reduce tip positioning error ($p < 0.05$).

Future work includes extensions from position-based virtual fixtures to orientation-based virtual fixtures for situations where the tip position is less important than where the tip is pointing. Such orientation fixtures would be useful to enforce a remote center of motion or in performing laser therapies, such as photocoagulation or ablation techniques. Another area of research is the application of the virtual fixtures to medical procedures, such as vitreoretinal microsurgery where extremely high precision is required within ranges of motion attainable by Micron. A new Micron prototype with a larger range of motion, presently under development, should reduce incidence of saturation and improve virtual fixture performance. Further testing in more realistic situations *ex vivo* and *in vivo* is necessary in order to fully validate the efficacy of the proposed work in realistic surgical experiments.

REFERENCES

[1] S. P. N. Singh and C. N. Riviere, "Physiological tremor amplitude during retinal microsurgery," in *Proc. IEEE 28th Annu. Northeast Bioengineering Conf.*, 2002, pp. 171-172.

[2] R. A. MacLachlan, B. C. Becker, J. C. Tabares, G. W. Podnar, L. A. Lobes, Jr., and C. N. Riviere, "Micron: An actively stabilized handheld tool for microsurgery," *IEEE Trans. Rob.,* vol. 28, pp. 195-212, 2012.

[3] H. Das, H. Zak, J. Johnson, J. Crouch, and D. Frambach, "Evaluation of a telerobotic system to assist surgeons in microsurgery," *Computer Aided Surgery,* vol. 4, pp. 15-25, 1999.

[4] A. Uneri, M. A. Balicki, J. Handa, P. Gehlbach, R. H. Taylor, and I. Iordachita, "New steady-hand eye robot with micro-force sensing for vitreoretinal surgery," in *Proc. Bio. Rob. BioMech.*, 2010, pp. 814-819.

[5] T. Ueta, Y. Yamaguchi, Y. Shirakawa, T. Nakano, R. Ideta, Y. Noda, A. Morita, R. Mochizuki, N. Sugita, and M. Mitsuishi, "Robot-assisted vitreoretinal surgery: Development of a prototype and feasibility studies in an animal model," *Ophthalmology,* vol. 116, pp. 1538-1543, 2009.

[6] B. Mitchell, J. Koo, M. Iordachita, P. Kazanzides, A. Kapoor, J. Handa, G. Hager, and R. Taylor, "Development and application of a new steady-hand manipulator for retinal surgery," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2007, pp. 623-629.

[7] C. N. Riviere, W. T. Ang, and P. K. Khosla, "Toward active tremor canceling in handheld microsurgical instruments," *IEEE Trans. Robot. Autom.,* vol. 19, pp. 793-800, 2003.

[8] Y. Ida, N. Sugita, T. Ueta, Y. Tamaki, K. Tanimoto, and M. Mitsuishi, "Microsurgical robotic system for vitreoretinal surgery," *Intl. J. Comp. Asst. Rad. Surg.,* pp. 1-8, 2011.

[9] S. E. Everett, Y. Isoda, R. V. Dubey, and C. Dumont, "Vision-based end-effector alignment assistance for teleoperation," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1999, pp. 543-549 vol. 1.

[10] A. Bettini, P. Marayong, S. Lang, A. M. Okamura, and G. D. Hager, "Vision assisted control for manipulation using virtual fixtures," *IEEE Trans. Robot. Autom.,* vol. 20, pp. 953-966, 2004.

[11] L. B. Rosenberg, "Virtual fixtures: Perceptual tools for telerobotic manipulation," in *IEEE Virt. Reality Ann. Intl. Symp.*, 1993, pp. 76-82.

[12] H. C. Lin, K. Mills, P. Kazanzides, G. D. Hager, P. Marayong, A. M. Okamura, and R. Karam, "Portability and applicability of virtual fixtures across medical and manufacturing tasks," in *Proc. IEEE Int. Conf. Robot. Autom*, 2006, pp. 225–231.

[13] M. Li, M. Ishii, and R. H. Taylor, "Spatial motion constraints using virtual fixtures generated by anatomy," *IEEE Trans. Robot.,* vol. 23, pp. 4-19, 2007.

[14] B. C. Becker, S. Voros, L. A. Lobes, Jr., J. T. Handa, G. D. Hager, and C. N. Riviere, "Retinal vessel cannulation with an image-guided handheld robot," in *Conf. Proc. IEEE Eng. Med. Biol. Soc.*, 2010, pp. 5420-5423.

[15] B. C. Becker, R. A. MacLachlan, G. D. Hager, and C. N. Riviere, "Handheld micromanipulation with vision-based virtual fixtures," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2011, pp. 4127-4132.

[16] B. C. Becker, R. A. MacLachlan, and C. N. Riviere, "State estimation and feedforward tremor suppression for a handheld micromanipulator with a Kalman filter," in *Proc. IEEE Intl. Conf. Intelligent Rob. Sys.*, 2011, pp. 5160-5165.

[17] J. Funda, R. H. Taylor, B. Eldridge, S. Gomory, and K. G. Gruben, "Constrained cartesian motion control for teleoperated surgical robots," *IEEE Trans. Robot. Autom.,* vol. 12, pp. 453-465, 1996.

[18] A. Kapoor, M. Li, and R. H. Taylor, "Constrained control for surgical assistant robots," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2006, pp. 231–236.

[19] B. L. Davies, S. J. Harris, W. J. Lin, R. D. Hibberd, R. Middleton, and J. C. Cobb, "Active compliance in robotic surgery—the use of force control as a dynamic constraint," *Proc. Inst. Mech. Eng., Part H: J. Eng. Med.,* vol. 211, pp. 285-292, 1997.

[20] R. Kumar, G. Hager, A. Barnes, P. Jensen, and R. Taylor, "An augmentation system for fine manipulation," in *Proc. Med. Image Comput. Comput. Assist. Interv.*, 2000.

[21] M. Dewan, P. Marayong, A. Okamura, and G. Hager, "Vision-based assistance for ophthalmic micro-surgery," *Proc. Med. Img. Comp. Comp. Astd. Intvn.,* pp. 49-57, 2004.

[22] M. Li, A. Kapoor, and R. Taylor, "Telerobotic control by virtual fixtures for surgical applications," *Advances in Telerobotics,* vol. 31, pp. 381-401, 2007.

[23] R. A. MacLachlan and C. N. Riviere, "High-speed microscale optical tracking using digital frequency-domain multiplexing," *IEEE Trans. Instrum. Meas.,* vol. 58, pp. 1991-2001, 2009.

[24] S. Hutchinson, G. D. Hager, and P. I. Corke, "A tutorial on visual servo control," *IEEE Trans. Rob. Autom.,* vol. 12, pp. 651-670, 1996.

[25] F. Chaumette and S. Hutchinson, "Visual servo control, part I: Basic approaches," *IEEE Rob. Autom. Mag.,* vol. 13, pp. 82-90, 2006.

[26] B. Becker, S. Voros, R. MacLachlan, G. Hager, and C. Riviere, "Active guidance of a handheld micromanipulator using visual servoing," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2009, pp. 339-344.

[27] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge, UK: Cambridge University Press, 2003.

[28] J. N. Weiss and L. A. Bynoe, "Injection of tissue plasminogen activator into a branch retinal vein in eyes with central retinal vein occlusion," *Ophthalmology,* vol. 108, pp. 2249-2257, 2001.

[29] L. A. Bynoe, R. K. Hutchins, H. S. Lazarus, and M. A. Friedberg, "Retinal endovascular surgery for central retinal vein occlusion: initial experience of four surgeons," *Retina,* vol. 25, pp. 625-32, 2005.

[30] Y. L. Ma and W. T. Hewitt, "Point inversion and projection for nurbs curve and surface: Control polygon approach," *Computer Aided Geometric Design,* vol. 20, pp. 79-99, 2003.

[31] X. Chen, J. Yong, G. Wang, J. Paul, and G. Xu, "Computing the minimum distance between a point and a nurbs curve," *Computer-Aided Design,* vol. 40, pp. 1051-1054, 2008.

[32] R. Richa, P. Poignet, and Chao Liu, "Three-dimensional motion tracking for beating heart surgery using a thin-plate spline deformable model," *Intl. J. Rob. Research,* vol. 29, pp. 218-230, 2010.

[33] *OpenCV Library*. Available: http://opencv.willowgarage.com/

[34] H. Hirschmuller, "Stereo processing by semiglobal matching and mutual information," *IEEE Trans. Pattern Anal. Mach. Intell.,* vol. 30, pp. 328-341, 2008.

[35] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Trans. Pattern Anal. Mach. Intell.,* vol. 14, pp. 239-256, 1992.

[36] *GNU Scientific Library*. Available: http://www.gnu.org/s/gsl/